INTEGRATOR'S JOURNAL

by Wayne M. Krakau - *Chicago Computer Guide*, March 1992

Wayne Krakau welcomes all calls, comments, and questions from basement businesses to corporate MIS. Wayne is a partner in Krakau Business Computer Systems, a systems integration firm and a Novell Gold authorized and Lantastic 5 Star Dealer. He has been working with computers for 15 years. He has an MBA in Marketing and a BS in Information Science and was originally educated in Computer and Information Systems Analysis Engineering. He holds Novell CNE and CNI ratings, and LANDA Certifications in LAN Installation, Architecture, Connectivity, Maintenance, and Program Development. Wayne can be reached for questions, comments, or topic suggestions at (708)298-7695 or through the Computer Guide.

This is a new column that will provide information on the products, people, news, and situations, encountered in the life of a systems integrator (yours truly). Since a large part of my work involves "rescues" of people after they've been victims of botched integration attempts or abuse or abandonment by vendors, the slant of this commentary may sometimes be a bit cynical. As an experienced street cop might tell you, that's a hazard of constant exposure to the seamier side of the beat.

Due to popular demand (via faxes to the Computer Guide -- Yes, we do read them!) LANs (Local Area Networks) and other connectivity issues will be prominently, though not exclusively, featured.

Upon reading this column, my dedication to defining terminology will soon become apparent. I would rather err on the side of over-explaining and have technically oriented readers skip over the explanations than leave some people in the dark. Even techies don't know every variation of every possible phrase (unless they cheat like I do and look them up).

I can still remember being baffled by the constant use of "BTW" in messages on BBS's (electronic Bulletin Board Systems). Eventually, I figured out that it was short for "By The Way". It was even worse for "OEM", an acronym constantly used in the computer press. It took an embarrassingly long time to find out that it means "Original Equipment Manufacturer". To add to my confusion, it's also used as a verb to describe two opposite sides of a transaction. Sometimes it refers to manufacturing and selling a product for another company's eventual resale under its name. Other times I've seen it used to describe the purchase side of that same transaction! I'd rather define terms up front to avoid this type of problem.

For my first news item, in the category of "You Can't Teach An Old Dog New Tricks", Novell has finally confirmed the rumors about the future (or lack thereof) of Netware 2.2. With the allegedly "brain-dead" Intel 80286 chip taking the part of the old dog and Novell taking the part of the unwilling teacher, Novell announced that there would be no further development or enhancements for version Netware 2.2. Maintenance and support will continue (for now), but the product is now officially a dead-end. (Mechanical life-support will continue for the comatose patient, but his records will be tagged "DNR" -- Do Not Resuscitate.)

A 10-user ($2495) version of Netware 3.11 was released last fall and a 5-user ($1095) version is expected soon. Prices on the 50- and 100-user versions of 2.2 are being raised by $500 (That's the stick). Special upgrade prices from 2.2 and below to 3.11 are expected (And there's the carrot). Those considering a purchase of or upgrade to 2.2 should stop and estimate the future costs involved in buying into an explicitly obsolete product. They also should consider purchasing XT's, AT's, and an occasional 386SX-16 to keep with the "retro" spirit of their system.

Vendors of LAN-related hardware and software have long treated 2.2 as the poor stepchild of the Netware family making it an unattractive purchase option for those who monitor the industry. It's good to see 2.2's fate confirmed officially so everyone will know.

My first cynical observation (or incoherent rave out, depending on your point of view) concerns the use of the term "systems integrator". A common remark in the LAN world is that it's easier to get Netware than it is to get a *Slurpee* since there are more Novell Authorized Resellers than there are *7-11s*. This is not a facetious remark - it's true! After hearing a client remark that there was a systems integrator on almost every street corner, I thought something must be wrong with the definition of systems integrator. Subsequently, I kept an eye out for an unbiased explanation of that term.

I found that explanation in a recent issue of *VARBUSINESS*, a publication for value added resellers (VARs), developers, and integrators. In it, 19% of the respondents to their "State of the Market" survey called themselves systems integrators while less than 1% actually qualified according to the magazine's criteria!

To use the term "systems integrator" the magazine required that at least 25% of the revenue of a company had to be from project management or systems integration. In their (somewhat pompous) definition, systems integrators "provide a comprehensive information processing solution through a unique combination of professional services and expertise in hardware, software, and communications technology." They emphasize that systems integrators must maintain full responsibility "for the entire project from design to post-implementation stages, including product selection."

That last sentence is the key. Computer systems, and in particular, LANs, are almost by definition, multi-vendor solutions.

For example, even IBM doesn't try to do everything anymore. They are currently abandoning their main applications programming effort, so applications software for their machines will have to be obtained from other vendors. Their main PC operating system (for now, at least), PC-DOS Version 5.0, was written by Microsoft with only minor tweaks by IBM. (Thank God! -- Their previous attempts at significantly modifying MS-DOS were versions 3.2 and 4.0, both unmitigated disasters.) That's another vendor. Also, most people end up with plenty of non-IBM hardware in or on their IBM computers, thereby adding more vendors.

When you can't get a single brand solution from a monolith like IBM the need for someone to take responsibility for getting all of these products to work together optimally becomes apparent. The government and big business figured that out decades ago, as evidenced by the success of EDS (Electronic Data Systems) and its kin. The appeal of systems integrators seems to be spreading -- why else would so many companies try to use (or misuse) this moniker? Oh, well, the computer industry has never been known for truth in advertising.

As mentioned above, questions, comments, and topic suggestions are always welcome.

BACKUPS, ANYONE?

by Wayne M. Krakau - *Chicago Computer Guide*, April 1992

Wayne Krakau welcomes all calls, comments, and questions from basement businesses to corporate MIS. Wayne is a partner in Krakau Business Computer Systems, a systems integration firm and a Novell Gold authorized and Lantastic 5 Star Dealer. He has been working with computers for 15 years. He has an MBA in Marketing and a BS in Information Science and was originally educated in Computer and Information Systems Analysis Engineering. He holds Novell CNE and CNI ratings, and LANDA Certifications in LAN Installation, Architecture, Connectivity, Maintenance, and Program Development. Wayne can be reached for questions, comments, or topic suggestions at (708)298-7695 or through the Computer Guide.

This is getting ridiculous! I'm now getting at least one new client per month who is backing up its data in what turns out to be an invalid or inconsistent manner. A false sense of security keeps these system managers (if you run a LAN or even a single computer, you're THE System Manager even if you're a "civilian" and don't use the title) thinking that they are doing their duty and safeguarding their computer systems (mostly LANs). In reality, they were really running either partially or fully unprotected.

One client had been using the same five 60MB (Note: MB=Megabyte, Mb=megabit) cassette-style tapes since 1986 to back up a 102MB network hard disk that held 97 MB! This is the proverbial ten pounds of potatoes in a five pound sack.

The dealer who sold them the system thoughtfully automated the backup process with a batch file. He (I know it's a guy - I asked) was so considerate, in fact, that he disabled the ability to automatically request continuation onto a second tape when the first was full. None of that time wasting stuff for his customers. No more of those annoying, error messages, either -- he suppressed displaying them. And knowledge of the error log file really should be on a need to know basis, so don't tell anyone about it, and certainly don't do anything radical like printing it out at the end of the backup. As for the verify (also called compare) function, well, that's like having a belt and suspenders. Let's skip it. Finally, the most confidential aspect of a backup system is the backup software itself. This, of course, means that you shouldn't reveal where it is hidden, and you absolutely can't divulge the ultimate of secrets -- how to execute it manually in an emergency.

While we're on a roll, let's also skip the lesson on multiple redundant backup tapes. Tapes won't ever wear out or get damaged. Besides, what's the few cents profit per additional tape sold versus the potential income in reconstructing a blown database. A side benefit is that you can also underbid the competition by including just a few tapes in your bid. And, again, keeping profit in mind, don't warn the customer about the effects of smoke and dust around a tape drive. Just wait for those lucrative repair calls. No lesson on retensioning or reformatting either -- time-wasters all!

The interesting thing about this is that the clients didn't even know that they were in trouble. They surely didn't know that critical portions of their main database, the one that the business depends on, were not included on their backups! They religiously followed the directions that they had, so no one could say it was their fault.

Through sheer luck, no catastrophe ever required the use of their backup tapes. Those tapes were completely worn out and could barely sustain the data that did fit on them. I outfitted them with a 250MB (to allow for planned expansion) DC2000 style tape drive (a Colorado Memory Systems Jumbo 250) and a controller with hardware compression (Colorado's TC-15). With a little added RAM (random access memory) and a memory management program (QuarterDeck's QEMM), this system could do a complete (really complete, this time) backup and verify in less time than the old tape drive could do its aborted backup without verifying.

This client now uses 18 tapes. The time it took to explain to him why the tapes were needed (and to get his eyes to pop back into his head) ate up more money than the minuscule profit on those tapes. He agreed with the backup/disaster recovery plan that I proposed. Actually, it's a modified version of a quite common one, devised on mainframe systems many years ago when the big iron folks (yes, I was one) first started to look upon computing as a potentially disciplined science, and not as an art (well, maybe a craft).

In this clients new system, there were three main sets of tapes, each consisting of a Monday, Tuesday, Wednesday, and Thursday modified backup (only new or changed files are backed up), and a Friday full backup. The active set is kept in a hallway adjacent to the room where the computer with the tape drive resided. The secondary set (from the previous week) is kept in the opposite side of the building, a considerable distance and several walls away from the first set. The final set (from two weeks ago) is taken offsite. In this case an actual vault would have been overkill, so the system manager just takes the tapes home. The other group of three tapes is a sort of fiscal monthly set. At the end of each three-week cycle (this client's fiscal month for backup purposes), a full backup is made and immediately taken offsite.

I have had several other clients running into the capacity problem (too much data for their backup system), one with a Bernoulli Box (from Iomega) on a single-user system, and the rest with tape drives on LANs. In each case, the system manager (again, all civilians) didn't realize that only part of the data was being saved.

By far the most common problem that I see is the use of too few tapes, or a poorly planned pattern of use. While not everyone needs to be as thorough as the client that I mentioned, the use of one, three, or even five tapes just doesn't make a system safe. It's very common to discover an error in a major database, examine your backups, and find that the problem actually occurred some time ago.

Another client of mine had problems with their main (and vital) database. The error that caused the data corruption had occurred three weeks prior to their noticing the problem. Sadly, they thought that it was better to make a full backup every night and reuse the same six (they worked on Saturdays) tapes every week than to follow the agreed upon backup plan. The "extra" tapes were held in reserve since their tapes had an unusually high failure rate (later determined to be due to damage caused by smoke - they eventually declared a nonsmoking policy because they got tired of paying repair and replacement bills on computers -- final tally: two keyboards, a power supply, and a motherboard dead, plus numerous repair calls, all from obvious smoking related problems -- they haven't made a hardware service call since the policy change). They had been fed bad information from the corporate MIS (that's Management Information Systems) department, but I still feel partially responsible since it was my duty to impress upon them the importance of sticking to the backup plan.

With no regular backup from an appropriate period, we had to go back to the last yearly tape (this was in March). We were lucky, since that tape had been put in the "extra" bin but had not yet been reused. The client had to have some of the data reconstructed by the software vendor (they used a vertical market package from a very small vendor who used a proprietary file format that they would not reveal -- if you wanted to reconstruct data, export it to a spreadsheet for further analysis, or even needed a simple report, you had to purchase their custom programming services -- the vendor was expensive, slow, and not particularly competent). The rest of the data had to be manually reentered from the original input forms and some transaction reports. This client immediately reinstituted the original backup procedures with some enhancements to make it even more thorough. It only takes one incident like that to "get religion".

NOVELL NEWS: Due to an outcry from existing Netware 2.2 users, Novell has partially recanted their statements regarding the impending obsolescence of Netware 2.2. They now say that it will be further enhanced with management and other features, but are still non-committal about a possible 2.3 version. It sounds like there will be a few improvements, but not enough to warrant a major version number alteration. Since most vendors providing Netware enhancements are still treating 2.2 as a second-banana at best (as per last month's column), my advice is still to stick with (or move to) Netware 3.11. It is the future.

As always, questions, comments, and topic suggestions are always welcome.

THE COMDEX DIET PLAN

by Wayne M. Krakau - *Chicago Computer Guide*, May 1992

Wayne Krakau welcomes all calls, comments, and questions from basement businesses to corporate MIS. Wayne is a partner in Krakau Business Computer Systems, a systems integration firm and a Novell Gold authorized and Lantastic 5 Star Dealer. He has been working with computers for 15 years. He has an MBA in Marketing and a BS in Information Science and was originally educated in Computer and Information Systems Analysis Engineering. He holds Novell CNE and CNI ratings, and LANDA Certifications in LAN Installation, Architecture, Connectivity, Maintenance, and Program Development. Wayne can be reached for questions, comments, or topic suggestions at (708)298-7695 or through the Computer Guide.

Operators are standing by! Have your credit card ready. We accept Visa, MasterCard, Discover, and American Express. In addition, we even accept Amoco, Sears, Field's, and Carson's cards (we are squeegee equipped). Yes, it's just what you've been waiting for -- The Comdex Diet Plan.

Forget aerobics classes. We'll teach you a modified race-walking step with an added lesson on broken field running. This is perfect for getting past booths crowded with greedy freebies seekers, ready to kill for a 50-cent pair of sunglasses, a 10-cent luggage tag, or that most valued of items -- a free issue of an obscure 2-pound technical journal covering a field in which you have no involvement (or interest).

Forget pumping iron. We use the latest in graduated weight training techniques. We'll show you how to start with two simple plastic bags (furnished with sharp-edged handles to toughen those wimpy office-worker hands) and bit-by-bit fill them to Schwarzeneggerian poundage while keeping your tilt to less than five degrees.

We'll even throw in a guide to warmups and increased flexibility including instructions on removing and replacing your show badge with your teeth (leaving the hands free to gather more literature).

We have two versions of the Plan. The Four Day Plan is for sprinters who desire breaks between exercises to see entertainment and demonstration sessions complete with magicians, comedians, singers, and dancers. (Warning: This area has been declared a Mime Free Zone!) The Three Day Plan is for marathoners who need to go the distance in as short a time as possible. So call now, and tell the operator which version you want.

Yes, I inadvertently tried the Three Day version of the Comdex (originally the Computer Dealer Expo) Diet Plan. After nineteen and a half hours of walking the corridors of McCormick Place (combined show hours for Monday, Tuesday, and Thursday, not including any pre-show walking) dragging more literature than I care to think about I really did lose six pounds!

I also obtained so much valuable information that will take months to fully digest it. For example, I have been searching (mostly in vain) for over five years for a software that could handle document and image management. I recently found a package that handles computer generated (normally with WordPerfect) documents quite well, PC DOCS (Tallahassee, Florida), but the selection in image-only management was sparse and products that handle both were quite rare and very limited in scope and useability.

At Comdex, I found LaserFiche from Compulink Management Center, Inc. (Torrance, California). It can organize pure images, transform images into documents via OCR (optical character recognition), logically linking the image with its corresponding document, and manage word-processor generated documents as well. They're even creating a Windows version to keep up with the latest trend in imaging and graphics. This is easily my pick for the highlight of the show.

Most of the other new and interesting products were in areas with more competition. Pen computing is an obvious example. Many firms demonstrated products with mostly similar features. I doubt that their enthusiasm will turn into sales for mainstream applications. While pen-based computing may be appropriate for certain niche markets (many examples of which were being demonstrated), none of the vendors present were able to provide convincing arguments to defend their claims of future near-universal acceptance of the pen as the perfect input device.

Random observation: How much money is being spent on screen-blankers for windows? I realize that flying toasters and the like are clever, but don't these people know how hard it is to burn out the phosphors (ostensibly what this software is designed to prevent -- they don't really push the security issue) on a color monitor, especially using a GUI (graphical user interface)?

Another random observation: I was telling a colleague that I had been pleased to discover only one booth with scantily clad models. We discussed the potential of such a display in directly offending women and indirectly offending men (by the assumption that they would fall for such a stunt). Just as I waxed poetically about the possibility that this industry may be maturing, the booth in question was over-run by hordes of show attendees (males, of course) jostling for position while attempting to get literature packets about an obscure (and somewhat lame) product from a tiny, even more obscure company. Our discussion ended at that point.

As to the Windows (Microsoft) versus OS/2 (IBM) battle, there was a split decision, but as with most professional fights, the champion cannot lose his title on a split decision. Microsoft gained mind-share through their brilliantly executed Windows World strategy. They showed their marketing savvy and, more importantly, the overall practicality of the Windows 3.1 (and future Windows NT) environment. IBM got points for pure technical excellence and for its improved interface (it out-MACs the MAC), but I doubt that most people were persuaded to abandon what is becoming an industry standard to accept OS/2 as a general purpose environment. IBM will certainly garner market share in specialty niches where OS/2's capabilities are absolutely required, but it doesn't look like it's ready to derail Bill Gates' favorite train just yet.

Speaking of Bill Gates, I found him lurking on the sidelines of a WordPerfect for Windows demonstration accompanied by two stiff looking gents in very severe suits. There was some speculation amongst onlookers as to their identity, with opinions varying from Microsoft vice-presidents to personal bodyguards. Considering that we were observing, essentially, more than six billion dollars on-the-hoof, either theory could have been true.

I noticed that Gates was somewhat shorter and a lot scruffier (Bill, you can afford to hire someone to carry your comb!) than I had imagined. He was dressed informally, wearing a plain blue cardigan -- aha! -- BLUE! -- Perchance a secret attempt at diplomacy with IBM (Big Blue)? -- Naaaah! It's probably the first ratty sweater that he pulled out of his closet while packing.

I did miss a chance at a Mike Wallace merit badge by not interrogating Gates on the spot. I guess I don't have the killer instinct for in-your-face-journalism. Oh, well -- at least I've kept my day job.

Questions comments, topic suggestions, and complaints from irate billionaire/geniuses are always welcome.

THE BROTHER-IN-LAW SYNDROME

by Wayne M. Krakau - *Chicago Computer Guide*, June 1992 - Newsware, August/September 1995

Wayne Krakau welcomes all calls, comments, and questions from basement businesses to corporate MIS. Wayne is a partner in Krakau Business Computer Systems, a systems integration firm and a Novell Gold authorized and Lantastic 5 Star Dealer. He has been working with computers for 15 years. He has an MBA in Marketing and a BS in Information Science and was originally educated in Computer and Information Systems Analysis Engineering. He holds Novell CNE and CNI ratings, and LANDA Certifications in LAN Installation, Architecture, Connectivity, Maintenance, and Program Development. Wayne can be reached for questions, comments, or topic suggestions at (708)298-7695 or through the Computer Guide.

It's becoming an epidemic. I've found yet another victim of the dreaded Brother-In-Law Syndrome. In this condition, someone with little or no background in an appropriate field, most frequently a brother-in-law, and almost invariably a male, provides computer advice.

I'm sure the government is commissioning a study as to whether this heavily skewed gender proportion is a genetic or a sociological phenomenon. Could it be the gene that causes men to freely dispense directions when they aren't familiar with the area? Or, maybe it's the same gene that makes guys with inadequate automotive aptitude freely disperse advice about a disabled car. (Oops, I've done that!)

I first discovered a variation of this syndrome while working on mainframes. Corporate officers would be approached by manufacturers' sales people. By giving a dazzling technical presentation to a non-technical audience, the sales people could avoid serious questions about the validity of their facts. They could freely use what's known as the FUD (fear, uncertainty, and doubt) factor to threaten catastrophic results upon the purchase of competitive products. The corporate officer would then override the results of expensive internal feasibility studies to purchase the product in question. The technical staff would be stuck supporting this product even though it wasn't the appropriate solution.

In the micro arena, the advice often originates, not with a sales pitch, but with a request for advice from someone who "knows" computers. The outcome of ignoring advice from family members (especially the aforementioned brother-in-law) is obvious. The advice is especially compelling since, in most cases, the advisor has no financial interest in the purchase (excluding, of course, those who bill themselves as "consultants"). Considering the lack of journalistic quality control in the microcomputer industry, the advisor can often quote magazine articles praising the suggested product.

The problem lies in accepting advice from a single source without appropriate justification, and, in not properly evaluating the source of the advice. For example, if I advise a client to purchase a product, and the client accepts my recommendation without explanation, I get scared. While it does wonders for my ego, it raises the possibility that the client is deferring to my (hopefully) greater knowledge and experience to such a degree that he or she is actually suppressing questions or filtering out possibly important details about his or her needs. This is the one last chance to catch any discrepancies between my mental picture of the clients requirements and what they really are. My justification of the product choice is a critical part of the analysis and specification process. Advice from a brother-in-law (or his surrogate) removes this critical step.

As to the source of the advice, most alleged consultants in the computer field are actually free-lance programmers. While programming is a perfectly honorable and creative field, having superior programming skills says nothing about your ability in business, product evaluation, hardware, and even software design. And, most of the "brothers-in-law" aren't even in the consulting business. Some are programmers in mainframes and minicomputers. Others are just microcomputer enthusiasts or hobbyists. Some merely use computers in their work.

The most recent occurrence of the Brother-In-Law Syndrome was with a company that purchased a network that included both new and existing machines. A member of the Board of Directors knew someone (maybe a real Brother-in-law?) who bought a particular brand of computer to tinker with at home. Based on that glowing accolade, the board member convinced the other directors to override the analysis of the company's in-house computer department and to purchase a dozen of these computers. It

was justified by stating that it was better to split the order so that the company wouldn't become too dependent on one vendor.

After watching the computer staff suffer with configuring these machines (we had preconfigured one example out of the computers that we sold them so that they could simply duplicate that configuration on the rest of our systems) I finally gave them some free advice (no money had been budgeted, they had bought a considerable amount from my company, and I'm kind of a soft touch). They eventually replaced all of the video cards and got the systems to work, but they never really fully optimized the configuration and they worked a lot of unpaid overtime (salaries only). Another victim of the "Brother-In-Law Syndrome".

For smaller firms without the corporate overhead and lacking a professional computer staff, the dangers are greater. One client of a colleague canceled the purchase of an incredibly sophisticated integrated accounting system when the brother-in-law (a real one) of the owner offered to write a custom application for very little money. His qualifications: he was a microcomputer user at work who dabbled in programming on the side. The result: the company's growth was severely hampered due to the time and money lost in attempting to fix the custom system. They eventually bought packaged software, anyway.

A related disorder illustrated by this example is that of "Programmers Disease". This is an overwhelming desire to solve any computing need by custom programming. The term used in the industry is "reinventing the wheel". A more efficient and effective line of attack would be to investigate possible software solutions in the following order of preference:

1. Single package.
2. Single package with application specific add-on packages.
3. Multiple packages linked with commercial packages.
4. Multiple packages linked with custom programs.
5. Modifiable single package with externally added custom modifications.
6. Modifiable single packages with internally added custom modifications.
7. Complete custom software.

While there may be some dispute about the exact order of this list, I believe that the trend of going to custom software only as a last resort is logical. However, if all you have is a hammer (programming expertise), everything begins to look like a nail (custom programming). Remember that the full custom programming method is the most expensive (potentially hundreds of times more then packages), most risky (you become the test site), and most lucrative (for the programmer).

While we do custom solutions, they are only a last resort after examining other alternatives.

The main advice that I can give to protect you from these maladies is to always get justification for a proposed solution. A qualified individual will not be offended and will, in fact, take reasonable questions as a sign that you are interested in being an active participant in the analysis and design of your computer system. That participation is the key to a successful project.

As always, questions, comments, and topic suggestions are always welcome.

GET PLUGGED IN

by Wayne M. Krakau - *Chicago Computer Guide*, July 1992

Wayne Krakau welcomes all calls, comments, and questions from basement businesses to corporate MIS. Wayne is a partner in Krakau Business Computer Systems, a systems integration firm and a Novell Gold authorized and Lantastic 5 Star Dealer. He has been working with computers for 15 years. He has an MBA in Marketing and a BS in Information Science and was originally educated in Computer and Information Systems Analysis Engineering. He holds Novell CNE and CNI ratings, and LANDA Certifications in LAN Installation, Architecture, Connectivity, Maintenance, and Program Development. Wayne can be reached for questions, comments, or topic suggestions at (708)298-7695 or through the Computer Guide.

It's happened again. I was bailed out of a potentially awkward situation by my favorite and most vital information source, Netwire.

A client called and described a random printing problem with her company's newly upgraded Netware 3.11 network. The first page of a document would be split into two pieces. The first piece consisted of a randomly sized portion from the upper center of the page. It would print on a page by itself, often mixed with gibberish caused by out of sequence printer codes. The rest of the page, minus an occasional word or two, came out as a completely separate page. This was usually perfectly formatted. Unless you noticed that first portion was on a separate page, you might not notice the few missing lines and might attempt to use the resulting printout.

As I worked to isolate and document the problem, concentrating particularly on the Print Server NLM (Netware Loadable Module, a way of implementing features directly within a 3.xx server so they become a part of the operating system, Netware - NLMs are direct descendants of the Value Added Processes or VAPs used in 2.xx versions of Netware), I received a similar complaint from another client.

This gave me the vital clue. This might be an internal problem (within Netware) rather than a blown hardware or software. That night I logged onto Compuserve, a commercial on-line communications network and entered the command "GO NETWIRE". This allowed me to access Netwire, the official electronic support organ of Novell.

I bypassed the on-line discussion area and went directly to the FILES section. Here, I scanned the descriptions of updates, patches, and fixes for the print server software. Within minutes I had located a file that contained updates to PSERVER.NLM (for 3.xx file servers), PSERVER.VAP (for 2.xx file servers), and PSERVER.EXE (for dedicated workstation-based print servers) to fix the exact problem that my clients had experienced. After entering a few simple commands, I had downloaded the files into my computer, ready to be transferred to the appropriate clients.

Since I was accessing Netwire, anyway, I also downloaded a few text files containing methods to fine-tune and/or debug Novell Networks and, even individual machines. One particularly helpful item summarized the latest hints on installing Microsoft Windows in a Netware environment. While parts of it presented old (for me) information about the Netware for Windows files - a must-have for any system administrator even thinking about implementing Windows on Netware - it also offered additional information based on results from the real world.

Then, I looked in on several forums - areas where individuals can interact via the posting of mostly public and occasionally private messages. I quickly scanned the system for questions that either looked interesting for my own use or had been left unanswered or inadequately answered.

After seconding a recommendation given by a systems integrator from New York to a systems administrator at a small firm in Iowa, I switched from Netwire to a manufacturer's area on Compuserve. I wanted to see if new drivers were available for its boards. Drivers are, generically speaking, programs that allow various types of hardware to hardware or hardware to software communications - in this case allowing a network interface card or NIC to talk to the computer in which it resides and to other computers over the network. I subsequently found and downloaded the latest drivers.

It was getting late, so I logged off of Compuserve. I wanted to get a good night's sleep so I would be able to rescue those two clients first thing in the morning.

Now, four months later, I just finished fixing the same problem using the same files that I had previously downloaded. It surprised me (though it really shouldn't have, production cycles being what they are) that Novell hadn't yet incorporated the fix into it's production product during the intervening months. Still, my ability to access to Netwire saved me.

In addition to critical files such as this fix, I have obtained loads of great "handy" files. For instance, I now use a special version of the NDIR (the Netware directory command). If used with the usual parameters, it runs just like any other version of NDIR, but, if used all by itself, it's menu driven! For those interested, it is in part three (out of four) of the DOS/Windows update series of files available on Netwire and is the only version that I have ever seen (previous or since) that works this way. That series includes updates of many of the standard Netware commands and, except for the applicable time-based charge for accessing the system, is free for the taking.

Many useful utility programs are available. PERFORM II, for instance is Novell's standard way of testing the overall performance and efficiency of a network. Many manufacturers publish figures based on this common (and free) benchmark. You can run it to find inefficiencies or outright errors on your network. You can even use it to develop before and after figures when trying to fine-tune your system. There are hundreds of examples like this one.

Frequently a problem will crop up with some new software or hardware. A call to the manufacturer (or, heaven forbid, actually reading the documentation) will often result in finding a requirement for the latest version of portions of Netware or its associated drivers. Often, the only way to get these files is from Netwire (optionally via your chosen reseller).

I cannot overemphasize the importance of having access to Netwire for active Netware system administrators. In addition to the files themselves, the ability to get answers is incredibly helpful. Novell personnel usually respond within 24 to 48 hours, but other "visitors" often contribute answers faster. These answers from out in the field are sometimes more accurate and detailed than Novell's answers. Keep in mind that these fellow "visitors" aren't restricted (short of libel and slander laws) in mentioning brand names as Novell personnel sometimes are. Since Compuserve also contains many manufacturer forums, you get an added bonus of access to them.

If you don't have a Compuserve account now, get one! Just make sure it's in your company's name if you don't want to pay for it yourself. (Unless, of course, you want to access the games, or one of the sexual discussion forums - then get a second account in your own name - those expense audits can be nasty!) Look inside the box of your communications software. Many publishers include a free startup kit or at least a discount certificate. If that doesn't work, call Compuserve at 800-848-8199.

As always, questions, comments, and topic suggestions are always welcome.

GROUP THINK

by Wayne M. Krakau - *Chicago Computer Guide*, August 1992

Wayne Krakau welcomes all calls, comments, and questions from basement businesses to corporate MIS. Wayne is a partner in Krakau Business Computer Systems, a systems integration firm and a Novell Gold authorized and Lantastic 5 Star Dealer. He has been working with computers for 15 years. He has an MBA in Marketing and a BS in Information Science and was originally educated in Computer and Information Systems Analysis Engineering. He holds Novell CNE and CNI ratings, and LANDA Certifications in LAN Installation, Architecture, Connectivity, Maintenance, and Program Development. Wayne can be reached for questions, comments, or topic suggestions at (708)298-7695 or through the Computer Guide.

Number one on my list of the characteristics of a poorly managed and difficult to maintain networks is the lack of attention paid to the importance of organizing by groups.

In my constant encounters with disorganized Netware LANs, this is the single most time-consuming (translation: expensive) problem to overcome.

Sometimes the issue comes up during the mending of a completely disabled LAN. Sometimes it's during the installation of a new application. Often its something as simple as the need to add new users or just slightly alter the responsibilities and capabilities of existing users.

It doesn't matter how the problem is discovered, the results are the same. The system administrator (often a "civilian" doing double-duty as "The Computer Person") spends valuable time at company expense trying to make just one more quick fix to a complicated series of rights assignments (rights being Novell terminology for designating who has access to what files and to what degree). Additional time is lost painstakingly examining personal login scripts (login scripts are a sort of autoexec file used to control and enhance the act of logging in to the LAN), looking for common patterns. After reaching a point of diminishing returns, the system administrator having better things to do, a call for help is issued.

Usually the first thing I'll notice is that the system login script has no EXIT commands. To understand the implications of this, you need to know that (excluding the optional Netware Name Service software) there are three types of login scripts, system, default, and personal.

The system login scripts are manually created by the system administrator using the SYSCON utility. When it exists, it is run by everyone who accesses the network with the LOGIN command. It resides on volume SYS (the mandatory name for the first disk portion that Netware controls) in the PUBLIC directory under the name NET$LOG.DAT.

If the system login script doesn't exist, as in a newly installed network, the default login script is run. Also, if the system login script exists but no EXIT command is encountered, a search is automatically made for a user login script. If that's not found, then the default is run after the system. The default login script is embedded ("hardcoded" in programming terms) inside the LOGIN command itself (LOGIN.EXE). Copies of that command are in both the LOGIN and PUBLIC directories on the SYS volume. One bit of warning - the actual contents (the programming code) of the default login script are inappropriate for anything past the first few logins on a newly installed network, and, from an experienced programmer's point of view, are written in a terrible style. So, don't even think of using the listing of the default script as documented in Netware's manuals as an example.

When the default (and sometimes even a personal login script) login script is run subsequent to the system login script, the result is often a login script gumbo - lots of ingredients mixed so that it's difficult to identify one aspect of the flavor as belonging to one particular ingredient. Perhaps a forensic pathologist (that's like "Quincy" for afficionados of late night reruns) would enjoy reverse engineering this gumbo, but I sure wouldn't.

The third kind of login script, personal, runs after either the system (if it exists) or the default (if there's no system). It's created manually by either the system administrator or the individual to whom it belongs. It's in a file called LOGIN in the subdirectory of the MAIL directory that is named for the

internal hexadecimal (base 16) number that Netware uses to identify that user. Each person potentially has a personal login script with that same name, LOGIN, the identifying factor being the location.

The hexadecimal number is found by selecting Other Information from within User Information in the Syscon utility. The personal login script is created and can be copied within the Login Script section of User Information of Syscon. It can also be manually copied by simply copying the appropriate LOGIN file from subdirectory to subdirectory.

This logic (or algorithm, for my fellow computer geeks out there) for deciding which of these scripts run can be short circuited by the use of the EXIT command. If an EXIT command is encountered depending on the use of IF statements, no further login script statements of any kind are run. If the EXIT command stands alone, then the current script stops and no further scripts are started. If the EXIT command has a parameter (up the 18 characters surrounded by regular quotes) the command described within the quotes is immediately executed. Hence, the following line:

IF MEMBER OF PRIMATES THEN EXIT "BANANA"

would cause all members of the group called PRIMATES to execute a program called BANANA, skipping all other lines in the script being run and bypassing any further login script processing. (Plantains, anyone?)

Finally we're back to the original question regarding the implications of having no EXIT commands in the system login script. No EXIT (or bypassing EXITs with IF statements) means get out the ladle, we're into the login script gumbo, again.

Think about this. You have a three-user network. Each person has a personal login script. How long would it take you to modify the scripts to add a new application that needed access to several different directories and maybe a couple of environment variables, all varying by department (or with three people the theoretical department). I could already make a valid argument that it would take longer than if you had used a system login script with IF statements and EXITs.

Now, calculate how long it would take if your network grew to nine users. No, don't simply multiply by three - try cubing. Hence, nine divided by three is three. Three (the original number of users) to the third power (the result of the previous calculation) is twenty-seven. Multiply that twenty times the time you estimated for the original addition to the three-user network. For example, if ten minutes was your original estimate, two hundred and seventy minutes, or four and one-half hours is the estimate for changing the nine-user network. If the person who originally designed the personal login scripts is not available, you may have to double or even triple that estimate.

If you haven't figured it out yet, the preceding calculations were complete gibberish designed to impress upon the reader the overhead of indiscriminate use of login scripts. I will continue with this subject in next month's column.

CORRECTION

Due to an error in the editorial process (something to do with a coffee achiever trying to hold a mouse steady at three in the morning) portions of last month's column "Get Plugged In" were omitted. While the random chunks removed from the middle of the article were annoying (doing wonders for the logical continuity and grammatical structure of the sentences) the two paragraphs lopped off the end did contain at least one vital tidbit of information - how to contact Compuserve! Their (voice) phone number is 800-848-8199. My apologies to anyone inconvenienced by the omission of that number.

As always, questions, comments, and topic suggestions are always welcome.

GROUP THINK II, THE ADMINISTRATOR STRIKES BACK

by Wayne M. Krakau - *Chicago Computer Guide*, September 1992

Wayne Krakau welcomes all calls, comments, and questions from basement businesses to corporate MIS. Wayne is a partner in Krakau Business Computer Systems, a systems integration firm and a Novell Gold authorized and Lantastic 5 Star Dealer. He has been working with computers for 15 years. He has an MBA in Marketing and a BS in Information Science and was originally educated in Computer and Information Systems Analysis Engineering. He holds Novell CNE and CNI ratings, and LANDA Certifications in LAN Installation, Architecture, Connectivity, Maintenance, and Program Development. Wayne can be reached for questions, comments, or topic suggestions at (708)298-7695 or through the Computer Guide.

Last month we covered reasons to avoid the indiscriminate use of Personal Login Scripts as part of a network management strategy that I call "Group Think". This month we will deal with the other half of this grouping approach, rights management.

Rights are the characteristics that govern who has access to directories and files and just how much can they do. The usual way to manage rights is to give a trustee assignment covering a particular directory (or file in Netware 3.0 and above) to a group or an individual. Experienced systems administrators also will encounter rights masks (maximum in Netware 2.2 and below - inherited in 3.0 and above), but they are not important to a discussion of the benefits of group management.

The most frequent situation that I encounter involves an ad hoc approach to trustee assignments. People are given individual trustee assignments as applications are installed or upgraded. As new employees come on board or existing employees have changes in responsibilities, more individual trustee assignments are handed out. Existing assignments that are no longer appropriate or even completely obsolete are often left behind.

This rights "hash" (analogous to the login script "gumbo" of last month's column) eventually reaches the breaking point when the management load of making these constant trustee assignment changes ties up too much of the system administrator's time. That's when the decision to reorganize is made.

Reorganization requires a painstakingly detailed examination of the ingredients of the hash. This task is even more difficult when the impetus is the departure of the system administrator. The only documentation often exists only in the gray matter (that's brain cells) of the system administrator. The new (often shanghaied) administrator has to try to reverse engineer the pattern of trustee assignments, usually without any prior knowledge of the underlying logic.

The solution here, as in the previous discussion of login scripts, is to plan from the beginning to use grouping as the key element. The groups can be organized by functional area. The group ACCT could be created for the Accounting Department. It could be given a trustee assignment to the accounting program directory (and by implication, its subdirectories) that allowed seeing and executing the accounting programs (R and F for Read and File Scan in Netware 3.11 terminology). The assignment for the accounting data directory (and subdirectories) could allow changing, creating, renaming, and deleting files (depending on the requirements accounting software) files (RWCEMF for Read, Write, Create, Erase, Modify, and File Scan).

An underlying concept to remember is to assign the least amount of rights that the application requires. If the accounting system doesn't create and delete temporary or report files, then don't give the corresponding rights. If it needs those capabilities only in a particular subdirectory, just give those rights explicitly in that subdirectory.

The groups also can be organized by application. All users needing access to the main wordprocessor could be put in a group called WP. This group could be assigned rights to see and execute the programs in the wordprocessing program directory (RF for Read and File Scan). The same rights could apply to a directory set aside for boilerplate (common reusable) documents.

In many companies, a combination of the two techniques (department and application organization) is suitable. Usually it is useful to create groups for even the smallest number of people

rather than giving individual assignments. If frequent transfers or employee turnover is an issue, groups of one (used, for instance, to give a manager additional rights for applications associated with his or her department), become practical.

Usually it is easier to process a transfer or a new employee by simply adding and subtracting group memberships from an individual's characteristics in Netware's SYSCON utility than it is to make additional trustee assignments for every directory affected by the personnel change.

Even the documentation overhead (You ARE documenting your network, aren't you?) is lower with the grouping method. One document would list groups with their respective trustee assignments. A second would list directories and their respective trustees (a group or individual having a trustee assignment). A third would list groups with their respective members. A fourth would list individuals and their groups. A final document would contain any exceptions where an individual was given a trustee assignment. These documents are much shorter with grouping than with the more typical individual-based method.

One additional benefit of this organization method is the reduction in the security risk to the network. It is much easier to leave behind an inappropriate trustee assignment when assigning by individual. These leftover assignments can come back to haunt you if they result in files being viewed or altered by those who have no business accessing them.

Grouping is a way to save both internal and external costs. Remember, not only is a system administer's time worth money, but, if the network is disorganized enough to call in outside help, the cost can be budget-busting.

RANDOM NOTES

Most people have discovered that Netware 2.2 and below is an 80286 specific operating system that will run perfectly well on a 386 machine. That's why they don't call it Netware 286 anymore. Some people have noticed that Netware 3.11 is a 386 specific operating system that changes to a 486 specific system (it loads 486 commands) when it detects a 486 chip. That's why they don't call it Netware 386 anymore.

Now as we anxiously await Netware 4.0 (originally announced as Netware 3.2), I've discovered that it follows Novell's tradition of utilizing the most advanced instruction set available. When it detects a P5 chip (the code name of Intel's as yet unreleased 80586 chip), it becomes an 80586 specific operating system. When rumors of errors in the original sample set of P5 chips circulated, it turns out that it was Novell who discovered the problems. Since Netware 4.0 is currently the only 80586 specific software available, it was the only thorough test bed.

So much for the constant complaints of hardware always advancing faster than software.

BACKUP SAFETY

by Wayne M. Krakau - *Chicago Computer Guide*, October 1992

Wayne is a partner in Krakau Business Computer Systems, a systems integration firm and a Novell Gold Authorized and Lantastic Authorized Dealer. He has been working with computers for 16 years. He has an MBA in Marketing and a BS in Information Science and was originally educated in Computer and Information Systems Analysis Engineering. He holds Novell CNE and CNI ratings, and LANDA Certifications in LAN Installation, Architecture, Connectivity, Maintenance, and Program Development. Wayne can be reached for questions, comments, or topic suggestions at (708)298-7695 or through the Computer Guide.

This follow-up to my April column on backups was inspired by my recent encounter with yet another outrageously inadequate backup system. The hardware was perfectly good, a Colorado Memory Systems Jumbo (one of my favorites), but the batch files that automated its use were worse than useless.

The client's initial call (via a referral) was to discuss possible inadequacies in general. This led to a discussion of their worries about backup system. As a courtesy, I offered to walk them through a verify (also called a compare) using their last backup tape. Before I got to the verify, I had them do a tape status request so I could compare the raw number of bytes saved with the CHKVOL (the Netware command for checking the space on a network volume) results we had just obtained.

The status check showed zero bytes saved! No verify was required. Later research showed that all of their tapes were empty. This meant that they hadn't had a valid backup for almost a year. I immediately talked them through a manual backup of their Netware volume. That backup included an automatic verify, and subsequently was proven valid. They now have a new batch file driving their backup procedure, and it works.

This episode leads to my first recommendation. Test your backup system. Any time that the hardware or software on the backup computer is changed, the backup system should be tested. If any significant changes are made to the network as a whole, the backup system should be tested. Even if no changes are made to either, tests should be scheduled on a regular basis with once a year being a likely interval.

These tests should consist of two parts. The first is manually running a verify (compare) operation. Verify results often include perfectly innocuous warning messages about various Netware system files. If someone has used the network after the backup was made but before the test, additional warning messages about changed files also could be generated. This is the perfect time to learn how to tell the difference between genuinely hazardous error messages and these normal warnings. That skill could eliminate a potentially billable but unneeded call for help in the future.

The second part of these tests is a test restore. This part is potentially dangerous, so you have to make an honest appraisal of your skill level before deciding whether to do it yourself or call for assistance in this test. The first test is harmless enough that even a non-computer person can safely attempt it with little risk, but a restore can conceivably destroy data, so be frank in evaluating your own abilities.

A test restore could be done by creating extra directories and filling them with copies of word processing documents. Then use Netware's SYSCON utility to make some test users. Use the SYSCON or FILER utilities (or the GRANT command for masochists) to make several Trustee Assignments to these directories for your test users. If you have Netware 3.xx (that's the buzzword for 3.0 and above), you can make Trustee Assignments at the file level, too. Use the FILER utility or the FLAG command to assign various combinations of attributes to these test files. Finally, do either a full or partial backup, depending on time constraints.

After the backup is finished, delete all of the test users and the extra directories. Then try to restore them. After the restore, examine the directories, files, and test users to see that all rights and attributes were restored.

Another tool in your backup arsenal is the BINDFIX command. Its normal use is to repair the bindery (the database of all Netware objects and their characteristics), but it also creates valuable backup files of its own. As the first step in attempting to fix the bindery, it creates a backup. The original files are

NET$BIND.SYS and NET$BVAL.SYS for Netware 2.xx and NET$OBJ.SYS, NET$PROP.SYS, and NET$VAL.SYS for Netware 3.xx. They are flagged as SYSTEM-HIDDEN. The backup files have the suffix "OLD" and are automatically reflagged as READ-WRITE.

These backup files can be copied to a floppy disk and stored for emergencies. In the future, if you get a bindery error that can't be fixed by BINDFIX, rather than going through the time and trouble of running a restore, just copy the BINDFIX backup files back to the SYSTEM subdirectory and run BINDFIX's companion command, BINDREST. BINDREST deletes the existing (presumably bad) bindery files and replaces them with the old ones. The entire process takes about a minute!

Please note that both the BINDFIX and BINDREST commands are usually run while logged in as the SUPERVISOR with everyone else logged out. If someone else inadvertently accesses the bindery while these commands are running, even a perfectly good bindery can be corrupted.

The final backup tool is Netware's NBACKUP utility. While it is designed to be a complete backup utility, it is limited to using only DOS devices such as disk drives, or an obsolete (and quite rare) model of a Wangtek tape drive. This severely reduces its usefulness as a general-purpose backup instrument, but leaves one simple but effective capability. This utility can backup the bindery and the directory structure (along with its attendant Trustee Assignments) to a single floppy disk in just two to three minutes and, more importantly, can restore it in even less time.

The value of that last ability can be seen if we examine the process of repairing a network whose entire hard disk was trashed. This means we are stuck with either a new, empty disk or an old reformatted (and also empty) disk. The first step is reinstalling Netware. The second it to do two restores. Why two restores - because the Trustee Assignments are dependent upon the bindery as well as the directory and file structure. It takes two passes using a slow (usually) tape drive to get it all straight.

NBACKUP provides an alternative. First, use NBACKUP on a regular basis to create special backup disks. This is done by first selecting a local drive for the log files. Then select "*" (Netware commands understand "*" as a synonym for "*.*") for directories to include and files to include. Leave the setting for directories to exclude at its default of blank (meaning none). Finally (and here's the sneaky part) set the files to exclude to "*". This will force NBACKUP to get the bindery, the directory structure, and the trustee assignments, but skip the files.

When recovering from the disaster, first run NBACKUP to get back all three at once. Then make a single pass (restore) with your backup tape. With all but the fastest tape drives, you will save a lot of down time with this method. Its only real limitation is its required skill level. NBACKUP is not automateable and its interface is kind of obtuse. I only recommend it to clients who have someone at an appropriate level of computer knowledge (as well as the spare time to do this manual task).

These recommendations can provide a safer environment in which to run your network. The real key to all of them is planning.

HIGHLIGHTS

by Wayne M. Krakau - *Chicago Computer Guide*, November 1992

Wayne is a partner in Krakau Business Computer Systems, a systems integration firm and a Novell Gold Authorized and Lantastic Authorized Dealer. He has been working with computers for 16 years. He has an MBA in Marketing and a BS in Information Science and was originally educated in Computer and Information Systems Analysis Engineering. He holds Novell CNE and CNI ratings, and LANDA Certifications in LAN Installation, Architecture, Connectivity, Maintenance, and Program Development. Wayne can be reached for questions, comments, or topic suggestions at (708)298-7695 or through the Computer Guide.

The trade show season is upon us. Over the last month I have attended FIVE trade shows! I've got to - it's my job.

The first show was the Chicago Netware Users Conference sponsored by CAN^2 (pronounced "can squared"), the Chicago Area Novell Networkers. This was a small, focused show, targeting Netware administrators. It was also the site of a couple of shots in the war of the network analyzers. Analyzers take the raw information going across the network wire, interpret it to various degrees, and then display it for troubleshooting and optimizing the network.

Network General exhibited its Expert Sniffer, the latest version of its network analyzer. They are trying to leap-frog Novell's Lanalyzer, which has historically presented its test results in an easier to interpret format. Their new system is based on an expert system that can reduce the volume of information collected from its analysis of the traffic on the network wire down to a minimum of data.

It is a kind of management by exception. The expert system determines what it "thinks" is the appropriate information required to diagnose any problem it discovers. It then displays only that pertinent data and its associated diagnostic messages. The results are rather impressive.

The user interface was straightforward enough that I was able to anticipate the demonstrator's keystrokes on the way to the final determination of the sample problem. When she switched into the old-fashioned manual mode, however, I could just barely follow along. The new expert system really makes the difference. I can't wait to see Novell's counterstrike.

On the lower end of the network analyzer spectrum, Intel was promoting (though sadly, not demonstrating) its new software analyzer, Netsight Analyzer as an alternative to Novell's Lanalyzer for Netware (when you're on top, everybody wants a piece of your action). This product has great potential for those unwilling or unable to invest in a Sniffer or Lanalyzer.

The second show was the Midwest Computer Show sponsored by the Illinois CPA Society. It had some general-purpose exhibits, but was really aimed at CPA's, either for running their own offices, or to get them to suggest products for their clients.

The biggest trend was toward electronic claims filing. Many programs were displayed that provided this capability to existing tax processing systems without doing any tax processing themselves. Some systems were designed to be used by service bureaus that would in turn market their services to accountants.

An interesting side issue was the presentation by Arthur Anderson & Company of a pre-release version of their completely rewritten tax processing software, A-plusTax. It's about time! After spending untold hours attempting to beat that program into shape for my clients, I'm glad to see them get that series of programs out of the stone age. The current version doesn't meet the (admittedly strict) standards that I used to evaluate mainframe programs back in the seventies, much less microcomputer standards of usability and reliability in the nineties. The new version looks much better and, as a bonus, runs noticeably faster.

One aisle down from the Arthur Anderson booth, I found what may be the motivating factor for their rewrite, CCH Computax's ProSystem. This system competes almost module for module with the A-plusTax system. It has a very modern user interface, extensive context-sensitive help, full network compatibility, and it is already in wide distribution. Considering the CCH's (Commerce Clearing House for the uninitiated) reputation as a tax authority and their long history in computerized tax processing (I

worked on their mainframes processing returns during the 1976-77 tax season and it wasn't a new service for them then.) they should add some much-needed competition to this narrow field.

The Third show was LANDEX, given by LANDA, the Local Area Network Dealer's Association in cooperation with its Chicago Chapter. It was a dealer-only show that revolved around a series of seminars with only a small exhibit area.

Both the technical and business oriented seminars were very valuable, with the WAN (Wide Area Networking) session being my personal favorite. Several vendors expressed their surprise at the degree of cooperation and information sharing among theoretically competing LAN integrators.

In the exhibit area, Networth presented their Netware Application Engine, their concentrator-based platform for Novell's Run-time Netware and other dedicated server applications. It is a complete 486/33 microcomputer with RAM (random access memory), a hard disk, two internal Ethernet connections (to connect to the rest of the concentrator) and two 16-bit ISA (Industry Standard Architecture or AT-Bus) card slots. It plugs directly into their Series 4000 Intelligent Hubs (Ethernet concentrators). This is an incredibly bright idea. It is a great solution for running database servers, gateways, communications servers, fax servers, email servers, and any other application that would benefit from being run separately from the file server. Now that Run-time Netware is available (a stripped down version of Netware that can run an NLM or Netware Loadable Module, but can't be a file server), this product should be in high demand.

Also of note were the optical disk and SCSI (Small Computer Systems Interface) offerings of Micro Design International (MDI). So far they are the only company that I've found that completely integrates optical jukeboxes transparently into a Netware 3.11 environment. Other products that I have encountered make the optical subsystem a separately managed add-on to Netware. MDI can make Netware treat the jukebox just like any other drive. This allows them to tap into all of the inherent speed-up techniques of Netware, including caching, and elevator-seeking (the ability to reorder disk access requests into the most efficient sequence -- absolutely vital in an environment where inefficient ordering can cause a cartridge to be reloaded unnecessarily).

The fourth show was Softeach, a dealer-only series of seminars given by Merisel, a large microcomputer products distributor. There were sixteen seminars and two receptions in two days!

The announcement with the most potential interest was the introduction of a line of token ring products by Standard Microsystems Corporation (SMC). Current token ring cards use either the IBM chipset or the Texas instrument chipset, with minor variations based on the manufacturer. The only room for real innovations has been in the driver software (which allows the hardware and software to communicate). SMC is using the technology it purchased from Western Digital to attack performance, reliability, and pricing problems via hardware. Because of SMC's history in the ARCnet and Ethernet markets, they have a good chance of succeeding. This ought to shake up the marketplace!

The fifth (and at least for a few of weeks, the final) trade show was the Laptop Expo. I didn't see any individual standout products, but rather a new movement that impressed me.

The PCMCIA (Personal Computer Memory Card International Association) standard for add-on function cards for laptop computers was evident everywhere. The cards fit into any computer that has a slot following the PCMCIA specifications. There were many different functions available at the show. I saw Ethernet cards, high-speed combined fax/modem cards, 3270 emulation cards (for IBM mainframe terminal), applications software cards (like Microsoft's Excell), and RAM cards. By eliminating the need for proprietary expansion slots, this standard will remove the risk of having an "orphan" laptop.

I am still looking for a program to calculate just how much shoe-leather or seat-of-the-pants fabric that was sacrificed to these shows and seminars in the never-ending quest to keep up with changes in this industry.

SHOW, AND TELL

by Wayne M. Krakau - *Chicago Computer Guide*, December 1992

Wayne is a partner in Krakau Business Computer Systems, a systems integration firm that is a Novell Gold and Lantastic Five Star Authorized Dealer. He has been working with computers for 16 years. He has an MBA in Marketing and a BS in Information Science and was originally educated in Computer and Information Systems Analysis Engineering. He holds Novell CNE and CNI ratings, and LANDA Certifications in LAN Installation, Architecture, Connectivity, Maintenance, and Program Development. Wayne can be reached for questions, comments, or topic suggestions at (708)298-7695 or through the Computer Guide.

At last, the final trade show of the major Chicago trade show season - PC Expo! This show wasn't as focused as in previous years, there being no central pressing issue driving the industry (at the moment), and providing a theme.

In hardware, the NEC Silentwriter Models 95 and 97 caught my eye (NEC Technologies, Inc., Boxborough, Massachusetts). They are six and ten page per minute laser printers with an extra trick. They are upgradeable to include both send and receive faxing. They also can be purchased with the fax feature preinstalled as the 95FX and 97FX.

If the printer is busy, incoming faxes are held in memory and then interspersed with regular print jobs. Outgoing faxes are generated with the supplied DOS and Windows compatible software. The software even works over a network, providing a shared fax/printing solution.

I'm a big proponent of computerized fax, especially on LANs. The print quality of received faxes is as high as the most expensive plain-paper faxes. In addition, the computer generated outgoing faxes skip the scanning step that is the weakest link in the process. This means that no matter what the eventual output device (fax paper or plain paper), the documents look great. It also eliminates waiting first for your printout, and then for the fax machine.

These fax printers have an extra feature that activates when two of them fax to each other. The output is a full 300x300DPI (dots per inch), standard laser resolution (at least until the HP Laserjet IV was released) as opposed to fax resolution, 200x200DPI in the fine mode.

On top of all of this, the printers are very functional "normal" printers, too. They have Adobe Postscript and HPCL5 (the Laserjet III language) emulation with automatic language and port switching. They have a 250-sheet feeder standard (a second is optional) and a 10-sheet envelope feeder.

The 95FX (6PPM - page per minute) is a good choice for single users or for small workgroup LANs while the 97FX (10PPM) has the speed for larger LANs.

On the software side, PC DOCS Open (PC DOCS, Inc., Tallahassee, Florida) has just been released. This is a complete rewrite of this prominent document manager. Instead of just supporting WordPerfect in a DOS environment like its predecessor, it now supports multiple word processors in your choice of DOS or Windows (though the DOS version obviously isn't as flexible).

It now uses Sybase's SQL Server as its profile database (as opposed to its full text index), boosting its speed over a network. This is a very practical application of client-server technology. This strategy splits a database into two parts. The workstation part handles initial editing, help, typing, and translating requests into a common language (SQL - Structured Query Language). The server portion actually finds the data and sends only the results of the query back to the workstation. An additional advantage to this splitting of functions is that (within limits) you can mix and match back ends (the server portion) and front ends (the workstation portion). PC DOCS, for instance, is capable of using Microsoft SQL Server, Gupta, Netware SQL, or Oracle as alternatives to Sybase SQL Server. More client-server applications are sure to follow.

The Windows version of PC DOCS is particularly attractive in that almost every part of the screen is modifiable with a tool called DOCS Designer. Since document management is such a horizontal (non-industry specific) application, the ability to customize this powerful product is vital.

SHELL GAME

The standard way to access Netware from a workstation (commonly called the shell) has changed. For some time, there have been two coexisting standard types of shell. The old version had to be generated using a utility called WSGEN (or the older SHGEN), which combined information from the Novell-supplied IPX.OBJ file (IPX.OBJ contains IPX, the Internetwork Packet Exchange program and SPX, the Sequenced Packet Exchange program) with information in a driver file (a driver is used to allow two or more items of hardware and software to talk to each other) supplied by the vendor of the Network Interface Card (NIC - pronounced "Nick" - Buzzword Alert: It is considered uncouth to say "NIC card"; the proper expression is "NIC"). The resulting file, IPX.COM, is executed in tandem with NETX.COM. A parameter file called SHELL.CFG was used to alter characteristics of the shell.

Note that NETX is not the same as NETx. NETx is an older program where the small "x" was replaced with a number indicating the major DOS version number. If your workstation ran DOS 3.1 or DOS 3.31, for example, you used NET3. For DOS 4.XX you used NET4, and for DOS 5.XX, NET5. These NETx's have not been fully supported since the NETX was written. NETX automatically adapts itself to any DOS version from 3.00 on up. Unfortunately, Novell chose to name these two separate programs in a confusing manner so that they are both pronounced as the two word combination "NET X". NETx was generated by WSGEN (or SHGEN), the while NETX is obtained from Novell, usually via Netwire.

Sometime late last year, Novell stopped testing the WSGEN-based IPX programs. At that time, without any announcement that I've seen, they started giving full support only to the newer way of doing things. I found this out the hard way while trying to fix a problem that two clients were having with remote printing.

The new (actually it's been around for quite a while, but wasn't mandatory) shell is called the ODI (Open Datalink Interface) shell. Where WSGEN was used to tell the software the NIC hardware settings, the new NET.CFG file takes over. It is a replacement for the SHELL.CFG file. It accepts all of the old SHELL.CFG parameters plus some new ones. The most important new commands involve the NIC hardware settings. The lack of a separate generation process such as WSGEN will make maintenance much easier. It also lifts the burden from the NIC manufacturers of rewriting huge complicated driver files every time Novell makes the slightest change in the shell. Now, their drivers change only when they redesign their NICs.

One warning about the NET.CFG file (and its predecessor, SHELL.CFG) is in order. It must exist in the directory where its associated files (IPX, IPXODI, LSL, NETX, the NIC drive file) exist and you must physically be in that directory. This makes the proper procedure - probably in a batch file such as AUTOEXEC.BAT - to first, go to the subdirectory where all of the shell-related programs reside, and second, to execute the appropriate combination of these programs. Then you can switch back to your "regular" directory.

Since it is often impractical to change all of your workstation shells at once, a multistep strategy may be used. The initial procedure would be to upgrade any critical or special-purpose machines. This would include file servers, fax servers, communications servers, print servers, gateways, and workstations that handle remote printing (via RPRINTER.EXE or a third party equivalent).

After that, upgrade the shell only of those machines that are being accessed by the system administrator for troubleshooting, other upgrading, or installation. If the user is helped or trained on a particular workstation, that could also provide an opportunity for a shell upgrade. Eventually the whole network will run under ODI.

Writing about the need to have up-to-date software in order to get full vendor support reminded me of a recent conversation I had with a systems administrator who I was training. His company had just signed a contract to get Netware 2.2 installed on dozens of networks at their offices nationwide.

After demonstrating how complicated it is to install and maintain Netware 2.2 (and worse for earlier versions), we discussed his company's plans for the future of their LANs. Most of the things they wanted to do were either available only on 3.11 networks or could be done on 2.2 LANs only with serious compromises in speed, reliability, features, and administration overhead.

This discussion elicited a few choice words from him (which I can't print in this publication) regarding his company's strategy and the ethics of the vendors involved. Notice that "vendors" is plural. He informed me that there were separate contracts for the software, the hardware, the associated communications systems, the installation and support, and finally, the training. (Actually, other than this person and one of his colleagues being trained as troubleshooters by me, no other training was scheduled.)

After he reverted to repeatable English, we jokingly formulated a litmus test for choosing a reseller, but ended up with a potentially useful rule to follow. If, through advertising or personal contact, a reseller has proposed to sell Netware 2.2 (or worse, an earlier version) within the last "X" number of months, one of the following is true:

A. The reseller is so out of touch with the real world that he or she doesn't even realize that Netware 2.2 is a complete dead end, long relegated to second class status by software and hardware manufacturers, and more recently by Novell. Therefore, why should you deal with them?

B. The reseller knows the above information, but wants to get rid of the leftover stock of Netware 2.2, and has no ethical problems with tricking a client into buying an inappropriate product. Therefore, why should you deal with them?

C. The reseller knows the above information, but wants to make a lot of extra money in installation, troubleshooting, training, maintenance, and upgrading fees by selling a more primitive product with high cost requirements in all of these areas. This reseller must have no ethical problems with withholding the fact that the cost difference between 2.2 and 3.11 can be justified by its lower costs in these areas. Therefore, why should you deal with them?

The only point of disagreement was the number of months to fill in the "X". He was thinking of the eighteen to twenty-four month range, while I was more inclined towards the twenty-four to thirty month range (to cover the tail end of the Netware 2.15 era), both being based on December, 1992 as the starting period for the calculation. He did mention, however, that his company's LAN contract was issued in the first quarter of 1992. Hmmmmmmmmm?

SHOCK TREATMENT - THE INITIAL THERAPY

by Wayne M. Krakau - *Chicago Computer Guide*, January 1993

Wayne is a partner in Krakau Business Computer Systems, a systems integration firm and a Novell Gold Authorized and Lantastic Five Star Authorized Dealer. He has been working with computers for 16 years. He has an MBA in Marketing and a BS in Information Science and was originally educated in Computer and Information Systems Analysis Engineering. He holds Novell CNE and CNI ratings, and LANDA Certifications in LAN Installation, Architecture, Connectivity, Maintenance, and Program Development. Wayne can be reached for questions, comments, or topic suggestions at (708)298-7695 or through the Computer Guide.

Does anybody out there have decent power protection? From what I have observed, not many do. While most people know that some protection is needed, most are being scammed by companies selling inadequate protection either as a part of a deliberate strategy for winning bids or from just plain ignorance.

At the low end of the power protection equation are surge suppressors. Don't even think of buying suppressors without Electromagnetic Interference and Radio Frequency Interference (EMI/RFI) protection. These disturbances can cause mysterious lockups or even undetected data changes.

Every computer and peripheral device on the network should have at least this level of protection. For machines with modems, use suppressors that include phone line protection. On networks that don't use fiber optic cabling, the underlying material in the cable is copper, a great conductor. If an inadequately protected PC gets hit by a major power disturbance, either via its power cord or from the phone line (through a modem), not only will that PC potentially get barbecued, but the electrical disturbance can get to the LAN cable via the network interface card (NIC). Once on the LAN cable it can get to every part of the network, including the precious file server.

The scam here involves differentiating between the various type of suppressors and other devices that resemble suppressors. Many computers are being "protected" by multi-socket temporary taps (translation: fancy extension cords) rather than actual suppressors. They are normally marked as such on the back.

Other computers are protected by low-end suppressors that aren't designed to provide computer-grade protection. They are good for protecting household appliances and some business equipment (try toasters and pencil sharpeners), but shouldn't be used on computers or any other delicate electronics. (I wouldn't even use them on a modern stereo system.)

A step up from these devices are the low-end computer-grade suppressors. These provide a limited but reasonable amount of protection (varying by brand and model), but are made with circuitry that weakens every time it protects. They are like the ablative tiles on the nose of the space shuttle. Every re-entry (or electrical variance) damages the tiles (suppressor), eventually forcing replacement. The estimated life of these suppressors varies depending on the quality of the electrical power available, but is generally about one to two years. When (not if) they fail, they do so in the closed position. This means that they continue to provide power, but without any protection, while giving no indication of the failure. The replacement cost plus the danger of undetected failure keeps low-end computer grade suppressors from being cost effective. This is a prime scamming opportunity.

The best surge suppressors cost only a little more than the low-end models. I tend to use sixty dollars (list price - real prices vary) as the lower limit for a six or seven socket suppressor with EMI/RFI protection and a six foot or longer cord. Anything priced below that makes me suspicious. Upon further investigation, it usually turns out to be a low-end model.

These high-quality suppressors are not damaged by normal operations, and, if stressed by extraordinary electrical events, fail in the open (off) position. Many of them can also detect and report (via LEDs - Light Emitting Diodes) on faulty wiring. Naturally, power is cut if wiring errors are recognized. Most manufacturers have models available with telephone line protection for those computers with modems or fax modems.

Surge suppressors are only fully effective if the underlying voltage from the wall socket remains somewhere near 120 volts. If, for example, you are in an older building, you may find that the building wasn't designed for the electrical load that modern lighting and office equipment demands. Your power might not fail outright, but may vary considerably from the 120 volt standard resulting in brownouts and sustained overvoltages. The fan on your PC might audibly change speed, ranging from an unnaturally quiet whisper to full tornadic proportions. Your monitor image boogies to its own bizarre rhythm or shrinks and grows as if inhabited by some disembodied heavy breather.

This is where a voltage regulators (also called line conditioners) and power conditioners come in. They can filter as well as the best surge suppressors, but can also regulate the output voltage, even when the input voltage varies far from the 120 volt standard.

The traditional mechanism for this regulation is a ferro-resonant transformer. This is basically a large chunk of soft iron set up as a transformer similar to those that provide power to an external modem. In this case, however, the windings on each side of the transformer match so no voltage conversion takes place. This transformer and its associated circuitry can keep the output at a smooth 120 volts even if the input power gets very dirty or varies considerably from 120 volts.

A disadvantage of the plain transformer used in a line conditioner is the large amount of heat generated as the transformer absorbs the ups and downs of the incoming power waveform. This is accompanied by noise from the required fan. From past experience, I can tell you that there is an extra noise - a kind of raspy growl that can be quite loud - that emerges from the transformer every time there is a major voltage variance. It can be very startling if you haven't been warned. A more worrisome problem is the extra defects in electrical quality that can be created by the transformer itself.

Newer (and more desirable) types of voltage regulators, the power conditioners use advanced electronics combined with a special transformer to match the protective abilities of a plain transformer, but without the electrical distortions.

Another possible drawback to both kinds of conditioners is that they can suffer in comparison to some simple UPSs (Uninterruptible Power Supplies). The prices of these two categories can be so close that you may want to skip the conditioner and go right to the UPS.

In surge suppressors, my choice has recently changed to the American Power Conversion (APC) "Plus" series of suppressors. They provide extremely high quality protection at a competitive price. They have seven sockets as opposed to the more common six. They have cleverly placed their sockets sideways in relation to the longest dimension rather than in line. This means that a transformer (power brick) placed in a single socket won't block one or two additional sockets as it would in other manufacturers' products. Their final advantage is a warrantee that allows for replacement of the suppressor if it fails and for the protection for up to $25 thousand of equipment that is plugged into the suppressor.

As for conditioners (line and power), I have usually been able to use a small UPS as an alternative. If I do need voltage regulation, I would opt for a power conditioner rather than the older line conditioner.

Tune in next month for a continuation of this POWER-full subject, when I will cover UPSs.

SHOCK TREATMENT - THE FINAL THERAPY

by Wayne M. Krakau - *Chicago Computer Guide*, February 1993

Wayne is a partner in Krakau Business Computer Systems, a systems integration firm that is Novell Gold and Lantastic Five Star Authorized. He has been working with computers for 16 years. He holds an MBA in Marketing and a BS in Information Science and was originally educated in Computer and Information Systems Analysis Engineering. He holds Novell CNE and CNI ratings and LANDA Certifications in LAN Installation, Architecture, Connectivity, Maintenance, and Program Development. Wayne can be reached for questions, comments, or topic suggestions at (708)298-7695 or through the Computer Guide.

Continuing with last month's subject, the highest category of power protection, the Uninterruptible Power Supply (UPS), is the most varied. It is an absolute requirement for file servers. It can also be used for critical workstations, wiring hubs, and even phone systems.

In spite of some vendors' snide remarks to the contrary, the term UPS does cover devices that do not draw power continuously from a battery. The basic standard (set prior to the development of the IBM PC) involves saving the computer from damage and avoiding data loss, but does not specify the exact method. The standard does allow for minute power losses, as long as they can be absorbed by the built-in capacitance (energy holding capability) of the computer's power supply. Given the generosity of this definition, you are still better off maximizing the protection within practical and monetary limits. The current pricing of the different divisions of this category permits the selection of devices that greatly exceed this bare-minimum standard without breaking your budget.

At the top of the UPS food chain is the true online UPS. All power is supplied continuously through its battery, with the battery constantly recharged in turn. Many of these (especially the room-sized ones found in mainframe and minicomputer installations) were designed based on the linear power supplies found in pre-PC computers and other electronic devices. They output power with square or other unnaturally shaped waves. Since PCs use switching power supplies that sip power only off the very top and bottom of the power curve, square or other oddly shaped waves are radically inappropriate for PCs.

This restriction applies to all the different types of UPSs, not just true online ones. Look for a UPS with true sine wave output. While you can make do with a very tight pseudo-sine UPS, current pricing for all types of UPSs have eliminated any price difference between sine and pseudo-sine so don't bother.

True online UPSs don't really provide extra benefits for PCs, given the design of the PC's power supply and the advances in other types of UPSs. They also have several disadvantages. The battery is constantly in use and therefore wears out fast. Converting electricity from AC (from the socket) to DC (for the battery) and then back to AC (for the PC) is not 100% efficient. There is a major loss of power due to this inherent inefficiency. That means higher electrical bills. The lost power is converted to heat. This means more noise (and even more wasted energy) due to the required large, fast fan. The extra heat also demands more air conditioning, wasting even more energy. Finally, true online UPSs are the most expensive type of UPS.

SPSs (Standby Power Supplies) wait for a major power drop (a blackout or very severe brownout) to kick in with battery power. As long as they are fast enough, they meet the theoretical minimum limits for protecting a PC. Generally, 4ms (milliseconds) is considered fast enough to switch to battery power. The problem is that this doesn't take into effect the time it takes to recognize the power outage. Power doesn't usually sharply drop to zero. It gradually drops randomly, getting more erratic as it falls. Only when it reaches a predesignated (by the manufacturer) limit will the switching circuits start. While power drops, the PC can be supplied with wildly varying and potentially damaging power. Also, the simplest SPSs don't even include surge suppression. They feed raw AC power straight through to the PC.

More advanced SPSs have added circuits to provide surge suppression and to improve detection of uneven power losses. They also synchronize better with the waveform upon both the switch to the battery and the eventual switch back to line power. Unsynchronized switching stresses the PC.

Hybrid UPSs are a combination of an SPS and a power conditioner in one box. The power conditioner portion provides regulation and suppression while the SPS portion protects against outages. The transformer allows the device to stay on line power much longer before being forced to switch to the battery. A regular SPS must switch even for a brownout. The transformer has one extra feature. Its inherent capacitance causes it to keep providing power for a short period even after line power has failed. This period is just long enough to smooth out the transition to battery, providing better quality power right up to the switch.

The latest incarnation of the UPS is the line interactive UPS. This uses advanced circuitry combined with a clever design trick to match the features of the hybrid and adds sophisticated monitoring and self-diagnostic features without bothering with a ferro-resonant transformer. The clever trick is to connect the invertor (the mechanism that converts DC power to AC power and vice versa) directly to the output line. This is my UPS of choice.

In choosing a UPS you must select a power rating in Volt-Amps (volts times amps), rather than the more commonly known watts. The formula for conversion is:

Watts = Volt-Amps X Power Factor

where the Power Factor for PCs is 0.6 (that's six-tenths).

Keeping in mind that as a systems integrator I can sell any brand, my current favorite brand of both surge suppressors (for the last year) and line interactive UPSs (for the last two and a half years) is American Power Conversion (APC). Their combination of quality products, great warranty ($25,000 worth of protection for damaged equipment), and outstanding tech support have kept me in their fold.

There is one final warning about this column. For the sake of brevity, it contains many oversimplifications and gross generalities about a very complicated subject. While engineers (including the professor who originally taught me wave propagation theory in Energy Engineering class) might nitpick, I believe that the basics are accurately presented. Now, where did I put that Jolt Cola?

DOWN FOR THE COUNT

by Wayne M. Krakau - *Chicago Computer Guide*, March 1993

Wayne is a partner in Krakau Business Computer Systems, a systems integration firm that is Novell Gold and Lantastic Five Star Authorized. He has been working with computers for 16 years. He holds an MBA in Marketing and a BS in Information Science and was originally educated in Computer and Information Systems Analysis Engineering. He holds Novell CNE and CNI ratings and all 5 of the original LANDA Certifications. Wayne can be reached for questions, comments, or topic suggestions at (708)298-7695 or through the Computer Guide.

It started with a phone call from a Novell Authorized LAN dealer asking for assistance in designing one of his first networks. He wanted us to work with him in interviewing his client, a law office, and planning a new network where only a couple of stand-alone PCs were already in use.

During the interview I discovered that many of the major decisions had already been made. The client's informal personal computer advisor (see my column from June of 1992 entitled "The Brother-In-Law Syndrome" for my opinion on amateur "advisors") had convinced the client to go with thin Ethernet using 8-bit cards in no-name (literally without a name plate of any kind) AT clones. Even the file server was to be an unsupportable clone of dubious heritage with only an 8-bit (XT class) NIC (network interface card). Multimate (YECHHHH!) was to be the main application program.

The senior partner in charge of the law firm's LAN project refused most of my advice. This included having a supportable (and Novell certified) machine as a file server, and equipping at least that machine with a 16-bit (AT class) NIC. It also included discussing the merits of UTP (unshielded twisted pair) wiring, especially in an office that wasn't laid out in a linear fashion - the only method supported by thin Ethernet. A requested check on local building codes regarding cable plants was also refused.

Considering the preeminence of WordPerfect in the legal market, I also questioned the choice of Multimate. It was invented to ease the transition from Wang dedicated word processing machines to personal computers and had no particular appeal as the choice for a firm that never had Wangs.

Later discussions with the LAN dealer were just as unproductive. He declined my offer of installation services. He also rejected my bid on the cabling. He told me that he had an arrangement with another firm that would do it for less than one-third of my price. When I cautioned him about the hazards of amateur wiring installations, he assured me that the firm he dealt with was a professional LAN cabling company. That ended our discussion.

A few weeks later I received another call from the same LAN dealer. He could not print over the network. He was also having difficulty getting Multimate and some other applications installed. I agreed to come to his client's office to fix the problems.

Using the P (Printer) and Q (Queue) commands on the main console of the Netware 286 server - in those days that name was still used - I discovered that no printers or queues had been defined. That was count one.

Prior to running Netgen, the installation program, to create a printer definition, I checked the underlying server configuration with the CONFIG command. That's when I realized that the server was configured to hold two NICs. Since this was a small, single-office LAN, this surprised me. A quick inspection of the back of the file server revealed only one card. The original installation had been blown. A second NIC had been defined where no actual card existed, causing a performance loss and providing a potential point for future system failure. That was count two.

The file server was the last machine in the daisy-chain of this thin Ethernet. Thin Ethernet demands that each end of the "chain" is terminated with a 50-ohm terminator, one of which (arbitrarily chosen) is grounded. The server had a home-made terminator consisting of a naked resistor poorly soldered to a short length of wire. It was a mess. That was count three.

Since this end of the network was ungrounded, I went to the farthest office to check the last machine in the original plan. It also had a bizarre terminator. That was count four.

Neither end of the LAN was grounded. That was count five.

Now, it was back to the file server to trace the rest of the network. The segment from the file server to the first workstation looked just fine, but the next segment was lying on top of an under-desk heating unit. The immediate problem was the electrical interference, with the long term problem being the heat induced degeneration of the cable. Naturally, I removed the heating unit. That was count six.

I decided to examine the first workstation. It booted off a floppy disk and automatically logged in to the network. The distinctive sound of a hard disk spinning up to speed emanated from the machine. Sure enough, it contained a perfectly good, bootable hard disk. The floppy boot disk was worse than useless - it slowed the boot process as well as access to DOS commands. That was count seven.

Based on a report that some of the workstations would randomly lose access to the network, I ran my favorite LAN diagnostic software, TXD, from Thomas-Conrad Corporation (Austin, TX). It uses Netware's built-in error detection capability that normal programs ignore. Upon initialization, it scans the network (or internetwork, if applicable) for all servers and workstations. In this scan, it couldn't find the workstation on which it was running! This was not a good sign.

Next month, the conclusion of this network horror story.

DOWN FOR THE COUNT - PART II

by Wayne M. Krakau - *Chicago Computer Guide*, April 1993

Wayne is a partner in Krakau Business Computer Systems, a systems integration firm that is Novell Gold and Lantastic Five Star Authorized. He has been working with computers for 16 years. He holds an MBA in Marketing and a BS in Information Science and was originally educated in Computer and Information Systems Analysis Engineering. He holds Novell CNE and CNI ratings and all 5 of the original LANDA Certifications. Wayne can be reached for questions, comments, or topic suggestions at (708)298-7695 or through the Computer Guide.

To review, we had discovered the following errors in examining the network in question:

1. No printers or queues had been defined;
2. Two NICs (Network Interface Cards) had been defined when only one existed.
3. One end of network cable had a messy, home-made terminator.
4. The other end had a similarly butchered terminator.
5. Neither end of the LAN cable was grounded.
6. The cable was resting directly on an under-desk heating unit.
7. The first workstation was being booted off a floppy disk even though it had a useable hard disk.

I had just run TXD (Thomas-Conrad Corporation, Austin, Texas), the network diagnostic program and was surprised to discover that it could not even find the workstation that was running it as it scanned the network. This indicated a distortion of the electrical signal on the wire so terrible that the NIC in this workstation couldn't detect a signal that it had just transmitted! What's the old saw about it being too noisy to hear yourself think?

Since I was still in the open main reception area of the lawyer's office, it was easy to trace the cable to the next workstation. It had no obvious problems, but something just didn't feel right. No, I'm not Yoda with some amazing ability to tap into the "Force" (my whiskers aren't as scraggly and I'm a lot taller), but I am alert enough to notice a tactile difference between cables when I have recently done a lot of work with that type of cable. The fact that I couldn't find any standard marking on the cable strengthened my suspicions.

This workstation actually booted off its own hard disk (Hey, that's one in a row!), so I continued the search, following the cable until it entered the wall. From there, I could deduce that it went straight up into the ceiling to pass over an intervening doorway on the way to the first office. Using a ladder, I removed the ceiling panels adjacent to the doorway. I located the cable and started tracing its path, hand-over-hand along the edge of the doorway.

That's where I found an inline connector (the kind used to connect two lengths of cable) that had a base with two holes. It was attached with two screws through those holes to the side of the office's massive heating unit in the ceiling! This metal to metal contact formed a ground. Thin Ethernet needs to be grounded at one arbitrarily chosen end, not in the middle! That was count 8.

Surprised, but undaunted, I continued the search. (Alright, it's not a quest for the lost Ark, but it does require some detective skills.) The cable continued through the ceiling. That's when I realized that this was still plain old PVC (polyvinyl chloride) sheathed cable even though it was now running through a drop ceiling. This was a major violation of the local building and fire codes which specify the use of either PVC cable inside of conduit or a plenum (a metal cable channel) or the use of plenum cable (Teflon coated instead of PVC) on its own. Some localities don't even allow Teflon cable.

Burning PVC gives off incredible quantities of thick smoke, filled with poison gasses (including that old executioner's favorite, cyanide gas). It also burns as fast as fuse wire. (Picture the introductory scene of both the old and new "Mission Impossible" series.) Violating these codes can bring both the obvious fines and the less apparent criminal and civil liability if anyone was injured or even killed in a fire. That was count 9.

Next, I found a T-connector. This is the cabling version of a three-way intersection. It was the same device that is used to attach a PC to the network cable via its NIC. What in the world was it doing in the middle of a length of a cable strung through a drop ceiling? Did someone expect to put a PC up there?

My questions were answered almost immediately when I pulled a rats' nest of cable out of the wall insulation consisting of multiple t-connectors strung together with short lengths of cable to form, in effect, a star configuration! It looked like something that an 8 year old put together while tinkering with an Erector Set. (Am I dating myself? Does anyone else remember them?)

I stared at the "star" while I tried to figure out just how much drugs and/or alcohol had been consumed before this cable plant was designed. Surely, I thought, no one in a sober state would do it, much less the allegedly "professional" cabling firm that the original LAN dealer hired.

I still could not believe what I had found, so I got off the ladder and checked each office. The confirmation was there. Each PC had its own terminator (with bonus points for each one being hand-made junk). It was true. The network had been designed as a star with one long leg (two workstations and the server were on the front office leg), with the center of the star made of the group of t-connectors. The end of each leg of the star was terminated (poorly). This was count 10 (TILT!!!!). Let's turn over the numbers, now. This game is lost.

Just to cap off my day, while I checked the individual offices, I finally found a length of cable that had a label. It was Radio Shack CB cable!!!!!! Yes, that's CB as in Citizens' Band Radio. Thin Ethernet cable is supposed to be RG-58 A/U rated at 50 Ohms impedance. Even most RG-58 cable isn't good enough. In a pinch, I have attempted used Radio Shack RG-58, and it wasn't the appropriate quality rating for Ethernet. (This was a while back, so they might sell good stuff, now.) Who knows what kind of signal you can get over this CB cable.

The LAN dealer (remember, he's Novell Authorized) referred me to the cable installer. The head (term used very loosely) of that firm stated that they specialized in robotics! I have no idea what they were doing with cable plants. He (you knew it was going to be a "he", already, if you had read my column on "The Brother-in-Law Syndrome") stated that since they knew wiring (their systems contain many wires) they were eminently qualified. (Riiiiiiiiight. And I'll let you have this genuine, real, honest-to-God Rolex watch for only $39.93 in three easy installments.)

When I explained that the cable plant didn't follow standards, he became very defensive and demanded to know who the "jerk" was that made up these stupid standards. When I mentioned that the IEEE (Institute for Electrical and Electronics Engineers, the governing body for many standards, including Ethernet, which they call 802.3), he asked me who "the hell" they were and why were they sticking their noses into his cable systems. This from a person allegedly involved in robotics, where electronics expertise should be paramount.

I tried to explain to him that the "rules" were not arbitrary, as he suggested, but were based on two ideas. The first was a bit of common sense in that the standards had to be commercially feasible. The second, and most important, was based on THE LAWS OF PHYSICS!!!!! As long as we can't "warp" the space-time continuum (ala Star Trek) we were stuck with these "laws", and he had better learn to live with them or he should leave this dimension for one in which his own pet laws applied!

In the end, the entire cable plant was yanked and restrung. Other than still using PVC sheathed cable, all standards were followed. I wish that I could say that this tale ended with everyone living happily ever after, but I can't. First, the lawyer attempted to get me to pirate Multimate across the network. I declined, installed it so that it could be used by only one key legal secretary, and suggested that he get the appropriate number of licenses (preferably not Multimate).

Later, I learned that the lawyer stiffed the LAN dealer to the tune of $3000 supposedly in a dispute over modems, but really to save money. He dared the dealer to sue him, stating that he could send a team of his in-house attorneys into the battle and keep him tied up in court for years, with no cost to himself, but with high legal costs to the dealer.

Subsequently, this same lawyer tried to alter the network configuration to allow him to pirate Multimate so everyone could use it. This attempt messed things up so badly that the network wouldn't even come up anymore. His request to the LAN dealer for help was met with a demand for prepayment at a premium rate. The lawyer refused. The last I heard, he was calling every LAN dealer and "consultant"

he could find to locate one who knew enough to repair his system, but was unethical enough to aid him in pirating Multimate and other software.

Somehow, I just can't bring myself to wish him good luck. Perhaps I should just refer him to the cable installer. That seems like a match made in heaven (or parts south).

TAKE AIIM

by Wayne M. Krakau - *Chicago Computer Guide*, May 1993

Wayne is a partner in Krakau Business Computer Systems, a systems integration firm that is Novell Gold and Lantastic Five Star Authorized. He has been working with computers for 16 years. He holds an MBA in Marketing and a BS in Information Science and was originally educated in Computer and Information Systems Analysis Engineering. He holds Novell CNE and CNI ratings and all 5 of the original LANDA Certifications. Wayne can be reached for questions, comments, or topic suggestions at (708)298-7695 or through the Computer Guide.

The AIIM Show hit town April 5th through 8th at McCormick Place. It was sponsored by the Association for Information and Image Management. The really unusual thing about this year's exhibition was that the biggest breakthrough was not technological - it was political.

Nien-Ling Wayman, president and founder of Compulink Management Center (Torrance, CA), the makers of the LaserFiche family of PC/LAN based image and document management products has been elected to the AIIM board of directors. This is the first infiltration by a representative of an open PC/LAN product into an organization dominated by manufacturers of large, expensive, proprietary systems.

AIIM is just starting to experience the same changes that publishing has gone through over the last few years. The shows were originally ruled by large proprietary vendors fighting to lock customers into their own unique (and incompatible) way of doing things.

This gradually transitioned into an intermediate stage wherein these manufacturers offered scaled down systems in an attempt to stem the tide. These shrunken systems contained standard pieces of PC technology with a different brand name slapped on the front and the price tag boosted to levels barely below the old systems.

In viewing publishing shows, I was often able to identify the real manufacturer's names and even the exact model numbers of these camouflaged items. There was usually a major price difference - multipliers of three to four were not uncommon. Eventually, the importance of the large system companies diminished. They still exhibit in some of the publishing shows, but only to cater to the most conservative of the attendees.

The same process is happening in AIIM. Ms. Wayman's election is but the first of what will certainly be an ongoing series of incremental changes in the balance of power. That's evolution. (The mammals are coming! The mammals are coming!)

Let's hope that she has the strength to refrain from laughing aloud when one of her colleagues tells her that "small" systems CAN'T do some task that she knows has been done on "small" systems for several years. That's always been one of the more difficult things for me.

This election also gives Novell an ally within AIIM, since Compulink is an early implementor of the new Novell-Kodak image enabled Netware system. This system allows image-oriented products to run as an NLM (Netware Loadable Module) inside of a Netware file server providing true client-server technology (where the file server does the searching, not the workstation) to image databases. Since Novell has declared that image management is the next great network application, this kind of clout is of great strategic importance.

There were some interesting technical developments at the AIIM Show as well. Optical jukeboxes (also called autochanger) were there in force. The devices allow access to multiple optical disk cartridges via software commands. They work very much like an old Wurlitzer jukebox, with cartridges in place of records. They do, however, allow multiple drives within the larger capacity systems. The drives can be WORM (Write Once, Read Many or Mostly), read/write, multifunction, or CD-ROM (Compact Disc - Read Only Memory).

The largest of these jukeboxes now holds 1.4 Terabytes (trillions of characters) using twelve-inch cartridges! In round numbers that's a one and a four followed by eleven zeros! Each byte represents the equivalent of one character. They don't even bother attaching them to the file server. They attach them directly to the Ethernet cable. High-speed magnetic media (that's means regular disk drives) are used to

cache the information from the optical disks. Before you jump to the conclusion that their capacity is too great to be useable, remember that Netware 3.11 can theoretically handle twenty-two of these drives with plenty of room left for software and indexes. Netware is designed to swallow thirty-two terabytes of storage.

The phenomenon of multiple brands and price ranges for the same device that happened in publishing was very apparent in another model of optical jukebox. It was sold under so many different brand names that I lost count. This jukebox holds up to one thousand and fifty-four cartridges while configured with one drive of your choice - it's made to hold different types and brands of drives. In a more realistic configuration, multiple drives would be implemented, with each additional drive reducing the cartridge capacity by ten. Depending on the drive chosen, each 5.25-inch cartridge can hold from six hundred megabytes (millions of characters) to one gigabyte (billions of characters).

Other breakthroughs were in the scanning area. Bell & Howell's forty-eight page per minute dual sided scanners scan both sides of a sheet at once, giving an effective throughput of ninety-six "faces" per minute. Kodak was showing single-side models with speeds in excess of one hundred and forty pages per minute. TDC was showing a wide variety of high-speed scanners with both single and dual sided capabilities. All of these brands can be connected to a PC via Kofax's newest line of high speed controller/compression cards.

As I finish this column, I have just received word that I won a scanner and its companion software by dropping my business card into a fishbowl during the AIIM Show. Hmmmmm. I started thinking about winning this prize after twenty-plus years of attending trade shows without winning anything. Now I know why Lotto works - there's always hope.

USEFUL STUFF

by Wayne M. Krakau - *Chicago Computer Guide*, June 1993

Wayne is a partner in Krakau Business Computer Systems, a systems integration firm that is Novell Gold and Lantastic Five Star Authorized. He has been working with computers for 16 years. He holds an MBA in Marketing and a BS in Information Science and was originally educated in Computer and Information Systems Analysis Engineering. He holds Novell CNI and Enterprise CNE ratings and all 5 of the original LANDA Certifications. Wayne can be reached for questions, comments, or topic suggestions at (708)298-7695 or through the Computer Guide.

YOU DO THAT VOODOO

The "Voodoo" series of books by Kay Yarborough Nelson (Ventana Press, Chapel Hill, SC) is amazing. Ms. Yarborough uses an overall theme of "Anything that I don't understand completely is magic, and I'll never understand computers. I just try to get some work out of them - and probably you do, too." This incredibly refreshing outlook permeates the whole series of books, presenting the reader with the "Tips and Tricks with an Attitude" as promised by the common subtitle.

Having read, reread, and - more importantly - used three of the series, "Voodoo DOS", "Voodoo Windows", and "Voodoo Wordperfect", I have become a great fan of the author. I have also skimmed, but not thoroughly read (YET!) "Voodoo Wordperfect", and "Voodoo Mac", and they hold the promise of similar quality.

My expectations for the first of the series that I read, "Voodoo Windows" were low, having worked with Windows extensively over the last few years. Wrong! I found so many valuable tips and tricks that I spent several hours implementing and experimenting with them. My Windows now looks and runs like no other. It is customized to the hilt. The amazing part is that I enjoyed reading the book. The "Attitude" mentioned in the subtitle came through in the form of irreverent humor and a genuine sense of wonder about the subject matter.

The author presents material that is complicated and obscure enough to intrigue the hard-core techno-geek. (Let's see - pen . . . penlight . . . utility knife . . . collapsible pointer . . . notepad. Whew! Thank God, no pocket protector - Wow! That was close. I guess the therapy is working.) The key factor is that the information is offered with no assumptions as to prior experience. A beginner can easily use the tips, too.

"Voodoo DOS" was only slightly less valuable to me. (I was experimenting with MS-DOS's predecessor, CP/M before MS-DOS was even a glimmer in Billionaire Bill's acquisional eye, so there was slightly less room for improvement.) "Voodoo Wordperfect for Windows" was extremely helpful, since I had just transitioned to the product. Both were as readable as the first. I can't wait to get to the rest of the series. At $19.95 a pop, these books should not be missed.

WINSOME WINCIM

All right. There's no printed documentation. Three-fourths of all of the icons and menu-based commands jump directly to dialing with no warning, without asking for confirmation, and with no sure way to abort the call short of turning off your modem. An inauspicious start, to be sure, but WINCIM (Windows Compuserve Information Manager) is a worthwhile program.

Anyone involved with Netware should already have a Compuserve account - it's the only way to survive. With WINCIM, you cannot only survive - you can thrive. It's actually fun to have simultaneous access to several sections within Compuserve and to switch among them at will. You can download multiple files, read your mail, and browse through forum message at the same time.

As you traverse the message threads you can display their pattern in graphical tree format. If you mark individual messages or threads, WINCIM will offer to download the marked items. This feature allows the money-saving technique of off-line reading. Offline message creation and replying are also

available. If you plan to return to the same section repeatedly, just add that area to your "Favorite Places" list. After that, just click on it and you'll be there.

WINCIM is a valuable tool for Windows-based access to Compuserve. Just order it while online (GO WINCIM) to save lots of time and money.

THE ULTIMATE UNDO

The idea behind Uninstaller (MicroHelp, Inc., Marietta, GA) is so obvious that I'm surprised that no one else has thought of it. Have you ever had a Windows application that you wanted to remove? Maybe you didn't like it. Maybe the upgrade installation program is buggy and you want to install the new version from scratch. You might have damaged files and want to reinstall a program from scratch to overcome errors. Or, maybe you wanted to test a program (for a review?) and get rid of it after the evaluation. Most windows programs won't automatically remove themselves and manual removal is normally incomplete and can even be hazardous. That's where Uninstaller comes in to play.

Uninstaller analyzes Windows program files to find what other program files they access and what segments of .INI (Windows initialization parameter) files affect them. It then offers to remove these programs and .INI options either one by one or in groups. Each step is carefully documented with context sensitive help accompanied by warnings when necessary. A final option to abort is given at any point where a change becomes unrecoverable. Even multiple programs from a single vendor that access common families of subprograms can be tracked down and safely eliminated.

My primary computer is constantly used to evaluate products, and is therefore littered with leftover pieces of Windows programs. Uninstaller found and removed them in minutes. It also easily eliminated some complete programs that were awaiting disposal. I netted many more megabytes of free space than I had expected, all without errors.

The program has some neat tricks for LANs, too. When uninstalling a LAN-based program, it can leave behind a dummy program (called a stub) to prevent individual workstations from getting error messages when trying to access a removed program. Instead of an error message, a warning is displayed explaining that the application in question is no longer available on the LAN. It then offers the option to clean up the Windows on the local hard disk. When all of the workstations have been cleared the administrator can manually request the removal of the stub.

This program is essential for anyone who evaluates Windows programs and quite desirable for any Windows users.

Strange Bedfellows

A few years ago, an organization was founded to meet the needs of LAN VARs (Value Added Reseller) and Lan-oriented systems integrators. The organization grew, to the benefit of its members. Then they started allowing anyone in who had any LAN authorizations - even those from tiny LAN companies who have no criteria for authorization save a state issued sales tax number. Later, they allowed stores that were little more than mail order houses and warehouse outlets to join the organization. Now I have learned that LANDA, the LAN Dealers Association is merging with NOMDA, the National Office Machine Dealers Association. At this rate I expect to hear of an upcoming merger of the American Cancer Society with a major tobacco company.

TOO SCUZZY A STANDARD

by Wayne M. Krakau - *Chicago Computer Guide*, July 1993

Wayne is a partner in Krakau Business Computer Systems, a systems integration firm that is Novell Gold and Lantastic Five Star Authorized. He has been working with computers for 16 years. He holds an MBA in Marketing and a BS in Information Science and was originally educated in Computer and Information Systems Analysis Engineering. He holds Novell CNI and Enterprise CNE ratings and all 5 of the original LANDA Certifications. Wayne can be reached for questions, comments, or topic suggestions at (708)298-7695 or through the Computer Guide.

It's pronounced "Scuzzy Too", written "SCSI-II", and means Small Computer Systems Interface, Version Two. It was designed as the successor to the original SCSI standard to be a common connection method for many different computer peripherals.

These peripherals include: hard disk drives; removable cartridge disk drives, read-write, WORM (Write Once-Read Many or Mostly), and CD-ROM (Compact Disc Read-Only Memory) optical disk drives; scanners; and even printers.

Though the SCSI circuit board placed inside the PC is usually called a controller, it really is a host bus adapter (HBA). The controlling circuitry is normally inside of or attached to each peripheral. The HBA just provides a means for the SCSI devices to communicate with the rest of the computer system.

Each HBA can handle up to seven peripherals, daisy-chained from each one's output port (connector) to the next one's input port with each port arbitrarily designated as input or output. On each end of the bus (the chain of devices) a terminating resistor (the original Terminator) is required. The HBA is equipped with ports to accommodate both internal and external devices.

The SCSI II standard was written to allow dissimilar devices from different (and often competing) manufacturers to coexist happily without unduly complicating the setup of the total computer system. That's the theory. Try to implement it, however, and reality smacks you in the face.

First, there are multiple types of HBAs. There are 8-bit HBAs, for XT class machines that also can be used in ISA (AT-class, known formally as Industry Standard Architecture) machines for low-performance applications such as with slower CD-ROM drives. There are standard 16-bit ISA HBAs. There are pseudo-busmastering 16-bit ISA HBAs for higher performance. Both also can be used in EISA (Extended Industry Standard Architecture) machines though the use of the pseudo-busmastering type will force a restriction to 16MB (megabytes) of random-access memory in the computer.

Then there are the specialty HBAs. These include EISA and Microchannel (IBM's proprietary bus structure used in some of their PS/2 machines) computers. Here are even more choices. While both busses have a 32-bit interface, not all cards with a 32-bit physical interface fully utilize it. (This warning holds true for other EISA and Microchannel boards, not just SCSI HBAs.) They may have only 16-bit internal circuitry. Remember, most computer performance calculations are based on the weakest link principle - one low-performance part (the 16-bit portion) working in conjunction with high performance parts (the 32-bit physical connector) effectively restricts overall performance of the system to its own limits. Occasionally the same effect is also seen with 16-bit boards (both SCSI and others) that contain chips that can only handle eight bits of information. VGA cards and sound boards have historically been notorious for this trick.

For these specialty busses, one other issue appears - busmastering. Both EISA and Microchannel allow for several true busmasters, as opposed to the single pseudo-busmaster (also called a DMA or Direct Memory Addressing busmaster) supported by ISA machines. Some boards for these machines (again including both SCSI and others) are not busmastering, thereby lacking the characteristics needed to take full advantage of the performance enhancing features of the specialty busses.

Besides the previously mentioned generic (my own term for being usable for theoretically any brand or type of SCSI device) SCSI HBAs there are many variations on the SCSI theme. There are still devices (HBAs and peripherals) available that follow the original SCSI standard, or to make matters even more confusing many that only partially support SCSI-II! Luckily, SCSI-II was designed to accept the older standard and the hybrids, but comparing performance can be ridiculously convoluted.

And now for the really fun part - proprietary and limited SCSI HBAs. Sound boards with built-in SCSI, scanners, older removable cartridge drives, and CD-ROM drives bundled with HBAs are all notorious for having half-baked SCSI implementations, often without making their limitations apparent. Many of them only handle their own devices. Some of them only work with a limited selection of peripherals. Most have non-standard software so they cannot be easily updated to support new devices.

Once you are past the basic hardware issues involved in installing a SCSI-II HBA, the software games begin. If you have a SCSI-II HBA and a SCSI-II peripheral, you cannot necessarily get them to communicate. Corel has some interesting driver software that alleges to work with almost anything within DOS or Netware, but I really don't care for its Netware version. Micro Design International's SCSI Express is the best I've seen in a Netware environment and does fairly well in a DOS environment. Adaptec's ASPI software has become the major defacto standard for everyday use, but their recent modification that keeps their software from working with any non-Adaptec HBA may kill their inertia. Their software now stops dead in its tracks and complains about a "foreign adapter" if it sees other manufacturers' cards. This xenophobic attitude is not going to do much for Adaptec's public image.

This list of the potential hazards of SCSI is not meant to dissuade you from purchasing SCSI. SCSI is the best way to go for file servers or for individual computers needing optical drives or scanners. It's just to warn you to be careful.

As I write this, I am in the process of completing the integration of a single-user document and image management system using a SCSI-II HBA to access a 1GB (a gigabyte is 1024 megabytes) internal hard disk, an external 650MB multifunction (using both read/write and worm cartridges) optical drive, a 2GB internal tape drive, and a ten page per minute scanner.

This nightmare project has been plagued with a combination of factually incorrect documentation, ambiguous documentation, physically unreadable documentation, missing documentation, mismarked jumpers (on the circuit board), memory management conflicts, disk-compression utility problems, cache problems, print spooler problems, keyboard accelerator conflicts, video accelerator problems, bad advice from multiple technical support personnel, factually incorrect advice from multiple support personnel, multiple occurrences of SCSI driver software problems, and, finally, an unusual applications software bug that has been unreproduceable elsewhere. This integration time was non-billable.

All this after checking with all the manufacturers involved to get their confirmation the proposed system would work. Their answer was uniformly - "No problem" - riiiiiiight. Systems integration is not always fun.

NOMBDA/LANDA MERGER UPDATE

As a follow-up to my comments on the upcoming merger of NOMBDA, the National Office Machine Dealers Association, and LANDA, the Local Area Network Dealers Association (now awkwardly referred to as NOMBDA/LANDA), I found the following quote in a NOMBDA publication as part of an article promoting the merger:

"...and this merger of NOMBDA's copier industry expertise with the LANDA networking know-how represents a cross-pollination of dealers, customers, and vendors that will surely lead to the long term growth of both industries."

I think I can remember overhearing a conversation amongst other dedicated LAN resellers in which they were commiserating on how unsuccessful their businesses were due to their lack of in-house copier technology expertise. Hmmmm. I guess I must have wasted the last ten years working with LANs and spending many thousands of dollars and hundreds of hours on LAN classes while I was falling behind in that critical high-tech copier stuff. I wonder if there is a CCE (Certified Copier Engineer) rating that I could study for?

SFT III-PEAT, Part One

by Wayne M. Krakau - *Chicago Computer Guide*, August 1993

Wayne is a partner in Krakau Business Computer Systems, a systems integration firm that is Novell Gold and Lantastic Five Star Authorized. He has been working with computers for 16 years. He holds an MBA in Marketing and a BS in Information Science and was originally educated in Computer and Information Systems Analysis Engineering. He holds Novell CNI and Enterprise CNE ratings and all 5 of the original LANDA Certifications. Wayne can be reached for questions, comments, or topic suggestions at (708)298-7695 or through the Computer Guide.

SFT III is here. The ultimate level of Novell's System Fault Tolerance (hence SFT) has been established as a practical alternative for those companies with the highest of up-time requirements or those with demanding performance needs. This is first hand testimony, since I have just put up my first production Netware 3.11 SFT III system at a client site.

SFT Levels I and II include such now common features as Hot Fix (really Dynamic Bad Block Remapping), an intelligent UPS (uninterruptible power supply) interface, duplicate FATs (file allocation tables), duplicate DETs (directory entry tables), disk mirroring (two disks with the same information accessed via a single controller), and disk duplexing (mirroring using two separate controllers).

SFT Level III adds mirrored servers connected by a dedicated high-speed link called (naturally enough) the Mirrored Server Link or MSL. As a secondary disk of a mirrored pair of disks transparently takes over the duty of supplying disk services to its server upon failure of the primary drive, the secondary server in an SFT III pair takes over upon the failure of the primary system.

Unlike the mirrored disks, the mirrored servers are smart enough to switch identities dynamically as one or the other goes down, is repaired or reconfigured as necessary, and comes back up. For instance, if the primary server is taken down for maintenance or enhancement, the secondary server not only provides server resources to the users on the network, it reassigns itself as the primary server. When the former primary server is reawakened, it will recognize that a primary server already exists and will automatically configure itself as the secondary server. At that point, the primary server will recognize that the new secondary server has its disk drives out of synch with its drives and will issue an internal remirror command.

As with mirrored disks, this command causes the secondary server's disks to be updated so synchronization is reestablished, but only after synchronizing the memory within the two servers. Then the other (now primary) server could be brought down for its own maintenance or enhancement.

A requirement for SFT III is a pair of matched servers. Though it is technically possible to have a pair that is different in architecture, memory, or disk characteristics, a brief scan of the SFT III Forum of Netwire (Novell's section within the Compuserve Information System) will convince you to stick with absolutely identical systems. It seems that everyone in the world (Compuserve is accessible in many countries) who has tried unmatched systems is on Netwire begging for help, unable to configure a stable network.

I chose a pair of DTK 486/66DX2 EISA (Extended Industry Standard Architecture) tower systems with 1GB (one billion bytes where one byte is needed for each character). The disk drives communicated via Adaptec 1742 EISA busmastering 32-bit SCSI-II (Small Computer Systems Interface, Version Two, pronounced Scuzzy-Two) controllers (really host bus adapters or HBAs). I chose an identical machine to run regular Netware 3.11 for another department. The only difference was that the mirrored servers had thirty-two megabytes of memory while the single server was equipped with only sixteen megabytes. The extra memory was to accommodate both a larger department running more sophisticated applications, and the extra four megabytes of overhead required by SFT III.

The next requirement is for a pair of Mirrored Server Link (MSL) network interface cards (NICs). Only three companies make NICs approved for use as MSLs, Thomas-Conrad, Eagle, and Plaintree. They use three separate proprietary methods for passing data at 100mbs (megabits per second, or millions of bits per second where eight bits are needed for each character of information before overhead is added). I chose Thomas-Conrad's TCNS (Thomas-Conrad Networking System) for my client since it has been

around the longest, and we have had successful experience with that company in general and the TCNS system in particular in the past.

Another requirement is for an equal number of NICs in each server. A Netware 3.11 server can have up to sixteen NICs (if you can find a box with enough slots), and each NIC attaches to a separate network (a network, in Novell terminology, is defined by the cable plant). If a secondary server is to take over all the duties of the primary server, it needs to access all of the same wire segments in order to avoid cutting off some portion of the internetwork.

In this system there are three networks, one with a regular Netware 3.11 server, one with the SFT III system, and one providing a backbone between the two servers. (Note that the mirrored pair looks like a single logical server to the rest of the network, and is referred to as such in any discussion of the internetwork as a whole.) This meant that each of the three physical servers needed two network connections. The client chose the popular 10Base-T standard (Ethernet over unshielded twisted pair wiring) for their networking method.

Since a single Ethernet connection is a waste of the capabilities of an EISA slot, I chose Standard Microsystems Corporation's (SMC) 3032TP dual channel NIC 32-bit busmastering EISA NIC. It has two complete Ethernet chipsets with two separate sockets to more fully utilize EISA's speed without wasting an extra slot. One socket in each physical server would attach to the concentrator group associated with its own department while the other one would attach to the single concentrator assigned to the backbone between servers.

Next month, I'll cover the inner workings of Netware 3.11 SFT III.

SFT III-PEAT, Part Two

by Wayne M. Krakau - *Chicago Computer Guide*, September 1993

Wayne is a partner in Krakau Business Computer Systems, a systems integration firm that is Novell Gold and Lantastic Five Star Authorized. He has been working with computers for 16 years. He holds an MBA in Marketing and a BS in Information Science and was originally educated in Computer and Information Systems Analysis Engineering. He holds Novell CNI and Enterprise CNE ratings and all 5 of the original LANDA Certifications. Wayne can be reached for questions, comments, or topic suggestions at (708)298-7695 or through the Computer Guide.

This is a continuation of last month's column on Netware V3.11 SFT III (System Fault Tolerant Level Three), Novell's method of using duplicate servers to increase network reliability.

SFT III splits the Netware V3.11 NOS (Network Operating System) into two parts. The first part, the IO (Input/Output) Engine, contains all programs that require direct manipulation of hardware. Even NLMs (Netware Loadable Modules), Novell's technique for adding features directly to Netware, are covered by this rule. NLMs that control hardware, such as NIC (Network Interface Card) and disk drivers, must be loaded on the IO Engine.

Since each of the Mirrored Servers is a separate machine, with its own unique hardware, there are two IO Engines, one for each machine. These are named by the system installer just like any other Netware server. While there are no real default names, the documentation refers to the two servers as LEFT_IO and RIGHT_IO, with the designations arbitrarily assigned.

The installer can give them more descriptive names, based on their locations, if she or he wishes. That would make for less ambiguity when documenting procedures. It is easier to remember that a server called COMP_IO is in the computer room than remembering which way you should be facing to locate a cryptically named LEFT_IO. In spite of this suggestion, it is easier to use the common, documented names of LEFT_IO and RIGHT_IO in a general discussion of SFT III characteristics.

The other part of netware is called the MS Engine where MS stands for Mirrored Server. It is a logical, not physical entity, but it is named like any other server. The rest of the network sees the MS Engine as the lone representative of the SFT III mirrored servers. It runs concurrently on both of the mirrored servers and handles those tasks that are independent of the underlying hardware details. NLMs that monitor or govern the more global aspects of the network must be loaded on the MS Engine. Note that the use of the singular term MS Engine is proper, even though it runs on both physical servers. An NLM that is loaded on the MS Engine automatically runs on both physical machines as part of the SFT III structure.

Any NLM that runs on the IO Engine must be separately loaded on each physical server. In some cases, depending on the vendor's licensing policies, two copies of an NLM may need to be purchased to avoid piracy allegations. It is not always obvious which engine is appropriate for a particular NLM. When in doubt, remember you can always RTFM (Read the F..... Manual)!

Like NLMs, many SET commands and some other procedures are specific to either the IO Engine or the MS Engine. A few, like the TRACK ON/TRACK OFF diagnostic procedures, can be run in either IO or MS Engines. For instance, you can have three sets of Tracking screens running at once, one each for the LEFT_IO Engine, the RIGHT_IO Engine, and the MS Engine. You can then switch among them using the standard ALT-ESC or CTRL-ESC keystrokes.

Since the two physical servers operate together, you can even switch amongst all available screens from both physical servers and the logical MS Engine while working on just one server's console. All commands, including the vital LOAD and UNLOAD operators are available from that single console.

An interesting and valuable side-effect of the split of the Netware NOS into two pieces is that you can easily run it on a single dual-processor machine, ignore the warning messages regarding the lack of a "twin", and have the workload of Netware split between the two processors. Only a few machines have the appropriate drivers for this trick, but the performance is said to be awesome. In theory, this feat could be extended to use two mirrored servers, each with two processors, but I haven't seen any performance results using that combination -- yet.

When starting up an SFT III system, the first server to load the MS Engine using the new ACTIVATE SERVER command automatically becomes the Primary Server, while the second becomes the Secondary Server. The choice of which to start first, LEFT_IO or RIGHT_IO, is completely arbitrary, though a "Preferred Primary Server" (my own term, not Novell's) could be assigned to ease documentation and training. It is awkward to teach someone to flip a coin to decide which machine to start first -- better to tell them to try a specific machine first.

The two servers share the same name and identity on the network via the MS Engines. The MS Engines communicate via the two IO Engines, since the IO Engines do all the talking to physical devices such as NICs. The main communication is over the MSL (Mirrored Server Link) NICs. This link allows the MS Engine to keep the memory and disk data of the two separate machines completely in sync. Both servers also monitor the regular network traffic to make sure the other is still active.

If the Primary Server fails, either fully or partially, the Secondary Server will detect it via either the MSL or the regular NIC (or NICs). It will then take over as the Primary Server. When the former Primary Server is reactivated, it will discover that a Primary Server already exists and will take up the duties of the Secondary Server. First, it will synchronize its memory with the Primary Server's. Then, the server's will compare dates and times on their volumes and begin a remirroring process from the up-to-date server to the out-of-date server. When the remirroring process is complete, full redundancy is available and the system is ready to incur another fault.

The SFT III initialization process is governed by three pairs of files that logically correspond to the STARTUP.NCF and AUTOEXEC.NCF in regular Netware V3.11. Each IO Engine has its own IOSTART.NCF and IOAUTO.NCF, and the MS Engine has an MSSTART.NCF and an MSAUTO.NCF. The IOSTART files are executed first. Then the MSSTART and MSAUTO run. Finally, the two IOAUTO files are executed.

A small complication is that the execution of the IOAUTO.NCF files requires an active MS Engine with a mounted SYS Volume (the mandatory name for the first disk volume in a Netware server). When starting a collective SFT III from a complete stop (both physical servers shut down), the IOAUTO won't execute since an MS Engine and a SYS Volume aren't already up and running.

The solution for this is to use a little-known but incredibly useful feature of Netware V3.11 - batch files. A raw ASCII text file with valid console commands (LOAD, BIND, SET, etc.) can be executed from the colon prompt on the server's console just like a batch file can be executed at the DOS prompt. The filename extension must be "NCF". The file should be stored in SYS:SYSTEM (the SYSTEM directory of the volume SYS), since that is the default directory for console commands.

If you put a copy of the commands from the IOSTART.NCF file into a file while following these rules (I called mine START.NCF), you can manually get around the initial start problem. After the MS Engine is alive, and the SYS Volume is mounted, just use either CTRL-ESC or ALT-ESC to get to the IO Engine console screen. Then type the first part of the filename (eg. START) to execute your batch file.

When one of the two physical file servers goes down without a catastrophic hardware problem, Netware is automatically restarted without falling back to a DOS prompt or fully rebooting the server. As long as the other server stays up, the newly restarted server sees an active MS Engine and an accessible SYS Volume. This allows the IOAUTO.NCF for that server to execute automatically, without human intervention.

Next month, I'll cover how to make Netware 3.11 SFT III really work.

SFT III-PEAT, Part Three

by Wayne M. Krakau - *Chicago Computer Guide*, October 1993

Wayne is a partner in Krakau Business Computer Systems, a systems integration firm that is Novell Gold and Lantastic Five Star Authorized. He has been working with computers for 16 years. He holds an MBA in Marketing and a BS in Information Science and was originally educated in Computer and Information Systems Analysis Engineering. He holds Novell CNI and Enterprise CNE ratings and all 5 of the original LANDA Certifications. Wayne can be reached for questions, comments, or topic suggestions at (708)298-7695 or through the Computer Guide.

On with the continuing saga of Netware V3.11 SFT III, Novell's maximum reliability system. This article goes beyond theory into the real world.

The first practical aspect of the design process for an SFT III LAN is obtaining identical computers to use as mirrored servers. When I say identical, here, I really mean it. If the machines have even the slightest difference in memory, architecture, or BIOS (Basic Input/Output System) chips, the system might not work. The machines that I chose for my client's LAN, DTK 486/66 EISA Towers, were actually made side by side, one right after the other. I know this because I watched them being completed! That's how I made sure that they were identical.

This undocumented requirement for absolutely identical machines precludes the conversion of an existing single server or workstation into an SFT III mirrored server. Considering the rate at which undocumented changes and fixes are incorporated into motherboards and bios chips during the modern PC manufacturing process, I personally wouldn't bother trying to use a pair of machines if I could not confirm that they came off the assembly line at essentially the same time. Every vendor whom I have dealt with during this project has mentioned this undocumented restriction as being essential. Violating it would leave you with an unstable system that might work now, but will inherently be unsupportable in the future when, not if, it starts acting "funny".

This restriction applies even more stringently for the individual components inside the servers. There had better be an exact match between the specific network interface cards, the mirrored server link cards, the disk controller cards, the disk drives, and the video cards. Note that this advice contradicts Novell's. Novell feels that mismatched systems are viable, but the evidence that I have seen refutes that theory. While there are reports of working mismatched systems, I personally wouldn't want to get involved with such projects.

The software that makes these devices work, called drivers, must be the very latest versions. This means that each individual manufacturer must be contacted to obtain the latest driver, either from Compuserve or from its corporate bulletin board system (BBS). Merely browsing through a BBS is not enough. Many manufacturers hide the latest releases until they have been "wrung out" by chosen customers. You must call to find out the latest file name and, if necessary, the password needed to download it. Only after you obtain the latest edition of these drivers can you begin the process of configuring your SFT III system. (This is also a pretty good idea for even a plain vanilla Netware system.)

The cable between the two servers, called the Mirrored Server Link, or MSL, will be carrying data at 100Mbs (100 million bits of data per second). Because of this, it must be held to a much higher standard of workmanship than a standard 10Mbs Ethernet cable or a 16Mbs Token Ring Cable. If copper (as opposed to fiber) cable is used, for instance, you must be aware that the standard for how much bare wire can appear prior to a termination point has much less leeway than the standards for slower networking technologies.

Typically, this link is a direct machine to machine connection without an intervening concentrator. That means that one connector must be reverse-wired in a manner similar to a null modem. That is, the send and receive conductors are reversed on one end only. Be careful, and follow the manufacturers' recommendations precisely.

The most interesting suggestion that I have, is to ignore one critical point in the documentation. The documentation explicitly states that the ACTIVATE SERVER command should never be placed in either of the IOSTART.NCF initialization files. On Netwire and via 1-800-NETWARE, I was advised to

arbitrarily choose one server and put the ACTIVATE SERVER command only in its IOSTART.NCF file. That is a major contradiction. My advice to you is - with careful planning and proper training - to place an ACTIVATE SERVER command in both IOSTART.NCF files!

The documentation was written with the idea that an embedded ACTIVATE SERVER command would cause a server that was supposed to come up as a secondary, to initialize instead as another primary server with exactly the same identity as the first primary. This would cause the entire internet (all directly or indirectly connected servers and workstations) to crash. The problem is, without automating the ACTIVATE SERVER command, the failure of both servers due to a sustained power failure would cause the system to stay down until supervisory personnel manually typed "ACTIVATE SERVER". The documentation assumes that both servers would never go down at the same time! I guess there are no sustained power outages in Utah (Novell's home).

The technical support people on Netwire and 1-800-NETWARE seem to have less fear of the possibility of duelling primary servers. However, they make assumptions about the predictability of future errors, believing that the system will react differently depending upon which physical server crashes. Again, the continuous availability of technical personnel, needed to restart the system manually, is assumed. That last point is my interpretation, since I can't believe that any company with data valuable enough to protect with SFT III would allow non-technical staff to get near their servers, much less lay hands upon them. The concept of a LAN running completely unattended (quite a common occurrence) seems utterly beyond the comprehension of the documentation writers and the technical support people.

Here is my plan. Adjust the allowed uptime and the recovery time parameters in your UPS (uninterruptible Power Supply) control software so that one server always goes down and later comes back up first. (You ARE using UPS's on all of your servers, aren't you?) This machine becomes your Preferred Primary Server (my term, not Novell's). The idea here, is to prevent the two servers from starting at about the same time, within a narrow range.

If they do start within a short period (about one minute when using 1G drives) both will come up as primary servers, crashing the internet. If, however, one is already up as the primary when the second wakes up, the ACTIVATE SERVER command will automatically abort and that machine will come up as the secondary server! This effect is dependent upon the existence of a valid mirrored server link, so if your link is defective in any way, the servers won't see each other and mayhem will result.

In addition to the adjustment of UPS parameters, procedural and training issues are raised. You must fully document the proper procedure for bringing up the system from scratch (when both servers are deactivated) and train your personnel appropriately. The first step is to start up the arbitrarily chosen Preferred Primary Server. Only after that server is completely awake, with all NLMs (Netware Loadable Modules) loaded, is it altogether safe to start the Preferred Secondary Server. This procedure eliminates the possibility of duelling primary servers.

Implementing my suggestion will result in having a system that will restart by itself under almost any conditions, a situation considerably more fault tolerant than that obtained by using either documented or the Novell-suggested procedures.

One final warning about Netware V3.11 SFT III is in order. It is far closer to the bleeding edge of technology than Novell would like to admit. The system that I have configured, and some others mentioned on Netwire, still don't work. The primary server has been chugging along happily for over two months, but the secondary server hasn't talked to it in weeks. The vendors involved are cooperating in finding the problem via Novell's Technical Support Alliance, but nothing has worked yet. I will publish the final results of this debugging effort in future columns. Meanwhile, I expect the real world to catch up to the theoretical soon, providing a premium fault tolerant option for critical systems.

THE LONG AND SHORT OF HORROR STORIES

by Wayne M. Krakau - *Chicago Computer Guide*, November 1993

Wayne is a partner in Krakau Business Computer Systems, a systems integration firm that is Novell Gold and Lantastic Five Star Authorized. He has been working with computers for 16 years. He holds an MBA in Marketing and a BS in Information Science and was originally educated in Computer and Information Systems Analysis Engineering. He holds Novell CNI and Enterprise CNE ratings and all 5 of the original LANDA Certifications. Wayne can be reached for questions, comments, or topic suggestions at (708)298-7695 or through the Computer Guide.

You've all heard about them. You've probably read about them quite frequently (including in the pages of this publication). They are downsizing and client-server horror stories.

The story lines usually run along similar tracks. A large firm with a well-run mainframe (or minicomputer) installation needs to cut costs or reduce the typical two-year backlog of mainframe applications. The potential solution - downsizing - using PCs and networks to replace functions previously performed by larger computers.

These days, the use of client server technology is often a major part of the solution. This involves separating the information requesting and displaying functions of a database from the searching and supplying functions. The workstation retains only the requesting and displaying functions while the server takes over searching and supplying.

(Note to techies: Stop shaking your heads and clucking your tongues - this simplistic definition is good enough for the sake of this discussion.)

Sometimes outside "consultants" are called in. Other times the in-house staff tries to "wing-it" (usually good for a few laughs).

After the transition to PCs and LANs is made, all hell breaks loose. Those nasty ol' PCs and associated client-server databases fail the company in a big way. The applications don't work right. The computers break down. Users have nervous breakdowns. Thousands (millions?) of dollars are lost. Dynasties fall. Peace talks fail. Talk shows are canceled. (You get the picture.)

What you don't see in these articles are the facts on the conditions prior to the downsizing adventure. The most important overriding fact: Implementing any new technology is a potential managerial nightmare. This applies to new mainframe and minicomputer systems (both hardware and software) just as much as it applies to PCs and networks.

Many mainframe computers crash (suffer a complete stoppage) on a regular basis. Mainframe users often state that they expect this! It has been this way for so long - since long before PCs existed - that they think this is standard operating procedure!

Also, mainframe applications are often so unreliable that I have encountered users who would actually pre-calculate the results of equations by hand so that they wouldn't accidentally enter a combination of data that would lock up an individual application, subsequently corrupting the underlying database. Convincing them that they can safely enter data at full speed after redesigning an application to work safely can be very difficult. They are so used to mistrusting computer people that they simply won't believe that a reliable application can be created.

It is not uncommon for mainframe programming projects to progress so slowly that they are declared obsolete and are canceled prior to completion - with up to multimillion dollar writeoffs. That, of course, assumes that the programs would have filled the users' needs in the first place. Many won't. Users are normally the lowest personnel on the programming design totem pole, so their opinions and true business needs get filtered out of the equation.

This is not to say that the data processing professionals involved are incompetent or that they just don't care. In fact, my personal belief is that a prime reason for the high turnover rate among data processing staff is their emotional and ethical dissatisfaction over being involved in inadequate and buggy computer systems.

The problem is that large, complicated computer projects (regardless of computer size) require teams of people from diverse backgrounds - users, applications programmers, systems programmers,

systems analysts, business analysts, assorted managers, etc. - to work in concert with limited resources. These include not enough personnel, inadequate training, unrealistic budgets, and - my own favorite - artificial deadlines made up to fulfill political, as opposed to business, needs. These factors work together to chip away at the reliability and practicality of computer systems.

Another factor is the overall size of the team. Research into software engineering has shown that problems in large computer projects are usually "solved" by throwing "warm bodies" (I am not making this term up - it's really used!) at the project. This same research has also shown, conclusively, that the more people involved in a project, the lower the productivity, accuracy, useability, and reliability of the project. There is even a point of diminishing returns where the project becomes impossible to complete!

The task at hand in this article is not to find a cure for these problems. Successful projects do occur, so we know that the problems are inherently curable. The real task is to step back and look at the reports of failed downsizing and figure out what the track record was before the downsizing took place. How effective was the company involved when it did its last mainframe (or mini) upgrade or system redesign? If they had problems in a field where they had major in-house expertise, they could easily have problems with a technology in which they had little or no experience.

Also, find out why the company is downsizing. Often, the reason is that the existing system is an unmanageable mess. Again, unless a major methodology change is undertaken, future projects, whether downsizing or not, are likely to end up in even worse shape.

THANKS FOR THE MEMORY

by Wayne M. Krakau - *Chicago Computer Guide*, December 1993

Wayne is a partner in Krakau Business Computer Systems, a systems integration firm that is Novell Gold and Lantastic Five Star Authorized. He has been working with computers for 17 years. He holds an MBA in Marketing and a BS in Information Science and was originally educated in Computer and Information Systems Analysis Engineering. He holds Novell CNI and Enterprise CNE ratings and all 5 of the original LANDA Certifications. Wayne can be reached for questions, comments, or topic suggestions at (708)298-7695 or through the Computer Guide.

Having memory problems? No -- not the kind they talk about in those thirty minute infomercials. Your computer's memory is at issue here. If RAM Cram is your disease, then Netroom is your cure. (Hmmmm. Perchance one too many Stallone commercials?)

Historically, the RAM (random access memory) manager market has been dominated by QEMM from Quarterdeck and 386MAX from Qualitas (not necessarily in that order). Helix's (Long Island City, NY, 718-392-3100) Netroom (originally called Headroom in the single-user version) was considered a close, but mostly unknown, third.

While QEMM and 386MAX alternated in holding the throne of the most efficient and effective memory manager, Netroom quietly kept up with their latest RAM-squeezing tricks. Usually, it was disqualified from the top spot in reviews not due to any technical deficiencies, but rather for a lack of sophistication in its original user interface.

My first encounters with Netroom were on LANs. The people at Helix discovered, early on, that the local area network market was potentially the most in need of memory management (hence the product name). They quickly grabbed the technological lead in manipulating the workstation software from Novell's Netware.

They followed up by being the first, and until recently the only, RAM manager with severely discounted network licenses. Their competitors required users to purchase a full single user copy for each workstation while Helix was selling ten-packs and fifty-packs and had special deals for higher quantities. It was this discounting that led me to sell Netroom to clients who otherwise could not afford serious memory management software.

The criticism of Netroom's old interface was a moot point for me. The reviews assumed that inexperienced civilians would run memory management software on complicated networks. As systems integrators, we take care of installing and fine-tuning systems before our clients see them. This negated the criticism, since the interface was easily learned by those with heavy manual memory management experience.

I still remember one of the first machines on which I installed Netroom two years ago. It was a dedicated fax server running Optus Facsys LAN-based fax software. The system had three Intel Satisfaxion cards, each with multiple CONFIG.SYS and AUTOEXEC.BAT entries. These drivers were in addition to the standard Netware Shell software. After running Netroom and hand tweaking the system, a CHKDSK command revealed over 710K total RAM with 638K available!

Until recently, I continued to recommend Netroom's competitors (leaning somewhat in favor of QEMM) for non-network and small network situations. Now that has changed. The latest version of Netroom supports a user interface that is arguably better than the competition. On top of that, it has leap-frogged over them to provide the most technically elegant and most effective RAM management software available.

"Cloaking" is Helix's word for its latest RAM-cramming trick - one that is inherently safer and less prone to mysterious compatibility problems than its rivals methods. Helix has licensed both a 32-bit system BIOS (Basic Input Output System) and a 32-bit video BIOS from Award Software Incorporated, a prominent BIOS developer, and included it within Netroom. With Netroom, you can effectively replace your system's BIOS chips on the fly during its boot-up sequence with the latest advanced BIOS. The BIOS occupies only 8K of system memory within the base 640K versus a more typical 96K adding the speed of 32-bit access and an API (Application Programmers' Interface) as well. The speed difference is

immediately humanly (as opposed to just performance testing software) detectable, even on a 486/66DX2 machine with a fast video card!

The API allows programmers to write utilities to use the full power of the 32-bit BIOS while occupying little or no RAM within the base 640K. Helix packages some utilities including a disk cache with Netroom that use this API with outstanding results.

Two of my recent implementations of Netroom demonstrate its effectiveness. In the first, I was configuring a LAN using all 486/66DX2 workstations running Microsoft Windows. Using Netroom's DISCOVER program to explore the characteristics of the computers and its CUSTOMIZE program to obtain optimum configurations, I was able to get an average 629K free. That was remarkable considering that network, anti-virus, LAN management, security, LAN inventory, and caching software were all being run on these machines.

In the second situation, a single user document and image management system was being run on a machine that was also a workstation on a LAN. This application software required approximately 580K to run with all of its features active and reliable. In addition to the network workstation software, this computer had special drivers for a SCSI-II (Small Computer Systems Interface, Version 2) card running both an external multifunction optical disk and a scanner. On top of that, it had drivers for a separate SCSI controller for a two gigabyte tape drive.

The original memory available was 475K. My initial attempt at improving that was with QEMM. After much hand tweaking, I finally managed to get 524K available memory. That wasn't enough. I switched to Netroom. With some more manual manipulation, I got up to 578K! Now, that's some improvement!

I was not yet satisfied, so I upgraded the system to DOS V6.0. Netroom's manual stated that it would automatically recognize DOS V6 and would be more efficient working with it. This effort garnered another 6K. The system now had 584K - more than the minimum required to run the applications software!

After testing the system for compatibility problems, I found that while all normal applications worked, some DOS utilities, like MODE and TREE, locked the system up with bizarre memory problems. Luckily, DOS V6.2 was released just as I came to this impasse.

If you own a copy of DOS V6.0, you are automatically licensed to upgrade it to V6.2. This client had just upgraded several machines from very old versions of DOS up to V6.2, so the diskettes were easily at hand.

This upgrade to DOS V6.2 successfully eliminated all compatibility problems. It also followed the rule about not trusting products with a version number having all zeros after the decimal point.

An interesting aspect of this last system was that, after the failure of my experiments with QEMM, I had been warned by the manufacturer of the optical disk drive that their drivers could NEVER be loaded high by memory management software. Oh well, never say "never".

THREE FOR THE ROAD

by Wayne M. Krakau - *Chicago Computer Guide*, January 1994

Wayne is a partner in Krakau Business Computer Systems, a systems integration firm that is Novell Gold and Lantastic Five Star Authorized. He has been working with computers for 17 years. He holds an MBA in Marketing and a BS in Information Science and was originally educated in Computer and Information Systems Analysis Engineering. He holds Novell CNI and Enterprise CNE ratings and all 5 of the original LANDA Certifications. Wayne can be reached for questions, comments, or topic suggestions at (708)298-7695 or through the Computer Guide.

SFT III UPDATE

Finally -- it actually works! My first Novell Netware SFT III has been running for ten days as I write this. The two servers are talking to each other like reunited twins who were separated at birth. As advertised, either server can be downed (the "DOWN" command is used to shut down a Netware server so this verb is commonly used) without hurting the network. The users don't even know that it's happening. They only notice when the downed machine is restarted because the network slows while the restarted computer is updated.

While this is convenient for maintenance and upgrades, the real value is protecting the network from server crashes. As I suggested to the client, you could take an ax to one of the servers, and the LAN would keep on ticking. Try that with your Timex!

There was a downside to this project. It took an incredible number of non-billable hours of research and experimentation to get the system working, but I sold the system, and I told the client it would work, so I had to make good on my promises. (Having ethics can be quite expensive!) Working in conjunction with the vendors, we went through two pairs of motherboards, three pairs of MSL (Mirrored Server Link) network cards, five separate sets of MSL cabling, and enough driver software versions (for talking to various hardware items) to fill a CD-ROM. That's all in addition to hundreds of dollars in Compuserve bills for tech support and software downloads.

The implied question is whether I would suggest SFT III for others. Yes, I would. As I have debugged this installation, others, as observed on Compuserve, have been successfully debugging other SFT III systems around the world. The technology isn't perfect, but it has become reasonably stable.

THE NAME GAME

During a recent LAN upgrade (new Netware, new server, new workstations), I noticed that the workstation that was doing network backups to an internal tape drive was hung up with an error message indicating that it had encountered an error accessing COM3. I restarted the backup and let it run while we configured the new server and workstations.

When I returned to that machine, it had again hung up with the same message. A reboot gave a message warning of major motherboard problems. Oh well, the computer was being replaced on that day, so I just chalked the error up to the final screw-up of a dying computer and went back to the new systems.

We left another computer running overnight to transfer data from the old server to the new server with the NCOPY command. NCOPY was used since it is a miniature client/server implementation of a COPY command. It lets the servers do the work, leaving the workstation to just display the results.

The next morning, we found the NCOPYing computer stuck, but not hung up, with a message about an error reading from COM3! We couldn't believe that two machines could fail with the same obscure error. This time, however, we could see the information on the last few files copied that remained on the screen. The last file name in progress was called -- you guessed it -- COM3! A simple response of "I" for "Ignore" was all that was needed to allow NCOPY to finish.

The reason for the errors was that someone had inadvertently named the third correspondence for a particular project COM3. That is one of the "magic" words that DOS reserves for its own use. Netware

doesn't care about it, so it allows that file name to be written. It is only when some internal DOS process tries to access this reserved name that errors occur.

Because of this incident, I thought I'd pass along a warning. DOS reserves the following names for its own use: AUX, for the auxiliary port; COM1, for communications port number one; COM2, for communications port number two; COM3, for communications port number three; COM4, for communications port number four; CON, for the main console, that being the screen and keyboard; PRN, a synonym for printer port number one; LPT1, for printer port number one; LPT2, for printer port number two; and LPT3, for printer port number three. These names should not be used for the first name of a file. Using a last name with them won't help. DOS still recognizes the first name and ignores the last.

An interesting trick to solve problems encountered with Microsoft Windows when using print spooler software (programs that rapidly save printouts to disk or in memory and feed it to the printer only as fast as it can accept it) takes advantage of these reserved names. You can disable printing directly to hardware and then redirect the printout within Windows to a file name like LPT1.DOS. DOS will ignore the last name and feed the output to the logical printer port. Since most print spooler software works by intercepting data sent to the logical printer port via standard DOS services, this allows Windows to cooperate with this type of program.

EXTRA!, READ ALL ABOUT IT

Attachmate Corporation (800-426-6283), makers of mainframe and minicomputer terminal emulation software for IBM environments, and inheritors of Novell's own products in that arena, have just released the beginnings of a series of products designed to make life much easier for people trying to automate IBM mainframe and MINI shops.

Attachmate makes the EXTRA! series of Windows-based emulators. Now they have created products that allow selected other programs to tap directly into their emulators and manipulate them as needed to automate tasks. One is for Microsoft Office. Office includes Microsoft's Word, Excel, Powerpoint, Access, and Mail products for Windows. Attachmate's EXTRA! Tools for Microsoft Office contains programs to allow easy, automatable transfer of data to and from these products and their Extra! emulators.

A product with even more interesting possibilities is Attachmate's EXTRA! Tools for Visual Basic. Microsoft's Visual Basic already contains features that make it one of the top tools for integrating and automating multiple products within the Windows environment. These Tools add the ability to tap right into all of the abilities of EXTRA! and tying them into almost any other Windows application. This is an amazing breakthrough for systems integration. This product can potentially eliminate many thousands of hours of manual programming for those trying to improve productivity and ease of use in IBM minicomputer and mainframe installations.

URGE OVERKILL

by Wayne M. Krakau - *Chicago Computer Guide*, February 1994 - *The Law Works*, August 1994

Wayne is a partner in Krakau Business Computer Systems, a systems integration firm that is Novell Gold and Lantastic Five Star Authorized. He has been working with computers for 17 years. He holds an MBA in Marketing and a BS in Information Science and was originally educated in Computer and Information Systems Analysis Engineering. He holds Novell CNI and Enterprise CNE ratings and all 5 of the original LANDA Certifications. Wayne can be reached for questions, comments, or topic suggestions at (708)298-7695 or through the Computer Guide.

It's happening again. People have such a great fear of falling behind technologically, that they are willing to delay purchases or make illogical buying decisions in a vain attempt to protect themselves.

This is most often seen with people who were caught in the switch between 8088/8086 (XT-class) and 80286 (AT-class) systems or in the demise of the 80286 in favor of the 80386 chip. They were stuck with a generation of chips that couldn't run the latest versions of software! Some of them went through both generational cycles. Imagine being forced to scrap thousands of dollars worth of precious XT's (Bought for $5,000 a pop!) and purchasing a load of hot new AT's, only to find that software is being written to demand an 80386 chip or higher. Now that's traumatic.

These people are now demanding the hot (literally) new Pentium-based computers. After negotiating for computers appropriate for a particular client's needs (and budget) I have been getting blindsided by sudden changes in the clients attitudes. Suddenly, the clients are desperate to get their hands on a Pentium, even though I previously had to twist their arms to get them to specify a 486/33SX instead of a 486/25SX.!

In order to pay for these systems, they start cutting back on other parts of their systems. Usually backup systems are the first to be scaled back. Less tapes and lower capacity tape drives start looking attractive. Reliability features such as RAID (Redundant Array of Inexpensive Disks), UPS's (Uninteruptible Power Supplies), and reliable servers are usually next on their hit list. RAM, quick video cards, fast hard disks, and fast or reliable network systems (network cards, concentrators, even the cable plant) are next - anything to get a faster CPU (central processing unit). I guess the Intel propaganda machine is really working.

Before directly commenting on the facts, I would like you to keep in mind that for some time now, I have considered it inappropriate to sell 80386 machines due to their lack of upgrade options. I've even declined opportunities in which clients demanded systems too weak to do their jobs. This, in spite of the fact that faster computers are more profitable to sell.

Now, on to the facts. When making purchasing decisions, first base your choices on fully optimized systems. It doesn't help to compare systems that are crippled by ignorant or lax installers. Running installation programs alone just doesn't cut it. I regularly receive calls from network users on the day after an optimization, asking if I had given them a replacement server since their system now runs faster.

It isn't magic, though it often involves some not-so-secret incantations and gestures, none of which can be repeated in this publication. It's just a fact of life that installation programs are meant for the lowest common denominator. That, combined with inept installers, makes many systems run even less efficiently than the way that standard installation programs left them.

After optimization, performance considerations get easier to track. In general, the "weakest link" theory applies. The key is to figure out the primary limiting factor - the weakest link.

In a server, memory is the critical component. Quantity of memory is followed by access time. That includes the speed of the chips, the size of the processor cache, and the underlying speed of the motherboard. That's why servers with 50MHz 80486/50DX (50MHz CPU and 50MHz motherboard) frequently outperform an equivalent server with an 80486/66DX2 (66MHz CPU and a 33MHz motherboard).

Next in line is disk access. Disk interface, speed of the disk, efficiency of the controller, and the bus interface of the controller all go into that equation.

After that, the network becomes the limiting factor. The network type, network card speed, and bus interface control this item.

Finally, the CPU speed comes into play. While there are situations where database server software or other NLMs (Netware Loadable Modules) change this equation, most servers are not CPU bound (That's tech-talk for a weak link). Only after the other factors, especially memory, are satisfied, does raw CPU power matter.

Windows machines parallel servers in their hunger for memory. More memory for them means extra speed and a lot of extra flexibility. You can run more simultaneous applications if you feed Windows more RAM (Random Access Memory).

The other major consideration for Windows is video performance. Video speed is improved by specialized video processor chips, faster or more sophisticated video memory chips (VRAM), and faster card bus interfaces. Video is such an important factor that a slow CPU combined with fast video can outrun a fast CPU with mediocre video performance.

In plain DOS machines, memory isn't as critical, but a couple of megabytes of left-over memory for caching and memory management sure helps. The demand for video speed is entirely dependent on the specific application involved. AutoCAD has completely different requirements from a character-based accounting system.

Once you have winnowed the choices down to the CPU, you have to be careful about your choices. Make sure that the 80486 chip really is one. An 80386 compatible with a cache added is not an 80486 compatible! If it is not an Intel chip, then ask for confirmation in writing that the "80486" you are getting will run the 80486 instruction set.

As for Pentiums, I am waiting for the low-voltage, cool running version before I suggest them for general use. I only use them if my client understands the potential drawbacks. This chip runs so hot that in initial tests for several manufacturers, it did what I call a "China Syndrome". That is, it melted right through the motherboard! With an added heat sink and more numerous fans, the motherboard will only just survive.

The problem is that heat is the ultimate enemy for motherboards. I wouldn't want to guarantee the long term viability of a motherboard equipped with an air-cooled Pentium. I am, however, intrigued by some new gadgets that provide for liquid cooling - like a Cray supercomputer! A cooling unit sits on the Pentium with a small tube snaking outside the computer to a small radiator unit. This radiator is supposed to be placed on top of your monitor. It's a clever solution, but a pain to deal with as well as a possible hazard. I wonder how an insurance company would react to a claim on a microcomputer that leaked!

Another concern is that current software is far behind hardware. Only Netware 4.x takes advantage of the full power of the Pentium. Only Netware 3.x and 4.x (though I've heard rumors about a couple of UNIX dialects) take full advantage of the 80486 chip. The rest of the industry is just now catching up with the 80386 chip. The theoretical advantages of an advanced CPU chip don't mean a lot when it is simply emulating a more primitive predecessor.

The final decision on CPUs has to take into consideration the cost effectiveness of the incremental increase in processing power that a more powerful CPU buys. What can be overkill for simple DOS-based word-processing might be just the thing, reliability questions or not, for CPU-intensive engineering or financial analysis tasks.

For those of you who don't follow rock music trends, the title of this article is also the name of one of the bands (along with "Smashing Pumpkins") that has brought Chicago into the limelight as the center of rock music. Now you know why the publisher originally wanted me to name my column "Wayne's World" instead of "Putting It All Together". Party on.

TRANSITIONS

by Wayne M. Krakau - *Chicago Computer Guide* March, 1994

Wayne is a partner in Krakau Business Computer Systems, a systems integration firm that is Novell Gold and Lantastic Five Star Authorized. He has been working with computers for 17 years. He holds an MBA in Marketing and a BS in Information Science and was originally educated in Computer and Information Systems Analysis Engineering. He holds Novell CNI and Enterprise CNE ratings and all 5 of the original LANDA Certifications. Wayne can be reached for questions, comments, or topic suggestions at (708)298-7695 or through the Computer Guide.

It's not surprising that civilians and inexperienced techies might be lulled into a sense of complacency on making major network upgrades. It's really shocking when resellers and "consultants", and in-house big system techies are negligent.

For resellers, "consultants", and the like, it's an opportunity lost. Major upgrades take planning, preparation, training, and testing. Except for that rare customer that has 100% of the needed expertise and spare personnel hours to boot, the extra labor and knowledge could be supplied by these outsiders - all for a profit. Either they lack knowledge of the real world or they feel that they can make more money cleaning up the mess after the fact rather than doing the job right in the first place.

For in-house personnel, a major upgrade is an opportunity to justify their existence, improve their status, and rack up points toward their next raise or promotion. These days, it might even provide insurance against the next series of layoffs. Or, it can provide a one-way ticket to the unemployment line.

As an alternative to the unemployment line, you can always argue that "small" computers are unreliable and not worthy to run corporate applications. Then you can get publicity by telling your horror story to reporters for mainframe and minicomputer oriented publications. This from people who normally need several approval signatures of their testing methods to put a program into production that had only a single line of code changed!

The latest series of upgrades that I've seen have been to Netware 3.12 and, to a lesser extent, Netware 4.0. After having one company pull the plug on Netware 4.0, and personally undoing the workstation portion of two upgrades to Netware 3.12, I have gained some insight.

Netware 4.0 was a very rough-edged product. It had problems both within the software and with the documentation. Netware 4.01 is much more stable and comes with improved documentation. While it doesn't follow through with all of the promises associated with Netware 4.x (you will have to wait for version 4.1 for a major improvement there), it is a viable product.

I believe that much of the bad reputation that hounds 4.01 derive from associated products, not the operating system itself. I have observed storms (hurricanes?) of protests in the Btrieve section within Netwire, Novell's own forum within Compuserve, the on-line information service. The protests center on problems running applications developed using old (5.x) versions of Btrieve with new (6.x) versions. I have also discussed these problems with developers within applications software companies. The consensus is that the changes in Btrieve along with some early instabilities caused many programs to fail.

If your applications failed with a new version of Btrieve, and you used Netware 3.x, all you had to do was to switch back to an old version until the compiler publisher modified their compiler (a program to translate from human readable language - or at least computer geek readable - into computer readable commands) and then the applications software vendor released a recompiled version of their product.

If, on the other hand your applications failed after you had upgraded to Netware 4.x, you've got big problems. Old versions of Btrieve aren't compatible! Your only choice is to revert to Netware 3.x. Now you know why I emphasized testing! Either that or getting your resume writing skills up to snuff.

The other related product is the new replacement shell. The shell is a generic name for the program that allows a workstation to communicate with the network. To accommodate all of the new features of Netware 4.0. a completely new shell concept was developed called the VLM, or Virtual

Loadable Module. Like the NLMs (Netware Loadable Modules) that make Netware 3.x and 4.x work, VLMs are loaded piecemeal as needed to make the shell work.

This new shell uses radically different methods to communicate with DOS. These methods are so different that VLMs include an emulation feature to simulate the older ODI (Open Datalink Interface) shell for those programs that need it. The problem is that this emulation is not perfect. Some programs just won't work with VLMs.

There is no easy way to tell if a given program will run. Calls to software vendors are one way to find out. Testing is the other. Remember, though, that Netware 4.x absolutely demands VLMs to access all of its advanced features. If your program proves unworkable under VLMs, you have no choice but to either switch to the bindery emulation mode in which Netware 4.x pretends to be 3.x with only 3.x's capabilities, or going back to real 3.x.

If you are running Netware 3.x, you can always go back to the ODI shell. If you are really a masochist, you can even go back to the ancient IPX/NETX shell. Ah, but here it gets complicated. Novell declared the IPX/NETX shell obsolete as of November 1991. At that time the ODI shell had been around quite a while and was mature, stable, and compatible with virtually all applications.

Now, Novell has declared the ODI shell obsolete. If you install Netware 3.12, the documentation tells you to install VLMs, only mentioning the ODI shell as an afterthought. If you didn't know any better, and religiously followed the directions, you could get hit with errors that are bizarre and difficult to track.

For example, applications written in Micro-Focus COBOL, as I write this, are not compatible with VLMs. Many accounting systems are written in this language. It is also very popular for downsizing projects. This is not the company's fault. They are just one of many caught in this trap. To run with VLMs (and by implication with a full-featured implementation of Netware 4.x), first you will have to wait for Micro-Focus to change their compiler. Then you will have to wait for the software developer or your in-house programming staff to recompile the programs. Finally, you will be able to run them.

Novell jumped the gun on this one. It's obvious that jumping into VLMs without planning and testing is dangerous. If you are upgrading to Netware 3.12, in particular, I suggest you don't use them until you either can prove that all of your programs will work, or you absolutely have to because you acquired a new application that absolutely requires VLMs.

Alternatively, I've got a great investment for your layoff settlement money. It's a 900 number for computers that are too shy to communicate with other computers over a network. You have your computer link up via modem with our computer and the two have a private one-on-one private conversation. In no time at all, your computer will have the confidence to communicate freely with its peers. We need start-up funds to place ads on all of the online services like Prodigy, Compuserve and especially the Internet. Note that a multiple-computer 900 line will also be available for those computers with "special" needs.

HOT-WIRED, Part One

by Wayne M. Krakau - *Chicago Computer Guide*, April 1994 - *NewsWare*, January 1996

Wayne is a partner in Krakau Business Computer Systems, a systems integration firm that is Novell Gold and Lantastic Five Star Authorized. He has been working with computers for 17 years. He holds an MBA in Marketing and a BS in Information Science and was originally educated in Computer and Information Systems Analysis Engineering. He holds Novell CNI and Enterprise CNE ratings and all 5 of the original LANDA Certifications. Wayne can be reached for questions, comments, or topic suggestions at (708)298-7695 or through the Computer Guide.

They are so common and so transparent to our daily lives that we ignore them. But, they are so vital to the functioning of our systems that even the smallest problem can disrupt an entire network. They are wiring systems, more properly known as cable plants. Here are some real gems that I have encountered.

I was called in by another LAN reseller to take a look at a 10Base-T (Ethernet over unshielded twisted pair cable) LAN at his client's office. They needed to get several new workstations connected to the LAN and have a server-based tape drive installed on their new network. In addition, they were having problems keeping the only existing workstation connected to the network.

The existing workstation was first. I checked the shell software (that talks to the network) and updated it to the latest version. Then I checked the configuration of the NIC (network interface card). It's always a good idea to see if a NIC is taking up some resource (IRQ, I/O address, high RAM address, DMA) that another device is using. This is one of the most common ways to get a network to act "funny". Reconfiguring the NIC and running tests against it didn't help.

Since the client was a telephone equipment and cabling company, and I was nearly surrounded by professional cablers, I didn't initially check the cabling. When the workstation kept losing its network connection in the middle of testing the NIC, I decided that cabling could be an issue. That's when I found out that the jumpers (terminated cables that go between things) that connected the workstations to the network sockets in the wall were made from what's commonly known as silver satin cable. In English, that's the same wire that goes from a phone to the wall socket in homes! It is flat - that is it has no twists. The IEEE 1
bring chaos to an entire network. This place had a silver satin jumper at every workstation! No wonder they couldn't keep their first workstation connected.

When I explained the standard to the manager (a non-techie), he reacted rather strangely. This made me suspicious, so I asked about the wire in the walls. He explained that since unshielded cable was good, his people had told him that SHIELDED cable would be even better! I re-explained the standard, but I could tell he didn't really believe me. All of his cabling people couldn't be wrong!

Then I asked to see their wiring closets. I confirmed that the wire coming out of the walls was shielded twisted pair (STP), not UTP. I also found that the cables going from the punch-down blocks (where raw wires are connected) to the patch panels (a kind of intermediate box of sockets similar in principle to an old-fashioned phone switchboard) and from the patch panels to the concentrators were also silver satin. Just for bonus points, the panels used the center four pins of RJ-12 6-conductor home-phone style sockets rather than the standard specified 8-conductor RJ-45 sockets.

Between the incorrect wire and the mismatching impedance caused by mixing different types of wire, a linked set of rusty coat hangers could probably transmit a better signal! No wonder they had problems.

I then attempted to get several more workstations connected to the network, but was unable to get more than three up at once. The cable plant just couldn't handle any more. Luckily, the tape installation worked perfectly, so I had something to validate my competence.

In an additional effort to establish that I wasn't making up standards, I called my cabling authority, Bruce Kahn of Telnet Communications Consulting, Inc. (708-215-0003), a local cabling

company, and confirmed my findings. Finally, I faxed the client a list of Ethernet specifications that I had compiled for the seminars that I occasionally teach.

Even before I left, however, I got some indication that I was convincing management that their cable plant had problems. I heard the distinctive sounds of one of the technical managers being interrogated - very loudly. Eventually, the company completely rewired the whole place. That solved their problems.

In another situation, I was called in directly by the system administrator of a LAN that occupied one floor of a fairly new high-rise building. We got together to discuss his expansion plans for the network. The client needed more workstations on their 10Base-T LAN, some of them MACs. In addition, the administrator was aware that due to his lack of training, the LAN was probably not optimized. Finally, he told me that he had some problems keeping workstations attached to the network.

The first thing that I noticed was that his recent Netware upgrade was to Netware V2.2, not the then current Netware V3.11. As I have stated before in this column, I consider that a dirty trick and grossly unethical. Resellers will dangle the small difference in price between the two versions in front of clients, but neglect to tell them that V2.2 will cost them thousands extra in initial installation costs and ongoing maintenance. On top of that, it is much slower and much more difficult for the system administrator to manage on a day-to-day basis.

Next, I noticed that they were using the old NETX/IPX twins rather than the ODI shell. The NETX/IPX combo was declared dead (and rightfully so) by Novell as of November of 1991, so I suggested that an upgrade was in order.

Then I ran TXD, Thomas-Conrad Corporation's (800-332-8683) diagnostic program. Just running the diagnostic program was enough to blow the workstation right off the network! I started to examine the wiring system. I found that one of the Thinnet coaxial cables used to connect the concentrators was physically damaged and two others looked to be on their last legs. The administrator immediately agreed to order replacements for all three.

The administrator then pointed out that several of the individual sockets in the concentrators were inoperable. Since they were using bargain-basement concentrators, I suggested that they replace all of them while they still had time. The administrator didn't have the authority to authorize that purchase.

Next, I inspected the cable from the wall to the punch-down blocks. It did not look like any data cable that I had seen before. While it was definitely UTP, it just didn't look right.

I traced the cable from the punch-down block to the patch panel and found that the patch panel end was terminated in an old 50-pin telco Amphenol socket. This socket takes a bundle of wires within one sheath and spreads them out into 50 separate contacts. While I have heard of people trying this trick with 10Base-T before, I noted that this was a violation of standards. Putting multiple high-speed data wires together within that sheathing and especially inside the actual Amphenol connector was asking for trouble.

The front connectors on the patch panel blocks were RJ-12s, not RJ-45s - here we go again. The jumpers from the patch panels to the concentrators and, as it later turned out, from the wall sockets to the workstations were all the funny looking UTP that I had found coming out of the wall.

I got out my LANcat, a test instrument programmed with the IEEE standards for the cables used with all of the common networking systems made by Datacom Technologies, Inc. (800-468-5557) and also sold under the Fluke brand name. I used it to test the cables running from various workstations back to the concentrators. None of the cable passed any of the tests at the Level 3 standard. Levels are a measure of the grade of cable ranging from 1 to 5, with 5 being the highest, not to be confused with "Type", IBM's way of designating different kinds of cable. Though we usually install at least Level 4 for use with 10Base-T to provide more leeway, and many clients now demand Level 5 to allow for later upgrades to high-speed networking, the actual standard specifies a minimum of Level 3 UTP. Failure in the Level 3 tests indicated that the cable plant was useless for 10Base-T.

I took down the numbers on the cables and connectors that I was inspecting and read them to Bruce. He later confirmed that the cables, punch-down blocks, patch panels, and wall sockets were all

LEVEL 1! This information was so surprising that I had him send one of his crew down to personally check the cable plant and do more testing. The results corroborated my initial findings.

This was a double surprise. First, Level 1 is basically very low grade obsolete voice-only phone cable. Levels 1 and 2 had not been stocked by most cable distributors since 1990 when the price of Level 3 dropped. This means that not only did the cabling company screw their client (the only appropriate word for it), but they must have gone out of their way to do it! Perhaps they had a left over spool that they wanted to get rid of, so they stuck this client with it. Either that, or they special ordered a spool just for the occasion.

The second surprise was that the cabler was a prominent local company and LANDA (Local Area Network Dealers Association) member who should have known better. Inexperience could not be one of their excuses.

At my suggestion, the client called in the original cablers to get what amounted to warranty coverage. That company claimed that the Level 1 cable was just fine for 10Base-T (and, I presume, that they had conditionally repealed the laws of physics).

They sent one of their technicians to look at the system. He replaced the jumpers at the three most troublesome workstations with Level 5 jumpers and declared the problem fixed. (This is your brain. This is your brain on drugs.) Five minutes after the technician left, the system administrator called me asking for help. All workstations had fallen off the net and none could get back on. (Aha - the dreaded Humpty-Dumpty Syndrome!)

The massive differences in electrical characteristics between the Level 1 and Level 5 cables stopped the network dead in its tracks. This was predictable. Radically different levels of wire put together in this way create a reflection point within the cabling system, scrambling transmissions. I had the administrator replace the new jumpers with the old, thereby clearing up the new problem.

The management of the client company rejected any talk of rewiring, even though the system administrator spent much of his day running from workstation to workstation getting machines back onto the network. They were unable to press home warranty claims with the original installer.

Management did the next best thing. They brought in a mainframe consultant with no cabling or even PC experience to recable all of the workstation and concentrator jumpers, using a reel of cable and a crimping tool from the local Radio Shack! Hey, why didn't I think of that! (Now I know why there is a waiting period for gun purchases.)

Need I tell you the results? The network has deteriorated to the point where the system administrator spends most of his time resurrecting dead workstations. (Picture the Ed Sullivan Show. Picture the plate-spinning balancing act on that show using six plates. Now picture the same act using seventy-five plates.) No further action is being contemplated.

Part Two of the wacky world of wiring will be in next month's issue. I've got to go on a late night rescue. Let's see now, Jolt Cola, Pepto-Bismol, aspirin, TUMs, pizza delivery coupons. Yep, I'm ready to go.

HOT-WIRED, Part Two

by Wayne M. Krakau - *Chicago Computer Guide*, May 1994 - *NewsWare*, February 1996

Wayne is a partner in Krakau Business Computer Systems, a systems integration firm that is Novell Gold and Lantastic Five Star Authorized. He has been working with computers for 17 years. He holds an MBA in Marketing and a BS in Information Science and was originally educated in Computer and Information Systems Analysis Engineering. He holds Novell CNI and Enterprise CNE ratings and all 5 of the original LANDA Certifications. Wayne can be reached for questions, comments, or topic suggestions at (708)298-7695 or through the Computer Guide.

In the last tale in this trilogy of horror (in the Stephen King tradition), another reseller called us in to find out why some newly added workstations on his client's token ring were mysteriously slowing down and sometimes losing contact with applications servers on the network. The need for help was urgent, and the problem sounded suspiciously like it could involve the cable plant, so I had my cabling integrator, Bruce Kahn of Telnet Communications Consulting, Inc. (708-215-0003), a local cabling company meet us there. We ran our usual software tests as well as the diagnostic software that the client had. Both programs presented unusual results.

One program could see all of the nodes (a node being any entity on the network, such as workstations, direct network connected printers, file servers, and applications servers), but showed repeated bursts of errors, especially when someone shut down a workstation. The other program missed some nodes that we knew were logged on and showed errors on the nodes that it could find. Finally, it would start mixing up token ring addresses and login names!

To track down the last and most bizarre problem, we created a set of login names that corresponded to the cable numbering scheme used at this site. With all of the workstations logged in with their cable plant number, we could cross reference office layout diagrams to find which workstations were having problems. We also recorded the token ring addresses so we could detect further discrepancies.

In the next round of tests, some nodes were producing thousands of errors while others were disappearing and reappearing randomly. Because of token ring's design, we couldn't tell if the errors were originating at those nodes or if the actual problems were elsewhere and we were only seeing the aftereffects. The network wasn't stable enough to get accurate software-based test results.

At that point, we started examining the cable plant. The quality of the wiring organization was very professional, but looks alone are not enough to evaluate a network. The most obvious initial problem was the use of segments of UTP with media filters (for conversion between UTP and STP) to connect between wiring closets on a system that used STP in one closet and the workstations associated with it while using UTP in the other. UTP should not be used to connect dissimilar media closets. Fiberoptic cable is the only safe way to connect them. We even looked up that standard, just to make sure that we remembered the rule correctly. With the new, less expensive and more reliable methods of installing fiberoptic cable available, there is no excuse for violating this rule other than a lack of knowledge and training on the part of the installer. A bid based on the older fiberoptic installation procedures can be high enough to encourage clients to cut corners.

Next, using the cable lengths originally supplied by the cable installer, we calculated the total length of the longest UTP lobe (a lobe being the cable going to a workstation or other node) and the cable between the most distant wiring closets. This number was greater than the allowed number as listed in the standard tables used for UTP cable in token ring networks. This meant that even if we ignored the mixed media problem, there was still a major standards violation.

We got out my LANcat, a test instrument programmed with the IEEE standards for the cables used with all of the common networking systems made by Datacom Technologies, Inc. (800-468-5557) and also sold under the Fluke brand name. We started testing individual cable runs. While the cables themselves (both UTP and STP) and the terminations were all of very high quality, all of the sample runs that we analyzed tested out as longer than the cable installer said they were.

To double check our figures, we used a tape measure on the shorter runs and counted ceiling tiles on the longer runs to get an independent reading of the cable lengths. These figures matched the LANcat's readings. This made the UTP length violation much worse than we originally expected. No wonder we couldn't get a stable reading from software-based tests. High-quality cable and manual dexterity can't make up for ignoring standards based on the laws of physics.

As a short term fix, the applications servers were moved to the same area as the workstations that used them the most. This reduced the occurrences of workstations losing contact with them, but did not clear up the underlying problem.

At this time, client management has refused to rethink its cabling strategy. A series of classic misunderstandings is slowing progress. The cabling company, probably not as experienced in data cabling as they presented themselves, warned the client of possible problems, but did not press the issue. While many cabling companies would have declined the project on ethical grounds, not wishing to burden a client with a cable plant guaranteed to cause problems, this company deferred to what they may have assumed to be the greater knowledge of the MIS department of the client corporation.

In a typical train of thought, the MIS people work with those big, scary mainframe computers, and therefore, must be the ultimate computer experts. PCs are computers, and LANs involve PCs, so they must be experts on them, too. Riiiiiight.

In reality, mainframe computer people may be the only group more strictly divided into specialties than the medical profession. As evidence, recruitment procedures are almost comical, with MIS managers demanding a match on industry, specific application within that industry, language, dialect, sub-dialect, database, database dialect, code generator, report generator, and programmer utilities - all just to qualify for an initial interview! Even within those narrowly defined specialties, the knowledge level varies according to a standard bell curve, with a few real experts at the high end, a few filter-feeders at the low end (named for the sea creatures that just sit there and gather food as it drifts by), and a great mass of just plain average folks in the middle - exactly like most other professions (including systems integrators).

The MIS people saw planning a cable plant as a cabling-only issue. They deferred to the people they thought that were experts in that field and assumed that the installer would refuse to install a defective cable plant. They may also have greatly underestimated the amount of knowledge needed for them to create a reliable preliminary plan prior to bringing in the cable installer. Remember, I've been beating on LANs for over 11 years and have had extensive related experience (see the bio at the beginning of this column), and I still bring in Bruce and his people to work with us on cable planning and any serious troubleshooting. I know the value of bringing in fellow professionals who are experts in fields that I can only dabble in.

Over the years, I have become very cynical about cable plants. That's because I see so many bad ones. The possibility of using spare pairs of telephone wires as network cables has never materialized for any company that I've dealt with. Even using data cable pulled by phone cabling specialists has been a complete disaster. The only systems that I've seen work were meticulously planned, professionally installed and warrantied by experienced data cable installers, and tested not just for continuity, but using sophisticated test instruments, for compliance to international standards appropriate to the specific type of network. That's the only sure way to avoid cabling nightmares.

DOWN TO THE WIRE

by Wayne M. Krakau - *Chicago Computer Guide*, June 1994 - *NewsWare*, March 1996

Wayne is a partner in Krakau Business Computer Systems, a systems integration firm that is Novell Gold Authorized. He has been working with computers for 17 years. He holds an MBA in Marketing and a BS in Information Science and was originally educated in Computer and Information Systems Analysis Engineering. He holds Novell CNI and Enterprise CNE ratings and all 5 of the original LANDA Certifications. Wayne can be reached for questions, comments, or topic suggestions at (708)298-7695 or through the Computer Guide.

Due to the unprecedented response to my last two columns, "Hot Wired," parts one and two, I decided that a follow-up was in order. The general theme for many of the inquiries I received was "I want to fudge the standards in the following way and have been advised by _____ (fill in the blank, usually with "cabling company", "consultant", or some variant of "reseller") that it will work -- can I get away with it?" That's a scary thought.

In one typical example, a 4Mb (megabit) token ring network was being upgraded to 16Mb. The exiting Level 3 (10MHz rated) cable was augmented by replacing all external portions (jumpers) in the wiring closet and at the workstations with Level 5 (100MHz rated) with the expectation that the cable plant as a whole would now be able to support 16Mb token ring.

I suspect that the math required to believe in the viability of this "upgraded" system is some kind of averaging: (Level 3 + Level 5)÷2 = Level 4 -- right? WRONG! Level 5 added to a Level 3 cable plant is usually worse than having all Level 3.

Each spot at which two different kinds of cable meet creates a potential point of distortion. Multiply that times the number of individual lines used and then again by two to account for one jumper at each end of the internal lines. Even if this effect wasn't a consideration, you have to remember that cables work on the weakest-link principle. Even one bad or inappropriately rated segment can wreak havoc on an entire network.

I think I may have inadvertently set myself up as some sort of father-confessor. People seem to be asking me for a special dispensation so they can use invalid cable plants. Sorry, I can't give one. The Laws of Physics keep getting in the way. Feel free to invoke the deity of your choice, and be sure to get back to me if you are successful. That would be in the same category as a weeping icon.

Here's the way to avoid messes like this one. First, deal with cablers who are familiar with the rules. This doesn't mean that they necessarily have memorized every rule, just that they know the general parameters well enough to recognize possible deviation and aren't afraid to look up the details (especially for the potentially complicated calculations required for token ring!).

Also, they must have the ethics ("Have pulpit, will preach", is my motto) to turn down cable installations that are inherently invalid. If you find discrepancies via other sources, and a cabler is still willing to install the invalid cable plant, look elsewhere!

Don't allow the cabler to get away with merely testing for continuity. They must use appropriate testing devices. While there is technically no such thing as a true "certifying" device, there are several instruments that can come as close as possible to this ideal. A bonus for this type of product is that most of them can print a cross-reference report listing the individual cables that were tested along with the details of the test results.

The final qualification of the cabler company is the willingness to back their work. Will the company send qualified technicians to fix problems? Will they reimburse your company for expenses due to having a LAN person track down mysterious errors that turn out to be cabling problems? This is the final and most important characteristic to look for. After finding errors in many existing cable plants over the years, I have found that most cablers won't support their work, even in the face of overwhelming evidence of a defective system.

The King is Dead, Long Live the King

The mournful sound of Taps wafts over the fields. A riderless horse with reversed boots in its stirrups passes at a slow walk. Artillery is aligned for a final salute.

No, I am not talking about deceased ex-presidents. I am talking about the recently announced demise of Netware 2.2. Novell has finally pulled the plug on the last direct descendant of what was originally known as Netware 286, itself a derivative of the old Netware 86 and in turn Netware 68. (Note that the digits were added after the fact for differentiation of each type from its successors.)

Like the recently expired ex-president, Netware 2.2's reign was a controversial one. Novell released it in the aftermath of a wave of protests about a proposed freeze on development on what was then known as the Netware 286 line. Novell succumbed to an immense backlash. People resented being "forced" to upgrade to Netware 3.11. That's like being forced to go from a hand-cranked engine to an electric starter.

While it carried on the legacy of its forebears (Netware 2.15c being its immediate predecessor), it really was a dead-end track on the evolutionary trail of network operating systems. Released at a time when it was generally acknowledged that the Intel's 80286 was dead, Netware 2.2 ended up being a boat to nowhere.

Unscrupulous and unknowledgeable resellers were happy to sell an operating system with a proven record of high installation and maintenance costs. (To this day, Novell's competitors continue to hold up this reputation as evidence that prospective customers should consider their operating systems, or just to get people to move to peer-to-peer systems. They neglect to tell people that a simple Netware 3.x LAN is actually easier to install than a peer-to-peer system for an installer experienced in both.) The resellers had the bonus of knowing that their customers would eventually have to upgrade their LANs when the artificially resurrected 286 line ultimately became extinct. Depending on the number of 2.2 boxes left on resellers' shelves, this opportunity will soon be eliminated.

Novell has taken a carrot and stick approach toward motivating the remaining Netware 2.x users (over one-third of all Netware customers, the last I heard) to upgrade. The carrot is a discount on upgrades to either 3.12 for 4.01 Netware with the promise of a free upgrade to the upcoming 4.1 for those who choose 4.01. Even 3.11 customers get the discount when upgrading. The stick is the threat of a gradual loss of support for older systems.

An interesting side note to this is that I just found out about an unpublicized upgrade opportunity for owners of older versions of Netware at least back to 2.0a, and probably including 2.0, too. Even though upgrades from these versions have long been unavailable through resellers, I recently found that they can be obtained at bargain prices directly from Novell. Just call 1-800-NETWARE and walk your way through the menu system to get the sales department involved with upgrades.

Since I have had multiple conversations with Novell personnel in other departments (mostly the reseller group) during which I was told that older versions of Netware were not upgradeable, I was quite surprised to hear of this option. I suspect that the impending doom of Netware 2.2 might have some effect on the availability of these old-system upgrades, so if you are interested, act now!

DOCUMENT AND IMAGE MANAGEMENT SYSTEMS - Part 1

by Wayne M. Krakau - *Chicago Computer Guide*, July 1994 - *The Law Works*, September 1994 - *NewsWare*, October 1995

Wayne is a partner in Krakau Business Computer Systems, a systems integration firm that is Novell Gold Authorized. He has been working with computers for 17 years. He holds an MBA in Marketing and a BS in Information Science and was originally educated in Computer and Information Systems Analysis Engineering. He holds Novell CNI and Enterprise CNE ratings and all 5 of the original LANDA Certifications. Wayne can be reached for questions, comments, or topic suggestions at (708)298-7695 or through the Computer Guide.

A full-featured document and image management system (DIMS) integrates the management of the images of text and graphic documents as scanned in from paper or film, text and graphic files through import routines, and the text that results from using optical character recognition (OCR) to translate images of documents consisting of dots back into logical text.

In this article, I will be mentioning several of my favorite products. I don't claim to give unprejudiced opinions. I do believe, however, that nobody gives unprejudiced opinions - not systems integrators, not VARs, not dealers, and especially not "independent" consultants. I just make it my duty to make my prejudices clear so that you may add as many grains of salt to my opinions as you think is necessary. I only ask that you keep in mind that I am in a position to get authorized to sell just about any product that I care to. This disclaimer is my way of avoiding penalizing companies with good products just to avert accusation of conflict of interest.

If you are interested in document and image management systems, you need to know about, or perhaps even join AIIM, the Association for Image and Information Management (301-587-8202). AIIM is just starting to experience the same changes that publishing has gone through over the last few years. Their industry was originally ruled by large proprietary vendors fighting to lock customers into their own unique (and incompatible) way of doing things. This gradually transitioned into an intermediate stage wherein these manufacturers offered scaled down systems in an attempt to stem the tide. These shrunken systems contained standard pieces of PC technology with a different brand name slapped on the front and the price tag boosted to levels barely below the old systems. Just to make things interesting, they often tweaked these systems just enough to make them just as proprietary as the old systems.

The biggest breakthrough in AIIM these days is not technological - it is political. Nien-Ling Wayman, President and founder of Compulink Management Center, Inc.(310-212-5465), the makers of the LaserFiche family of document and image management software (my favorite), has been elected to the AIIM board of directors. This is the first infiltration by a representative of an open PC/LAN product into an organization dominated by manufacturers of large, expensive, proprietary systems. I have commented previously that her most difficult job was to refrain from laughing aloud when one of her colleagues tells her that a particular project either CAN'T be done or that it can be done - but it will cost a MILLION dollars!

First, here are some examples of existing document and image management solutions, just to give you an idea of what they're being used for. The uses are not always obvious. (In case you are wondering, these are LaserFiche installations.)

The Los Angeles Department of Water and Power was sued for $24.5 Million. (Does anybody remember the movie "Chinatown", and how complicated water allocation is in California?) They purchased a document and image management system to store all documents and research material associated with the case. They used this system as leverage to negotiate an out of court settlement for only $2 million. That was less than the estimated legal fees for a multi-year courtroom battle. The opposing attorneys were so impressed that they later purchased a system for their own firm's use. This is an

example of how lawyers can use document and image management systems. Many such systems were originally designed with legal clients in mind.

Northrop uses a document and image management system to store their Material Safety Data Sheets. They had a legal requirement to have information on the chemicals that they use available for immediate access in an emergency. They now store 8000 active documents and 10,000 inactive documents, all available for instant access. It eliminated fines for non-compliance while reducing paperwork handling costs.

The IRS is using a system to track the documents used in going after money launderers. They store both the legal filings and evidential documents for instant access and long term storage.

20th Century Fox stores their script library in a document and image management system. It includes every script that they own, including the thousands that were never produced. They consider this database one of their greatest assets, in that any new idea can instantly be compared with this database so that the ideas contained within the old scripts can be utilized and potential copyright violations can be avoided.

Best Foods started putting research materials into a document and image management system after they accidentally duplicated an expensive report which had been done years before in a different department. Now they scan their document database prior to beginning any research project.

Merisel, the giant computer distributor, uses a document and image management system as the base for their fax-on-demand system that is used by many resellers. They call on a touch tone phone or directly via their fax phone and can get product information faxed back to them.

As you can see, document and image management systems are very horizontal (as opposed to a vertical market) in nature. That is, they are useful across many different types of businesses.

Before I cover the technical aspects of document and image management systems, I want to give a warning. Most hardware and system software used in document and image management systems were designed for and perfected on networks. The single user versions of drivers that make the sophisticated hardware work are frequently unstable and inefficient. Be prepared to be an unwilling beta site for these single-user systems.

The first step in planning a document and image management system is to contact a lawyer who specializes in the legality of electronic documents. Unless you have no interest at all in the legal standing of your stored documents, this is a must. It is the only way to protect yourself. Also make sure that the product you choose absolutely forbids the alteration of the underlying original image. You should remove any image editing capability from scanning workstations - including something as simple as Window's Paintbrush. Even the possibility of altering the original may invalidate the legal standing of your images.

I will continue with Document and Image Management Systems, Part 2, next month. It will cover the technical aspects and will include examples on how my clients are using the technology.

DOCUMENT AND IMAGE MANAGEMENT SYSTEMS - Part 2

by Wayne M. Krakau - *Chicago Computer Guide*, July 1994 - *The Law Works*, October 1994 - *NewsWare*, November 1995

Wayne is a partner in Krakau Business Computer Systems, a systems integration firm that is Novell Gold Authorized and Novell IBM Host Master Authorized Reseller. He has been working with computers for 17 years. He holds an MBA in Marketing and a BS in Information Science and was originally educated in Computer and Information Systems Analysis Engineering. He holds Novell CNI and ECNE ratings and all 5 of the original LANDA Certifications. Wayne can be reached for questions, comments, or topic suggestions at (708)298-7695 or through the Computer Guide.

This is the second part in a series of columns covering Document and Image Management Systems (DIMS). These systems integrate the management of the images of text and graphic documents as scanned in from paper or film, text and graphic files through import routines, and the text that results from using optical character recognition (OCR) to translate images of documents consisting of dots back into logical text. Part 1 of the series covered the politics of this particular corner of the computer industry, some samples of completed systems, and the legal implications of implementing a DIMS.

The first technical aspect of DIMS that you must worry about is to develop a plan to accommodate or avoid the throughput problems that usually follow the installation of a document and image management system. I am on call for one manufacturer of document and image management systems to help them plan and troubleshoot systems. So far, most of my work for them has been to come in after a marginal network has collapsed from faulty cabling due to the increase in traffic from the new document and image management system. Many systems are already teetering on the edge of disaster. The extra traffic from a DIMS can be enough to put them over that edge. Either response time becomes excessive from the multiple retries necessary to get a message over the network, or the system breaks down to the point where individual workstations start dropping randomly off the LAN. (Refer to my April, May, and June columns for more information on the care and feeding of cable plants.)

For LANs without cabling problems, response time can still be an issue. The sheer quantity of traffic may be enough to cause untenable delays. A new DIMS is useless if nobody can get their work done.

The most obvious way to overcome throughput problems is through the brute force method. For instance, I am currently designing a network for an engineering firm. While they don't plan on an imaging system, they do have a similar need for a high-speed LAN. I am going to use Madge Networks, Inc. (800-876-2343) upcoming 100 megabit unshielded twisted pair network cards along with a Synoptics Communications, Inc. (800-776-6895) concentrator. This system will follow the soon-to-be-approved MLT-3 standard, a derivative of FDDI (Fiber Distributed Data Interface, a 100-megabit descendent of Token Ring using optical fiber). It's fast but very expensive.

The second way to attack this problem is to split the network. Netware 3.x and higher can handle up to 16 network cards. I often use Standard Microsystems Corporation (800-SMC-4YOU) Elite Ultra 32T card. It's a dual channel Ethernet EISA card with two complete Ethernet chipsets and connectors. The server sees it as two different cards. With SMC’s card or separate cards of your own choice you can split the network several ways.

You can just split it arbitrarily 50/50 (or into smaller segments with more additional network cards).

You can isolate critical workstations. In one case, the engineering department of a client required extra throughput, so we are going to give them a dedicated concentrator attached to one of the Ethernet ports. The other port serves the rest of the network.

Another way is to isolate those devices that could overburden the network. This would include any scanning stations, network printers, and fax servers. This is particularly applicable to imaging, since

scanning can suck up loads of network bandwidth. The regular users could be on a separate segment from these high-bandwidth devices.

The third way to overcome throughput problems is my favorite. That is to use client-server technology to reduce the amount of data going over the network. Client-server technology splits a database into two parts, with the client portion running on the workstation and the server portion running on the database server (which may also be the file server). The workstation software presents the visual interface, supplies a help system, does some preliminary error checking, and formulates requests for information or update requests based on the operator's input. The request for information is submitted over the network to the server software.

The server software acts upon the request by searching its hard disk for the appropriate information or by making a final check of the update request. For an information request, just the data needed is sent back over the LAN. For an update request, only an acknowledgment is sent back.

Without client-server technology, the individual workstation is completely responsible for searching the underlying database and processing any updates, causing potentially enormous quantities of data to clog up the network.

You can use a simple client-server database to track indicative (document identification) data. Full-featured systems use OCR, or Optical Character Recognition, to create ASCII text out of the files scanned into the DIMS. They also allow queuing of the OCR task for processing by multiple OCR computers. (A handy technique is to scan the documents during the day and OCR them using all available idle computers during the night.) The resulting text is the base for a searchable full text index and also needs to be client-server.

Finally, the database tracking the images themselves absolutely must use client-server technology. Since I cover the Novell-centric end of the business, I prefer systems that use NLMs, Netware Loadable Modules, to allow client-server access. To further reduce traffic, the actual images should be stored in a compressed format and be passed over the wire still compressed. They can then be expanded at the workstation. There are two other benefits to the NLM method.

First, as an NLM, no actual person opens the database. You may have experienced having a workstation go down in the middle of a transaction, thereby blowing up your database. With properly implemented client-server technology, workstations merely submit requests and receive replies. Only the server software opens and manipulates the database. So, a workstation failure will not jeopardize the database. It is, of course, implied here that you have protected your server with an Uninterruptible Power Supply. The UPS combined with the inherent reliability features of Netware such as Read after Write Verification, Dynamic Bad Block Remapping (Hot Fix), and the Transaction Tracking System (TTS), make the file server the most dependable machine on the network.

Second, if the imaging database grows so that its NLM takes up too many server CPU cycles - that is, it overtaxes the server's Central Processing Unit, you can get a product called an NLM Server from a company formerly called IWI and now named CommVision (415-254-5720). It is a complete 486 computer on a single card, including RAM. It comes bundled with Run-Time Netware, so you can off-load the imaging NLM onto this card. The NLM Server is placed within the file server, so it maintains bus-speed access to the disk subsystem and the server's network cards. This is often a better solution than having a totally separate machine as a database server that then has to talk to the file server via a LAN connection. Of course, the NLM Server can run other types of NLMs, too, and you can use more than one in a single server. If you need to have your imaging server, your communication server, your fax server, your E-mail server, and your mainframe or minicomputer gateway (or even some less-complicated combination) all in one box, then check out CommVision's 20-slot monster (my description, not theirs) computer. It has triple redundant, hot-swappable power supplies, a 100-megabit backbone, and a segmentable 20-slot backplane. That will cut down substantially on network traffic.

This series will continue with Document and Image Management Systems, Part 3, next month. It will cover storage methods and integration issues.

DOCUMENT AND IMAGE MANAGEMENT SYSTEMS - Part 3

by Wayne M. Krakau - *Chicago Computer Guide*, September 1994 - *The Law Works*, November 1994 -

Wayne is a partner in Krakau Business Computer Systems, a systems integration firm that is a Novell Gold Authorized and Novell IBM Host Master Authorized Reseller. He has been working with computers for 17 years. He holds an MBA in Marketing and a BS in Information Science and was originally educated in Computer and Information Systems Analysis Engineering. He holds Novell CNI and ECNE ratings and all 5 of the original LANDA Certifications. Wayne can be reached for questions, comments, or topic suggestions at (708)298-7695 or through the Computer Guide.

This is the third part in a series of columns covering Document and Image Management Systems (DIMS). These systems integrate the management of the images of text and graphic documents as scanned in from paper or film, text and graphic files through import routines, and the text that results from using optical character recognition (OCR) to translate images of documents consisting of dots back into logical text. Part 1 of the series covered the politics of this particular corner of the computer industry, some samples of completed systems, and the legal implications of implementing a DIMS. Part 2 covered the technical aspects of planning a DIMS. This column will cover storage methods and integration issues.

The conventional storage technique for a DIMS is to put the searchable database and index information on your fastest magnetic drive and put the images on optical drives, using jukeboxes as more capacity is needed. I prefer using Micro Design, Incorporated's (407-677-8333) SCSI Express NLM (Netware Loadable Module) to manage optical drives. MDI's product, unlike most of its competitors, taps straight into Netware's native capabilities, such as caching and elevator seeking.

Elevator seeking is particularly valuable for jukeboxes, since it can substantially reduce the need to switch cartridges. For example, assume that three read requests come in and that request number one is for cartridge ten, side A, request number two is for cartridge twenty-three, side B, and request number three is for cartridge ten, side A. Without elevator seeking, the system would mount cartridge ten, then cartridge twenty-three, and then cartridge ten, again. With elevator seeking, the system would mount cartridge ten, satisfy both requests one and three, and then mount cartridge twenty-three. No matter how fast the jukebox (or the unfortunate person assigned to swap cartridges on a single drive) changes cartridges, it is many thousands of time slower than a simple read of the currently mounted cartridge.

While I am covering optical storage, I did want to mention that distribution via CD-ROMs is becoming very popular for imaging systems. A notebook computer with a CD-ROM drive makes a great mobile office. One of my clients, for instance, is planning to produce a thousand CDs every few months containing all of their sales information and product specifications, and distribute them to their worldwide sales force.

Remember, however, that CD-ROMs are designed for mass distribution. WORM (Write-Once, Read Many or Mostly) and Read-Write optical drives are appropriate technology for primary storage in a DIMS. Considering the number of questions I have been getting lately, these two separate technologies are commonly confused.

Some of the latest developments in optical jukeboxes include 2.4TB (a terabyte is 1,000,000,000,000 bytes or 1,000 gigabytes) Ethernet-connected jukeboxes that use 12-inch cartridges. Also several companies are using giant jukeboxes modeled after mail sorting cabinets, with multiple rectangular slots in a large array. Slots are accessed by an arm attached to a pair of crossbars that move like a giant "Etch-A-Sketch". (Does anybody else remember them? They had the original Graphical User Interface long before the Xerox PARC - Palo Alto Research Center - invented them and way before Apple copied Xerox's technology.) Each slot can carry either a rack for multiple cartridges or an optical drive. Any cartridge-holding slot can be later retrofitted with another drive, thereby negating the one big drawback of jukeboxes. That drawback is the idea that an individual drive cannot have a cartridge

swapped until the existing cartridge is finished being used. Even if the cartridge-changing time is very short, everybody waits if the drive is in use.

Another storage method that has just become financially feasible is RAID. RAID means Redundant Array of Inexpensive Disks - as opposed to SLED which means Single Large Expensive Disk.

RAID technology is broken down into several categories called levels as follows:

0 - Data striping without parity. That means data is spread out over multiple disks for speed.

1 - Mirrored disk array. For every data disk there is a redundant twin. Also includes duplexing, the use of dual intelligent controllers for additional speed and reliability.

2 - Bit interleaves data across array and reads using only whole sectors.

3 - Parallel disk array. Disk striping with dedicated parity drives. Drives are synchronized for efficiency in large parallel data transfers.

4 - Independent disk array. Reads and writes on independent drives in the array with dedicated parity drive using sector-level interleave.

5 - Independent disk array. Reads and writes data and parity across all disks with no dedicated parity drive. Uses parallel transfers. Multiple controllers optionally used for higher speed. Usually loses only the equivalent of one drive of array for redundancy. This system is the most popular these days due to both speed and cost effectiveness.

Micropolis (415-964-8300) has a series of RAID systems that they call Radion LT. The largest of these provides 3.52GB (gigabytes) of storage and is expandable to 54GB by stacking up proprietary modules like oversized Leggo blocks! As you add more disks, you can add more SCSI adapters. If a drive fails, the system keeps on running. You can hot-swap the drive without shutting the system down. As the system is expanded, you can even lose multiple drives and still keep running! In conjunction with this software-based system I have been using Adaptec's (800-934-2766) 2742AT twin EISA SCSI adapter. It puts two separate SCSI adapter chipsets on one EISA card. Netware sees it as two separate adapters. Note that the better RAID systems allow the use of multiple controllers to increase throughput.

Pacific Micro Data (714-838-8900) has taken a different tack. They use a 7-slot tower with sliding trays to accommodate anyone's disk drives. They suggest using a Mylex (510-796-6100) controller. It has three separate SCSI ports and hardware RAID on a single EISA card. This offloads the overhead of RAID from the server's processor. Note that I prefer dealing with Mylex products through third parties such as Pacific Micro Data, since Mylex seems to be set up more to support OEM agreements than direct dealer and end-user relationships.

Compatibility is the next thing to watch for. Just because a program runs as a part of Microsoft Windows, doesn't mean it's really a full Windows product. Proprietary routines are common, especially in a market like DIMS where many of the key players are transitioning from selling the old-style proprietary systems to supplying PC and LAN oriented products. I've even received complaints from dealers that some imaging products that they currently carry purposely force incompatibilities so that you have to buy into their proprietary solutions for such things as printing, network faxing, general communications, remote control, E-mail, and even hardware such as scanners, disk drives, and jukeboxes. Their clients are really mad at them for setting them up with products that have these restrictions. Some systems also require a separate imaging server, often running an operating system that would force you to seek outside support. If you are not an OS/2 shop, for instance, why should you be forced to have an OS/2 box? Keep in mind that the separate server technique slows things up both because the servers are usually linked only by a slow network cable and because of the relative speed of NLMs.

Integration among programs is one of the reasons that people go to graphical environments like Windows in the first place. To integrate programs within Windows, a high degree of compatibility is needed. Only programs that use standard Windows routines can be easily integrated using macro programs or a language such as Microsoft’s Visual Basic.

One client I am working with needs to integrate an IBM AS/400 minicomputer application used to run the business with a smaller related LAN-based database that holds a subset of the AS/400 data, as well as a Compulink Management Center's (310-212-LINK) LaserFiche DIMS. We are planning on using either Attachmate's (800-426-6283) EXTRA or Wall Data's (206-881-5995) Rumba terminal emulation products along with their respective Visual Basic interfaces to feed data to the AS/400. We are using Microsoft Access with Visual Basic's built-in capability to manipulate Microsoft's Access databases to handle the LAN data. The document and image management software will handle the imaging and the document-based indicative information. The user will enter the data once into an Access database, and then Visual Basic will copy the appropriate fields into the AS/400 program and into the imaging system's indicative fields while Access will retain only the subset that it needs to satisfy those users who don't have access (note the small "a") to the AS/400.

Another client wants to use Lotus Notes, but is unimpressed with the limited imaging add-on available for Notes. We are going to tie LaserFiche and Notes together to satisfy their needs.

It is very common to tie accounting systems and contact managers to DIMS. In a typical use, a customer calls up to ask why a delivery has not arrived. You bring up the customer's information in your accounting software. Simultaneously, an automated link to your DIMS brings that program up in the bottom half of your screen. It then grabs the customer number from the accounting system and passes it to the search field in the imaging program. This brings up the documents associated with this customer. Then you can view the receipt that shows that Joe on the loading dock signed for the package at 4:00PM, yesterday. Then, while you are still speaking to the customer, you could fax the receipt to them using your network-base (of course!) fax system.

The last item I want to mention is input. It is a great help to be able to batch-import text and graphics that already exist in electronic form. Even though a fuzzy search feature can help get around OCR inaccuracy, direct import of text will always be more reliable. Direct import can also retain the quality of graphic images. There can also be a great savings in labor costs. Additionally, with batch-import of graphics, even old microfiche and microfilm stock can be easily input.

Document and image management systems are a great solution to many potential problems, including storage costs, retrieval speeds, misfiling, and data analysis. It does, however, take careful planning to fully integrate all of the new technologies that make them run best.

LAN DESIGN - Part 1

by Wayne M. Krakau - *Chicago Computer Guide*, October 1994 - *NewsWare*, April 1996

Wayne is a partner in Krakau Business Computer Systems, a systems integration firm that is a Novell Gold Authorized and Novell IBM Host Master Authorized Reseller. He has been working with computers for 17 years. He holds an MBA in Marketing and a BS in Information Science and was originally educated in Computer and Information Systems Analysis Engineering. He holds Novell CNI and ECNE ratings and all 5 of the original LANDA Certifications. Wayne can be reached for questions, comments, or topic suggestions at (708)298-7695 or through the Computer Guide.

This article will cover Network Design, and, by implication, management and troubleshooting. I am going to cover some nuts and bolts techniques and also give some more esoteric advice, generally concentrating on a proactive approach. It's always easier to avoid trouble than fixing it and it's better to be prepared for trouble than to be surprised by it.

In real estate the general theme is "Location, Location, Location". In networking the general theme is "Planning, Planning, Planning". There is a direct relationship between the amount and quality of planning that goes into a network and its eventual stability and reliability. Just because Local Area Networks look small, don't underestimate their complexity. Once you have more than one person sharing a computing resource, it's just as complicated and potentially hazardous as if you have one hundred. Only the economies of scale are different.

The esoteric aspect to the planning process is the application of human resources. In the corporate world, in particular, it is very common to have upper management hold the opinion that their programmers and associated personnel are computer "experts" and that since PCs are computers, these programmers must also be experts on PCs and LANs. This is a logical trap that can compromise the planning process. An active programmer has the wrong mindset and generally an inadequate background to handle the sophisticated and specialized planning required. The person assigned is given the responsibility but not the tools to complete the project.

An allowance must be made for additional training, considerable time for research, and for outside help as necessary. As an independent CNI (Certified Netware Instructor), I constantly see people taking CNE (Certified Netware Engineer) courses who are inadequately prepared for the course material. Tinkering with a PC at home or even at work while your main job is on other computer architectures doesn't equip an individual for LAN design. Even programming PC's doesn't prepare you. It just qualifies an individual to be a potential quick learner and makes it less likely that the individual will be afraid of the technology. Even though such an assignment plays to the programmer's ego, it is inherently dangerous and the turnover rate for LAN administrators gives evidence to that fact.

This is not meant to be a slam against programmers. If I ever forget to wear my seatbelt and end up going face-first through my windshield, and someone offers me a referral to a doctor who is world-renowned in the field of viral research at the Centers for Disease Control, I would decline. Most likely, I would request an experienced reconstructive plastic surgeon - preferably board-certified. It's not that I don't respect the intelligence or competence of the virologist, it's just that virology, while still a vital part of medicine, isn't even close to the specialty that I need. The mindset is wrong and the proper skills are not present. The same situation exists in my suggested approach towards programmers in a LAN design scenario. They are decidedly the wrong people for the job. I won't go into the ethical ramifications involved when a programmer fraudulently uses a "Consultant" title and undeservedly claims expertise in LANs.

A trap that some firms fall into is to hire a high-end person. The problem there is that you end up spending premium dollars for an individual who is essentially idling 98% or more of the time with skills degrading all the while. One possible method out of this trap is to use a medium-skilled individual to supervise one or more lightly skilled, but ambitious staff members. Then, develop a relationship with

your favorite systems integrator so that a high-end person can be brought in as needed. This saves a lot of money while building in-house expertise. It also avoids the problem of having someone with an inappropriate background being forced to "wing-it". Since this is a potentially self-serving suggestion - my firm has been a systems integrator since 1983 and I worked for a mainframe-based systems integrator in my second-to-the- last big iron job - feel free to take it with whatever grains of salt that you feel necessary. However, please keep in mind that, for maximum personal profit, I could have suggested that you become completely dependent on a systems integrator for even day-to-day network management rather than maximizing the use of in-house personnel.

An additional consideration regarding programmers is the choice of software to use on your network. Programmers will normally lean towards the most programming-intensive choices. My suggestion on how to prioritize your software search is to start with unmodified packaged software, right off the shelf, at the top of the list, and end with completely custom code at the bottom. With Microsoft Windows becoming so popular, the most common and successful systems that I see use Visual Basic (or a similar tool) as the “glue” to combine multiple packages into a logically united system. By using the appropriate toolkits and the built-in flexibility of Windows, programming can be minimized while providing a virtually custom solution.

One of the most common problems that I see is a faulty cable plant. A bad cable plant has the potential of becoming the most time-consuming, labor-intensive, and incredibly expensive debugging session that you have ever experienced.

To even briefly summarize Ethernet cabling standards requires at least two pages of tiny type - and token ring is much worse! It is so complicated that I would have to write a booklet, as many manufacturers of token ring products have done, to get the same level of detail as two pages on Ethernet. Manufacturers’ literature can be a handy reference if you are knowledgeable enough to pick your way around proprietary features and limitations.

While I have heard rumors that such a thing is occasionally discovered, I personally have never seen a successful conversion of old telephone wiring into a LAN cable plant. I haven't even heard about one from other systems integrators and VARs either. Since most alleged "LAN cablers" can't even get their act together, how can you expect telephone cable installers, trained in much lower standards, to get it right? If you are thinking of using existing cable, even if someone tells you it was installed to high-speed data standards, have qualified personnel test it before planning on using it. We are not talking about simple continuity tests here. A rusty coathanger can pass one of those. We are talking about using test instruments specifically designed for a LAN environment. Even though my company uses another company to install cable, I encounter bad cable so often that I carry a LAN cable testing instrument with me as part of my standard toolkit. There are several brands and models available - just make sure that you get the cables tested by someone who can interpret the results. Depending on the size of your cable plant, you may want to purchase a test instrument and get training on how to use it.

Pick qualified installers. A select few telephone cable installers have successfully made the transition to LAN cabling, but most are just winging it. Try asking questions about the standards. Request an installation that doesn't follow standards and watch for the reaction. If the cabler declines the project on ethical grounds, that's a good sign. Ask about fiber. If they give you information that indicates that they are only familiar with old, expensive methods of putting in fiber and aren't trained in 3M's Hot-Melt system, that's a bad sign. They should also be testing with appropriate instrumentation and be willing to back their work.

I will continue this article on Network Design next month. Meanwhile, there are LANs out there awaiting rescue.

LAN DESIGN - Part 2

by Wayne M. Krakau - *Chicago Computer Guide*, November 1994 - *NewsWare*, May 1996

Wayne is a partner in Krakau Business Computer Systems, a systems integration firm that is a Novell Gold Authorized and Novell IBM Host Master Authorized Reseller. He has been working with computers for 17 years. He holds an MBA in Marketing and a BS in Information Science and was originally educated in Computer and Information Systems Analysis Engineering. He holds Novell CNI and ECNE ratings and all 5 of the original LANDA Certifications. Wayne can be reached for questions, comments, or topic suggestions at (708)298-7695 or through the Computer Guide.

This article is a continuation of the Network Design theme of my previous column. Last month I covered human resources, software, and cable plants. This month I will start with reliability issues.

The ability to keep on running after a failure is called fault tolerance. There are a near-infinite number of shades of fault tolerance that you can use, depending on how risk averse you and your organization are. Risk aversion is a term commonly used in the insurance industry and in stock trading to describe what Clint Eastwood characterized a bit more graphically on film. While pointing a gun at a criminal's head, Clint (portraying "Dirty Harry") said "Do you feel lucky?". That's what it boils down to. How lucky do you feel and what is the cost of guessing wrong? (That sound you hear is the distinctive click of the hammer coming down on an empty cylinder.)

The ultimate in safety is the use of redundancy. For instance, using twin servers and Novell's SFT (that's System Fault Tolerance) Level III allows one server to take over from the other automatically and transparently during a system failure or even during routine maintenance and upgrades. There are several less expensive aftermarket options available that will allow this trick with a somewhat lower level of transparency. With these alternate products, you have to physically intervene to get up and running on the redundant server.

Many of my clients implement a Secondary File Server/Workstation. It runs day-to-day as an unassuming workstation, but, in times of crisis, can be rebooted as - a file server! (You were expecting Superman?) Most of these secondaries have an extra hard disk with the same capacity as the one in the main file server. Some have just enough capacity to run critical applications, though that makes them much more difficult to maintain. Either way, one just needs to apply the previous night's backup to be up and running in case of a main file server crash. For an even quicker, response, apply the backup every morning as part of the daily routine. In an emergency, just shut down the malfunctioning main server and reboot the secondary as a file server. You will only lose an average of a half day of work with minimum downtime. Since the computer is normally used as a workstation, the cost of having this standby capability is quite small. Just add in the cost of an extra hard disk, maybe a little extra memory, possibly an upgrade to the next faster CPU or underlying architecture, plus a little extra labor to set it up. As I said, this option has been used successfully by many of my clients for the past several years.

Another way to increase reliability is to protect the part of the system most likely to fail - the hard disk. Mirroring means using a redundant second hard disk, controlled by the CPU of the file server. If one disk fails, the other takes over.

A few years ago, I received a phone call for help with a routine WordPerfect problem from a client. Just as we were both hanging up after resolving the problem, the client suddenly yelled my name. Luckily, I heard him. He explained that "By the way" (like it was no big thing), he had noticed a "FUNNY" message on the file server screen. After my pulse dropped back below 180 and I regained the ability to speak coherently, I calmly (or at least as calmly as possible), noted that, while using the phrases "funny message" and "file server" in the same sentence was linguistically and grammatically correct, it was, perhaps, not an acceptable combination for use in polite society.

I encouraged him to read the message to me. The message was "Disk 0 failed"! I asked him how long the message had been there. He told me it appeared three days prior to this conversation! After another pause, to regain my composure and carefully choose my words, I told him that he owed me a dinner. Why? Because I had used every resource short of physical violence to convince him to buy into a redundant disk system (in this case mirrored) and he had given me an incredible amount of grief on the issue every step of the way! The mirrored disk had saved his company the many thousands of dollars per hour that it would have cost to have the LAN down. I replaced the bad disk on the next Saturday, while encouraging the early reporting of any future "FUNNY" messages. (I never got the dinner, but I did get a reasonably good lunch out of the deal.)

Duplexing is the next step up from mirroring. It means using two intelligent controllers with two hard disks, so that the controllers relieve the file server of the extra burden of tracking the redundant disk operations. Not only does this remove the extra processing overhead - it actually speeds up the system due to the ability to do split reads, in which each read request goes to the disk that is less busy. This performance boost is in addition, of course, to providing redundancy for the controllers as well as hard disks.

The best way to protect the disk channel while getting an incredible boost in performance is to use a RAID system instead of a SLED system. SLED means Single Large Expensive Disk while RAID means Redundant Array of Inexpensive Disks. Note that mirroring is really a simple form of RAID.

RAID devices are divided into categories called levels as follows:

Level 0- Data striping without parity. That means data is spread out over multiple disks for speed.

Level 1- Mirrored disk array. For every data disk there is a redundant twin. Also includes duplexing, the use of dual intelligent controllers for additional speed and reliability.

Level 2 - Bit interleaves data across array, reading using only whole sectors.

Level 3 - Parallel disk array. Disk striping with dedicated parity drives. Drives are synchronized for efficiency in large parallel data transfers.

Level 4 - Independent disk array. Reads and writes on independent drives in the array with dedicated parity drive using sector-level interleave.

Level 5 - Independent disk array. Reads and writes data and parity across all disks with no dedicated parity drive. Uses parallel transfers. Multiple controllers optionally used for higher speed. Usually loses only the equivalent of one drive of array for redundancy. This system is the most popular these days due to both speed and cost effectiveness.

Due to recent speed and efficiency improvements in RAID Level 5, I now recommend it to most my clients. In general, using RAID 5 means that you only lose the capacity of one drive within the array. In an array of seven disks, for example, you would lose only one-seventh of the total array capacity to redundancy (though the actual redundant data is spread across all of the drives - there is no single redundant drive). For an added boost in performance, you can use multiple controllers with a single array, switching from software-based (using the file server's CPU) to hardware-based RAID (using controllers with their own processors) when maximum performance is required.

Using external RAID systems in conjunction with the Secondary Server/Workstation concept mentioned previously results in a very efficient disaster recovery plan. When the main server goes down, merely plug the RAID system into the back of the secondary server, reboot, and you are on line immediately with files that are only minutes old!

An important factor that can affect LAN reliability is power protection. Can I assume that everyone by now knows that you must protect a file server with a high quality uninterruptible power

supply and connect that UPS to the file server with an intelligent communications link? Well, if you didn't know it before, consider yourself informed. It is required!

Now for a little quiz. Can you name the best electrical conductor? (Jeopardy song.) Gold! How about the second best? (Jeopardy song.) Silver! Now, what's the third best? (Jeopardy song.) Copper! Now, can you name the material, excluding fiberoptics, inside LAN cables? (Jeopardy song.) Copper! The question for bonus round is based on that last answer. Can you guess what you call a file server protected by a high-quality UPS on a LAN where all of the workstations and printers are properly protected by high-quality surge-suppressors, and there is one unprotected phone line going into a modem? (Jeopardy song - long version.) The answer is "toast". Power protection works on the weakest link theory. Even one weak link can destroy a system.

If lightning strikes close to a phone cable near your building and induces a current in the line going into that modem, the power can easily go through the modem, into the serial port, into the motherboard, into the network interface card, jump onto the network cable, and spread out over the LAN. Lest you think I am exaggerating the danger, this happened a few weeks ago to a new client of mine. After repeated warnings of impending doom, they finally bought a UPS to protect their file server, but they refused to heed my advice about their inadequate surge suppressors. The day after they installed a new, faster file server, they called to tell me that it had died overnight. After asking for additional clues, they mentioned that a workstation had also bit the dust. It turned out that the building's power had problems overnight. The resulting overvoltages cooked the workstation. The file server was next in line on the LAN cable, and it got fricasseed, too. Please don't make the same mistake. Good power protection is relatively inexpensive to implement and very expensive to omit.

I will continue with the ongoing theme of LAN Design next month. Note that all contestants receive the home version of the LAN Design game as a parting gift.

LAN DESIGN - Part 3

by Wayne M. Krakau - *Chicago Computer Guide*, December 1994 - *NewsWare*, June 1996

Wayne is a partner in Krakau Business Computer Systems, a systems integration firm that is a Novell Gold Authorized and Novell IBM Host Master Authorized Reseller. He has been working with computers for 18 years. He holds an MBA in Marketing and a BS in Information Science and was originally educated in Computer and Information Systems Analysis Engineering. He holds Novell CNI and ECNE ratings and all 5 of the original LANDA Certifications. Wayne can be reached for questions, comments, or topic suggestions at (708)298-7695 or through the Computer Guide.

This article is a continuation of the continuing saga Network Design. Last month I covered fault tolerance, and electrical protection. This month I will begin with backup strategies.

First, plan a reliable pattern of backups. This sounds simplistic, but in many cases it's not being done. I regularly encounter cases of people religiously following a prescribed backup schedule without having a valid backup.

Three way rotating backups with one of the three groups kept offsite is the minimum needed for a reliable system. A simple example of this is a pattern consisting of three weekly sets (Monday through Friday for most businesses) and three of what I usually call "fiscal month" tapes. Each weekly set is used in turn, with one set always kept off site. The fiscal month tapes are also used in turn, every three weeks. While some people think this 18-tape strategy is overkill, even this pattern is something that many mainframe and minicomputer people would sneer at. They are used to much stricter and more thorough patterns. On those systems, mission critical (Don't you just love dramatic-sounding buzzwords?) data is backed up between the individual steps of a given processing task!

(For those of you who may nitpick about my use of the phrase "data is" instead of "data are", please remember that English is a dynamic language, changing to accommodate common usage. Common usage for over 30 years has put "data" in the same category as "herd" and "flock" when used in this context in spite of its Latin roots. In some circles, it is considered an affectation to say "data are" when referring to computer data in this way!)

The second backup consideration is to make sure that you have an appropriate physical backup method. This usually means some type of tape system with a capacity as large as your largest single volume. What's that you say? You can't afford a device big enough to back up your giant disk system? Try this one. An 8GB DAT drive at an approximate street price of under $2100! The tapes cost about $36 each due to the current shortage, but that should drop substantially when the shortage abates. As a bonus, the compression routine is in hardware, not software, so it doesn't revert to half-size (4GB) drive when you attach it to a file server. What's that? You say you say that's not big enough? Try a six-tape autoloader based on the same drive for a street price of around $3600! It even fits in a five and a quarter inch full-height drive bay! There are also much larger specialty systems, including 12-tape and 24-tape autoloaders in the same family. Great bargains are available for smaller systems. No matter what your backup needs, there is an affordable means of reliably backing it up. Realistically, if you can afford 48GB of hard disk space, buying something like this 6-tape autoloader shouldn't be a big problem.

Third, as Nike says in their ads, "JUST DO IT!" Designing a backup system and the logic behind it is useless if the pattern chosen isn't followed. Make it easy to do. Server-based backups can be executed without intervention. For workstation-based backups, assign supervisory rights to a dummy user ID with no password. Restrict that ID to the physical workstation that has the tape. Alter the system login script to automatically turn BREAK off. Execute the backup. Do a verify (also called a compare). Print the results. Use DIR>TEMP1.FIL and DIR TEMP1.FIL to provide time and date stamps both before and after the actual backup. Logout, and then immediately reboot (via the public domain WARMBOOT command) to avoid any possible security problems. If possible, tie the special dummy user ID into the system's automated backup capabilities. If the process is not fully automatable, make sure that the last one out the

door logs in under that ID. Obviously, make sure that someone reviews the results and replaces the tape the next business day.

Communications needs are next. At minimum you need a modem and a remote control program on the administrator's workstation. If you use Windows, obviously, Windows capabilities are required to match the resolution that you use. With this system, when you call for help, you don't have to wait for someone to get to your site, or even to laboriously talk through the problem and its solution over the phone. I had enough of that nonsense when I was beating on mainframes! There is no excuse for it, these days.

Just connect with the remote control program, demonstrate the problem, and then sit back and watch it being fixed. In the worst case, at least the appropriate support person can be sent out with exactly the right tools to fix the problem. The dollar savings on the first call alone can often cover the cost of the modem and the software.

The next step is workstation-to-workstation remote control. From your desk you can support your users and install or upgrade software. In larger sites this is a great time-saver. It also impresses the users - a great help in the political environment that surrounds LANs. If the modem software is either the same brand as the workstation software or at least is compatible with it, a support person can take over the administrator's workstation via modem and then take over any workstation on the LAN to help an individual user. This works well in a multi-site installation. The administrator can call another site and provide support right to the desktop. Note that Netware's RCONSOLE, the program that allows supervisory personnel to take over the main console from a workstation in Netware 3.x and 4.x networks, is compatible with remote control software. For those of you that have specialized communications servers running Netware Connect or one of its competitors, you can get remote control software that will run via your communications server.

Tune in next month for the final (at last) episode of LAN Design and find out if the intrepid systems integrator is felled by the evil Bugmaster!

LAN DESIGN - Part 4

by Wayne M. Krakau - *Chicago Computer Guide*, January 1995 - *NewsWare*, July 1996

Wayne is a partner in Krakau Business Computer Systems, a systems integration firm that is a Novell Gold Authorized and Novell IBM Host Master Authorized Reseller. He has been working with computers for 18 years. He holds an MBA in Marketing and a BS in Information Science and was originally educated in Computer and Information Systems Analysis Engineering. He holds Novell CNI and ECNE ratings and all 5 of the original LANDA Certifications. Wayne can be reached for questions, comments, or topic suggestions at (708)298-7695 or through the Computer Guide.

This article is the final installment in my LAN Design epic. Last month I covered fault tolerance, and electrical protection. This month I will begin with network organization issues.

It is vital to plan the organization of your LAN. For instance, managing and debugging LANs is an arduous task when individual login scripts are used. One small change in the network that would be trivial to accommodate in a system login script can result in a mad scramble to change individual login scripts. Even Novell, after years of teaching that user login scripts are practical, has finally changed their opinion and now strongly encourages the use of a flexible system login script with little or no employment of user login scripts.

People who say that they can't avoid user login scripts are like the programmers that I used to encounter who always came up with excuses as to why they couldn't write a particular routine using structured coding techniques. It's all a matter of mindset. As you get more practice with the system login script, you find that you eventually don't need the user login scripts. The underlying logic becomes apparent.

The workstation-based backup system that I suggested in Part 3 of this series is a good example. I used to implement that type of system as a user login script. Eventually, I figured out how to get around the limitations of the early versions of backup software available years ago and have since used the system login script

Amazingly, many people who were originally exposed to login scripts in old (2.x and below) versions of Netware still haven't figured out the improvements that have been made to the flexibility of login scripts over the years. The addition of an "ELSE" statement (as part of the "IF" logic), the accessibility of DOS environment variables, and the enhanced "ROOT" option of the "MAP" command are the three that I've found the most useful.

One technique that makes using only a system login script easier, is to organize users into groups. Groups also make the assignment of rights much easier. By extensive use of groups, you avoid the temptation to neglect the security aspects of the network. With groups, there is much less overhead involved in administering security.

Planning security is a tricky issue. I have found that the best way is to use the "need-to-know" method. Just assign the rights that are required to get everyone's job done and no more. Every file that someone cannot needlessly change is one less potential "oops" that can be avoided. Also, remember that a virus gets the same access rights as the user who inadvertently executes it. If the user can't change a file, then most viruses can't. In Netware, you can even make a file "execute-only". I have not heard of anyone being able to break that attribute. As an added bonus, it makes files impossible to copy, thereby limiting the potential for piracy.

Another interesting aspect of security can be found in applications programs. I strongly recommend that you carefully analyze all directions and documentation from applications programs regarding rights assignments. The vast majority of software companies, including the giants, are far too liberal. Work out their suggestions on paper and, if possible, experiment to try to find out how much extra protection you can provide without affecting the reliability of the program. You will find that, although

the documentation often instructs you, for instance, to give all rights to all users of a given program, that the program will actually run fine in a very protected environment. Creative use of trustee directory assignments, trustee file assignments, and individual file attributes may be necessary to effectively secure the program's files.

Now comes the "fun" part - documentation! Everything on the LAN should be documented. Anything that's not will have to be discovered on-the-fly, during a debugging session - with the LAN down and users and management screaming at you! If you have to call in outside help, their meter will be running while you try to gather information, or worse, they have to do it. Creating this documentation will also help you avoid configuration mistakes.

All hardware needs to be tracked right down to the smallest detail. IRQs, DMA channels, low memory address ranges, high memory address ranges, CMOS settings, ROM and board version numbers, and all associated details have to be recorded. Serial numbers are also necessary to fill out registration cards. There are many utility programs available to collect hardware and software configuration information. Most are not really comprehensive, so you may have to use multiple utilities. I carry several of them with me.

Much software information has to be tracked. Directory structure, configuration files, login scripts, version numbers, licensing limits, and individual software settings, need to be recorded.

The method used to record this data is immaterial. Paper documentation works fine if it is well organized. A database or even a spreadsheet will do for an electronic record. If you want to get fancy, there are several commercial LAN inventory programs available. Just gather the information any way you can. By the way, if you use a computer to record this stuff, make sure you have a backup so you can get at it during a catastrophic system failure!

Since I just mentioned version numbers, I want to emphasize that the driver software that you get with network operating systems, network interface cards, disk controllers, and just about anything else that needs drivers is usually out of date. Always check with the manufacturer's bulletin board or its forum within an online service to see that you have the latest drivers.

For Netware administrators, a Compuserve account is absolutely vital, since Novell's online support is on Netwire, within Compuserve. When I teach Netware classes, I only half-jokingly tell my students to throw out all of Netware except SERVER.EXE, and to download the rest.

I also want to warn you that using VLMs, or Virtual Loadable Modules, may be hazardous to your network’s health. Make sure that you test all of your applications thoroughly before switching. I discovered one client in the middle of uninstalling Netware 4 and reverting back to version 3.11 because too many of their applications were incompatible with VLMs. Netware 4.01 was not very useful for them without VLMs, so they just pulled the plug. Also, remember that even though Netware 3.12 specifies VLMs, it doesn’t really need them. It will run just fine with the older ODI Shell. I have several clients who have switched back to ODI from VLMs after upgrading to version 3.12. I expect that both compiler publishers and Novell will eliminate these incompatibilities over time. Until then, make sure you test prior to your implementation.

In this series, I have given you some ways to prepare for troubleshooting and to try to avoid the need for troubleshooting. By implication, I hope that you realize that I chose these suggestions based on my observations as to what kind of problems actually occur. To do the troubleshooting, you need to look at these proactive methods and probe for weaknesses in their implementations through what can be best be described as detective work. By sneaking this information to you via a discussion of management and planning, I hope that you will be better prepared to do your detective work. Good hunting!

RISING STARS

by Wayne M. Krakau - *Chicago Computer Guide*, February 1995

Wayne is a partner in Krakau Business Computer Systems, a systems integration firm that is a Novell Gold Authorized and Novell IBM Host Master Authorized Reseller. He has been working with computers for 18 years. He holds an MBA in Marketing and a BS in Information Science and was originally educated in Computer and Information Systems Analysis Engineering. He holds Novell CNI and ECNE ratings and all 5 of the original LANDA Certifications. Wayne can be reached for questions, comments, or topic suggestions at (708)298-7695 or through the Computer Guide.

Whew! Another trade show season is over. I like to measure the quality of a trade show by how much you have to bend over to scratch your knee. For the best shows, the quantity of valuable literature is enough so that you don't have to bend at all! This is, of course, assuming that you remember to switch your bag from side to side on a regular basis. If you don't, then you end up doing Marty Feldman impressions (as he appeared in "Young Frankenstein"). Of course, even if you do remember, you may be eligible to be an extra in "Gorillas in The Mist, Part 2" - without using any prostheses.

One product I spotted at these shows that absolutely fascinated me is the Pereos tape drive from Datasonix (Boulder, CO, 303-545-9500). It is a ten-ounce tape drive (the whole drive, not just the tape) that holds approximately 1.25GB after compression. That weight includes the two AA batteries that power it! (An AC adapter is also available.) The tapes are about the size of my thumbprint and about one-fifth of an inch thick!

This drive uses that same basic technology as the microminiature personal note-taking recorders made by Sony. It uses a non-tracking helical scan recording method (essentially an elegant simplified version of what 4MM and 8MM tape drives use) with a very low theoretical error rate. Lest you worry about saving data using a technology derived from the consumer recording industry, remember that 4MM DAT tapes evolved from Digital Audio Tape standards and 8MM tapes are a direct offshoot of 8MM video cameras.

The current version uses a removable parallel port interface, with PCMCIA and SCSI interfaces in the planning stage. The drive is a separate module that plugs into the interface, so one drive can be used with multiple interfaces.

My initial reaction to seeing this drive was disbelief. It looked too small to hold 1.25GB! The salesman had a little plastic box in his shirt pocket, smaller than a cigarette case (but not as deadly). He pulled it out and handed it to me and said "Here, have 8GB!" The box contained eight of the tiny tapes, each in its own slot. I probably could have carried six of them in my shirt pocket. This brings new meaning to the term "Offsite Backup".

The Pereos is currently being marketed with an eye toward people with laptop computers or both desktop and laptop computers. It has features in its software that allow fuzzy search for file names, file synchronization between computers, and the ability to restore to a specific date and time. (Sherman, crank up the Wayback machine!) These extra features (and its price - $695) should make the Pereos popular in many other market niches. The only thing that they are missing (for now) is a way to backup Netware Rights and Bindery information.

Another item of more immediate use in Netware environments is Vinca Corporation's (Orem, UT, 801-223-3100) StandbyServer. This product allows you to configure a backup server that is available at any time with just a few keystrokes, using the clever device of tricking Netware into thinking that the second server is simply a mirrored (that is a duplicate) disk drive. Because of this trick, the only restricting requirement is that the sizes of the Netware volumes being mirrored are identical. All other aspects of the servers can be varied.

You can think of this system as a sort of "SFT III Lite". Where Netware SFT (System Fault Tolerance Level) III automatically switches to the secondary server upon failure of the primary, the

StandbyServer requires a few manually entered keystrokes to confirm the demise of the primary and give permission to start the secondary server.

This disadvantage is balanced by the compatibility and flexibility of the StandbyServer. SFT III has serious restrictions on the type of NLM (Netware Loadable Module) that can be run. Any NLM that directly interacts with the underlying hardware of the server is probably ineligible for use on an SFT III system. The StandbyServer doesn't care about the type of NLMs that it runs. It doesn't try to match the two servers on a CPU (Central Processing Unit) cycle by CPU cycle basis. It merely mirrors the disks and coordinates some surface aspects of the two servers' operations.

Since the secondary server runs as a sophisticated DOS program, only a single copy of Netware is active at any given time. This means that you don't need a second copy of Netware. StandbyServer comes bundled with two MSL-compatible (Mirrored Server Link) cards (from $1695 to $1895 depending on bus structure). Bundling reduces the possibility of the type of multivendor recriminations that impeded my original implementation of Netware SFT Level III, and maintains the possibility of upgrading to SFT III in the future.

Remember this product the next time you replace your server. It may be a good idea to keep the old one around and use StandbyServer to provide backup server capability.

For the security-conscious, Veritel Corporation of America (Mt. Prospect, IL, 708-670-1780) offers the Logon Verification System and Caller Verification System. I've been interested in security verification systems for several years. Historically, they use a combination of a personal code or password, a centralized software or hardware based challenge generator, and a calculator-sized response generator (that often doubles as a real calculator).

Typically, the user enters a login ID. The challenge generator replies with a numeric or alphanumeric challenge code. The user enters this challenge code and a password into the calculator. An algorithm (Don't you just love these buzzwords? I could have used "formula" but I have to put buzzwords in wherever I can to comply with the minimums set by the Computer Writers' Guild.) inside the calculator derives a response - assuming, of course that the password is correct. The user then enters the response code on the keyboard. If everything works right, only someone who both discovers the password and steals the calculator/response unit can gain unauthorized access.

In practice, it's a royal pain. If you make even the slightest error in transcribing either the challenge or the response, you don't get in - this is assuming you have the calculator/response unit with you. No unit - no access.

Veritel has come up with a way to eliminate the extra codes and the response unit from the equation. They use your own voice as your entry key. (They were probably frightened by the self-destruct sequences in the original Star Trek series as children - or at least by the simplicity of the codes required.) It recognizes your voice and key phrases as necessary to allow access to computer systems, PBXs, voice mail systems, and even the internal 800-number dial-out systems used by many corporations.

The system adjusts well to day to day variations in voice quality so that allergies, a cold, or a sore throat is not likely to confuse it. It is designed with an algorithm (another point towards my quota) that detects the waveform distortion generated by even the most sophisticated recording system, so recording someone's voice won't get you past its security. (Is it live - or is it Memorex?) It can also adjust for phone versus direct microphone input so the two are not interchangeable. It even supports fallback to touch-tone codes if you lose your voice (and, of course, if your system administrator activates that function).

This is the ultimate in easy to implement security. Try it the next time you get a multi-thousand dollar fraudulent bill on your outgoing 800 lines.

Now that we've explored these rising stars, assume the lotus (not Lotus) position, close your eyes, and recite our new mantra - "Netware 4.1 lives, Netware 4.1 lives, Netware 4.1 lives . . . "

MORE RISING STARS

by Wayne M. Krakau - *Chicago Computer Guide*, March 1995

Wayne is a partner in Krakau Business Computer Systems, a systems integration firm that is a Novell Gold Authorized and Novell IBM Host Master Authorized Reseller. He has been working with computers for 18 years. He holds an MBA in Marketing and a BS in Information Science and was originally educated in Computer and Information Systems Analysis Engineering. He holds Novell CNI and ECNE ratings and all 5 of the original LANDA Certifications. Wayne can be reached for questions, comments, or topic suggestions at (708)298-7695 or through the Computer Guide.

Yes, it's time for another exciting episode of the Trade Show Tango. My desk is overflowing with imaginative new products and major revisions to old ones, so I decided to devote another column to the rising stars of the bunch.

Before I mention the first product, a technical explanation is needed. Switched Ethernet is a technique of increasing the amount of data that can be handled by an Ethernet network. Ethernet is essentially one big party line. (Remember the hand-cranked phones in the original Lassie Series?)

When one computer is about to send a message, it first listens to the party line to hear if anyone else is already sending one. Assuming the line is free, it tries to send its message. When it talks, all of the other computers hear it, not just the intended receiver. If more than one computer tries to talk at the same time, the other computers notice the increased-strength signal caused by the crisscrossing original signals. The first one to notice this signal increase puts out a special jam signal to tell everybody to back off and try again. (Think of the buzzer that sounds off when someone answers a question incorrectly on a game show.)

When the other computers hear the jam signal, any of them that were in the middle of trying to send a message stop immediately, wait for while, and then try to send again. The definition of "a while" is determined by a random number generator built into the Ethernet card.

For those who aren't fans of programming, a random number generator gives a different result number for each "seed" number that it is fed. For this process, the seed is the Ethernet identification number of each network card. Barring a major manufacturing error (at least one of these errors has occurred), every Ethernet card ever made has a different Ethernet ID. A different ID means a different seed number and, therefore, a different result number. Since the result number determines the definition of the "while" that each card waits before resending its message, each card will wait a different amount of time before resending, thereby avoiding another simultaneous send.

This process is complicated by the fact that each computer's ability to resend is affected by its own speed and by the relative efficiency of its Ethernet card, but in practice it works fairly close to the theoretical model. The overall process is called CSMA/CD, meaning Carrier Sense Multiple Access with Collision Detection. (Love those buzzwords!)

Ethernet provides a theoretical maximum of ten million bits per second (where eight bits make a byte which is the equivalent to a single character), but this is shared via the aforementioned process by all of the computers on the network. Switched Ethernet is a sneaky way around this sharing problem.

Switched Ethernet uses the same basic principles used by the phone company before fully automated switching was available. Upon request, an operator would have to manually place a jumper cable between the sockets leading to the sending and receiving phones, creating a temporary complete circuit. (Think of Lily Tomlin's Ernestine character.) After the call was complete, the jumper would be removed.

In the modern version of this example, a temporary circuit is created between the sending and receiving computer so fast that the two essentially have the entire ten million bits per second to themselves for just long enough to get the message through. Then the line is freed up for the next pair of computers to communicate.

If you make the assumption that this particular operator in our example is handling the connection point between a single trunk line and multiple local lines, the comparison with switched Ethernet on a LAN is more accurate. (Though I am sure that techies reading this are already squirming in their seats over this gross oversimplification.) Since most of the traffic on a LAN is between the workstations and the file server, the line to the file server becomes the "trunk line" in this example. It is the line through which most data must pass and is the one that we must optimize the most.

This leads me (finally) to the first of the hot new products, a clever idea developed by XNET Technology (Milpitas, CA, 408-263-6888). They have developed a product that they call a ParallelSwitch, which encapsulates an Ethernet switching hub on a single ISA (Industry Standard Architecture) or EISA (Extended ISA) board. It contains one basic Ethernet chipset to talk to the file server's motherboard (the trunk line) and six switched Ethernet ports to service the network (the local lines).

While switching is not an original concept, the idea of putting what amounts to a complete switching hub on a single board is. At $1,988 for ISA and $2,988 for EISA 10Base-T models, this product is a major price breakthrough in this category. In addition, it is arguably more efficient than an external switching hub since the switching circuitry is on the same board as the servers own Ethernet chipset.

Each time the file server communicates (in either direction) over one of the six switched lines, it momentarily sees the full bandwidth (essentially, how much room for signals) of Ethernet. This raises the theoretical effective throughput (how many signals really gets through) to sixty million bits per second!

Of course the real world will intrude and keep us from getting the theoretical ideal, but we can maximize our throughput by selecting EISA as opposed to ISA. With an EISA interface, we are less likely to run into bottlenecks in throughput due to limitations within the underlying bus architecture. Within those limitations, XNET's solution to a maxed-out Ethernet is very useful, especially for the network administrator with a limited budget.

The second product on my list is really two similar products, 3Pack from Funk Software, Inc. (Cambridge, MA, 800-828-4146), and NetSqueeze from the LAN Support Group (Houston, TX, 800-749-8439). Both are file compression NLMs (Netware Loadable Module) for Netware 3.11 and 3.12 file servers. Both mimic the configurable compression features built into Netware 4.x, even to the point of using licensed STAC compression.

These products provide a way for an administrator to take advantage of compression without making the transition to Netware 4. While I suspect that both companies may be getting hate-mail from Novell for encouraging users to delay upgrading, I have to admit that I regularly find companies that are not quite ready - either technologically or financially - to upgrade. Aftermarket compression products are a method of letting an old system limp along until the inevitable (in my opinion) upgrade becomes feasible.

The last two products on my list are actually upgrades to ones that I have previously praised. Micropolis Corporation (Chatsworth CA, 800-395-3748) has made major improvements in its Radion series of RAID (Redundant Array of Inexpensive Disks) systems. They have designed a controller that they call Gandiva, to give hardware RAID to their systems. You can even start out with software-based RAID and move up to hardware-based later. This improvement will make their Radion Series much more competitive with the high-end RAID systems currently available, as well as providing flexibility due its upgrade capability.

Another enhancement that Micropolis introduced is a rack-mounted version of its RAID system. This RAID will allow expansion to 81GB using 3.5" drives or an amazing 245GB using 5.25" drives! The system should be handled by the Gandiva controller to enhance speed (though technically you could use software). This should keep staunch mainframe advocates up worrying all night!

The final product is the new CommSwitch Series 2000 scalable computer system from CommVision (Mountain View, CA, 415-254-5270). I have written about CommVision's (formerly IWI) eight and twenty slot computer systems before. Now they have a rack-mountable very expandable version

available. They have versions with up to 180 32-bit, hot-swappable slots! Up to 160 of these slots can be filled with processor boards. Talk about multiprocessing!

Since Netware is already divided into segmentable pieces called NLMs, it is quite easy to split Netware's load (using Run-Time Netware) among many processors without the efficiency bottlenecks and reprogramming necessary to fully implement symmetrical multiprocessing systems. Each of these processors can be a 486 or a Pentium with up to 128MB RAM for the 486s and 256MB RAM for the Pentiums. Each board is hot-swappable and has a full 32-bit, high-speed interface. The overall packet-switching bus that connects these boards has an aggregate peak transfer rate of up to 8.5Gb (that's Gigabits) per second!

While CommVision is billing this system as the ultimate communications server, I would suggest combining it with hardware RAID (maybe from Micropolis) and turning it into an incredibly fault tolerant superserver. I suspect that it would run circles around conventional superservers with only four or so processors, and throw in the added reliability as a bonus.

Finally, I just wanted to make it clear that there is no truth to the rumor that the Justice Department has been purchased outright by Bill Gates in return for settling the national debt.

BUS WARS

by Wayne M. Krakau - *Chicago Computer Guide*, April 1995 - *NewsWare*, August 1996

Wayne is a partner in Krakau Business Computer Systems, a systems integration firm that is a Novell Gold Authorized and Novell IBM Host Master Authorized Reseller. He has been working with computers for 18 years. He holds an MBA in Marketing and a BS in Information Science and was originally educated in Computer and Information Systems Analysis Engineering. He holds Novell CNI and ECNE ratings and all 5 of the original LANDA Certifications. Wayne can be reached for questions, comments, or topic suggestions at (708)298-7695 and via E-mail to wayne@mcs.com.

ISA, MCA, EISA, VL, PCI - it's alphabet soup time again! First let's get to the definitions. A bus (it's supposed to be "buss", but I'll bow to current common usage) in a PC is the slot into which you plug circuit boards to add functions.

ISA means Industry Standard Architecture. In English, that means a specification taken from reverse-engineering the original IBM AT bus. (Since IBM never published a thorough description, reverse-engineering was required to derive this specification.) It is a limited bus, but it served the industry well as a standard for growth. It is now only used for limited speed applications such as simple serial and parallel ports.

MCA originally meant Music Corporation of America, but, after some legal wrangling (after the fact) by IBM lawyers, it also became the legitimate abbreviation for Micro Channel Architecture. This bus was IBM's attempt to create another proprietary standard, and, basically, screw up the PC industry for their own benefit. (Remember FUD - Fear, Uncertainty, and Doubt, IBM's unofficial strategy for larger computers?) The only thing it screwed up was IBM's PC division (which is taking massive financial losses). Even IBM has given up on this architectural dead-end.

EISA means Extended ISA. EISA was invented by a committee of manufacturers to combat IBM's MCA. It has done very well in file servers and other specialty machines, but has not made many inroads in the regular PC (single-user and network workstation) market. If you need a machine full of high-speed slots, such as a file server, EISA is currently the only choice.

The VL bus is another committee design. It was meant to be THE next generation of bus, but has proven to be only an interim step. While its base performance is adequate (it often can outgun its main competitor, PCI), its lack of expandability has doomed it. Each additional VL slot puts more stress on the overall computer system, so your won't see computers with more than 3 VL bus slots. This makes it acceptable for now, but not for the future.

The PCI bus is THE bus for the future. Even Apple has chosen it! The problem is that "future" is the operative word. Though it is theoretically possible to stuff a machine full of PCI slots, current Intel-based motherboards usually support only two. This limitation is built into the existing Intel chipsets. Some manufacturers have decided not to wait for Intel's new chipset which will allow three PCI slots (still insufficient for servers) by tweaking their current Intel-based systems to support three. This means that anyone buying these computers is gambling on a kludgy modification to the Intel's design. ("Kludgy", as in "kludge", an ad hoc change or patch - a quick, inelegant fix.)

Some motherboard manufacturers have rejected Intel's chipsets completely and have designed their own. The viability of the resulting motherboards is dependent upon their accurate interpretation of the PCI specification - essentially an Intel invention. These manufacturers offer multiple PCI slots and even combinations of EISA and PCI in a single machine.

PCI has more problems. It has been said (with admitted exaggeration) that when adding a PCI board to a system that already contains one, you have about a 50% chance of having the combination work. Supposedly, this is not due to the PCI specification being inherently bad, but to board designers misinterpreting the definition! Great! That's just what I want to hear when I'm on the phone to a technical support person while an anxious client looks over my shoulder, expecting me to perform miracles. All I

need is to have one side say that the definition is ambiguous and the other side say the definition is precise, but that the interpretation is faulty. (Less filling - Tastes great!) I DON'T CARE! I just want these people to get their act together and make products that will work together every time.

Another PCI problem is with communications between the PCI interface circuitry and the Pentium interface circuitry on certain motherboards based on older Intel chipsets. Transparent data changes have been reported in these motherboards in the national press. "Transparent" in this case means data is being changed without any warning! That's more scary than an outright system failure! I'm sure that the IRS will be thrilled to hear that your company is reporting a year-end total of $1.25 worth of withholding for your 100 employees! (Maybe this will scare people into using the integrity checking routines built into almost all modern accounting software.)

Note that this problem is separate and distinct from the more widely reported errors inside Pentiums. Unfortunately, I haven't found a specific reference as to the exact model numbers or precise vintage of these "older" chipsets and motherboards.

An additional, though temporary, problem is the lack of availability of specialized server-oriented boards for PCI. These devices are currently made only for EISA systems. Luckily, that will change over time.

My current bus selections revolve around the limitations and real-world conditions that I have listed here. For servers, I usually specify EISA. I am in the process of switching to combined EISA/PCI systems, and will eventually use pure PCI.

The reason I am willing to use the "combo" systems, is that I will only work with PC manufacturers who design and build their own motherboards, regardless of whose chipsets they use. This gives me a way out when any question of compatibility arises. For example, my research into problems with the original implementation of Novell's Netware SFT Level III (System Fault Tolerant, Level 3, using dual servers), was extremely shortened because I was able to have a discussion of EISA bus timing issues with one of the people who designed the motherboards that I was using. That conversation provided enough information to immediately eliminate motherboard compatibility from my list of possible errors and get on with researching other possibilities.

I am still selling mostly VL-based workstations, but am gradually increasing the use of PCI. I have never had (or even heard of) problems with the VL bus. You just put a board in, and it works - fast. If future expansion is vitally important and the client understands the issues mentioned above, then I am willing to sell PCI.

There are some additional motherboard issues that need to be brought forward. The first concern is that there are no standards for "sideways" slots of any type. Sideways slots use some type of riser card mounted in the motherboard using either a proprietary slot or a tweaked standard slot. The riser itself contains slots. Circuit boards are placed in these slots parallel to the motherboard. Machines with sideways slots have a long history of compatibility problems and are usually difficult and expensive to work on, so I discourage people from using them.

The second motherboard issue is the number of slots. The trend in computing is to add multiple complicated hardware-based features. These can include sound boards, accelerated video cards, speech cards, telephony-related cards, and many more. Because of this, I prefer to use computers with the old standard of eight slots. This leaves room for expansion and provides a way out if you need to replace the motherboard due to malfunction, obsolesence, or compatibility problems. This standard size motherboard can be replaced by almost any other standard-sized motherboard. The lack of a built-in video controller, disk controller, or even input/output ports means that in case of failure or obsolescence, you don't have to purchase an overpriced proprietary motherboard when all you need is a single inexpensive component.

And you thought the only bus you had to worry about had a large picture of a skinny dog on its sides!

SPEEDY ALKA-SELTZER

by Wayne M. Krakau - *Chicago Computer Guide*, May 1995

Wayne is a partner in Krakau Business Computer Systems, a systems integration firm that is a Novell Gold Authorized and Novell IBM Host Master Authorized Reseller. He has been working with computers for 18 years. He holds an MBA in Marketing and a BS in Information Science and was originally educated in Computer and Information Systems Analysis Engineering. He holds Novell CNI and ECNE ratings and all five of the original LANDA Certifications. Wayne can be reached for questions, comments, or topic suggestions at (708)298-7695 or via E-mail to wayne@mcs.com.

No, this is not a nostalgic tribute to classic advertising campaigns. (Besides, if I remember correctly, I was still wearing pajamas with feet when that campaign was in its prime.) The "speedy" in the title refers to LAN speed at the physical level - just how fast can you pump data down the pipe. The "Alka-Seltzer" refers to the medication that LAN administrators (and systems integrators!) are going to need when they find out that their already complicated lives just got a lot more complex.

FDDI (Fiber Distributed Data Interface), a 100Mbps (millions of bits per second) descendant of Token Ring, has been available for several years, but is still too pricey at $1300 and up per card to make much of a dent in the workstation connection market. It has been used for high-speed backbones, but is still a bit expensive even for that.

ATM (Asynchronous Transfer Mode, not Automated Teller Machines) provides a 155Mbps link, but it is too new to be in widespread use and makes FDDI look inexpensive. IBM's special 25Mbps version, called ATM25, is seen as being too immature, even though it is in the process of being approved as a secondary standard. While it is less expensive than the 155Mbps version of ATM, it suffers in comparison with other methods in a cost per megabit comparison.

Administrators of Ethernet LANs are being swamped with new options to wring more speed out of a network connection. The real challenge is to decide amongst them.

The most spectacular methods are two that involve raw speed, 100Mbps versus the 10Mbps Ethernet standard. These two methods are called 100Base-T and 100VG-AnyLAN, and are designed to work, initially, over unshielded twisted pair (UTP) copper wire. If you know that Ethernet, when using UTP, is called 10Base-T, you can get an idea of these names' derivations.

Each of the two competing standards (or should I say "almost" standards, since as I write this, neither is technically finalized) is supported by its own consortium of manufacturers, sixty-two for 100Base-T and twenty-six for 100VG-AnyLAN, with five companies straddling the fence by supporting both types (by my count in mid April 1995). There was supposed to be only one standard called "Fast Ethernet", but the political maneuvering got so messy that we ended up with two.

100Base-T is the true descendant of 10Base-T, so it inherited the moniker "Fast Ethernet". It follows the same basic methods as 10Base-T, in particular CSMA-CD (Carrier Sense Multiple Access with Collision Detection), its way of arbitrating which computer gets access to the LAN. The main difference is that 100Base-T's raw speed is ten times as high and the maximum span of the LAN is correspondingly only one-tenth of its predecessor's. That gives you between 200 and 210 meters to work with, depending on the type of concentrator or hub that you use. You are also restricted to only two repeaters or hubs.

Currently, Fast Ethernet works only with Level 5 (the highest quality) cable using two pairs of a four-pair bundle. A standard (Yikes, more standards!) for using the older, more common Level 3 UTP cable by utilizing all four pairs of a four-pair bundle is on the way. Concurrently, the 100Base-FX standard for fiberoptic cable is being finalized. Using fiber segments will help alleviate the problems caused by the severe distance limitations of Fast Ethernet.

As they say on *Monty Python*, "And now for something completely different." The other standard, 100VG-AnyLAN, really doesn't have a lot to do with the original Ethernet standard. It is a new hybrid design with bits of Ethernet, Token Ring, and various other methodologies. It maintains the typical

Ethernet distance limitations and can work with Level 5 (two pairs) or Level 3 (either two or four pairs) UTP cable, and soon, with fiberoptic cable. Because VG eliminates potential collisions within its hubs, it tends to do better in applications where a continuous stream of data is needed, such as video or sound transmission. Its "Ethernet" name is basically a marketing ploy, implemented to make this new standard more palatable to LAN managers and upper management types. (Note that the preceding sentence is a disputable political statement of my opinion. It's considered the "purist" point of view.)

An older technology that is just hitting its stride is Ethernet switching technology. For a detailed explanation of this technology see my *More Rising Stars* column from two months ago. In summary, the hub gives the full Ethernet 10Mbps of bandwidth to one port at a time by switching between ports at very high speeds. At the other end of those ports, you can put other concentrators, or for maximum throughput, wire them directly to individual workstations. Keep in mind, however, that Ethernet switching is totally proprietary - each vendor has their own unique way of doing it.

Yet another technology (Hey, I told you this would make your life more complex!) is Duplex Ethernet. It allows simultaneous two-way communications over the network cable. This concept is similar to the comparison between half-duplex and full-duplex within your communications software. This is advantageous when there is a lot of data being sent by workstations. Since most networks have 95% or more of the data coming from the file server, this isn't necessarily a great advantage. It is advertised as providing a theoretical 20Mbps channel, but your mileage may differ. Just to keep things interesting, Duplex 100Base-T is on the way. It raises the ante to a theoretical 200Mbps. Luckily for those evaluating this technology, Duplex Ethernet is interoperable among different brands of hubs and network cards.

Finally, (Now that I've got your head really spinning.) multiport Ethernet cards for file servers are going to get very popular due to the potential speed of the PCI bus and the burgeoning 100Base-T market. Manufacturers are developing four-port PCI Duplex Fast Ethernet cards with a theoretical aggregate throughput of 800Mbps! Each of these cards takes up only one slot in the file server, so you can use more than one. Network card design has advanced to the point where several of these cards placed in one file server are efficient enough that they won't bog down the filer server's processor.

Once all of the hardware is available, a maximized system could combine these technologies, but the best use of these technologies is probably in some combination based on the availability of new products and tempered by real world cost considerations. For example, I am currently designing a system in which a compromise solution has the best price-performance ratio. It will include two four-port 100Base-T PCI cards leading to separate backbone lines to each wiring closet. Within each wiring closet there will be an Ethernet switching hub with a single 100Base-T port on its back and multiple 10Base-T ports on its front. The workstations will use standard 10Base-T cards. One great advantage of this design is its expandability. With the 100Mbps backbone lines, we can upgrade to 100Base-T concentrators if they're ever needed in the future.

Now that I've finished covering these new toys, I'm going to the drug store to stock up on more of my favorite analgesics and antacids.

DÉJÀ VU

by Wayne M. Krakau - *Chicago Computer Guide*, June 1995

Wayne is a partner in Krakau Business Computer Systems, a systems integration firm that is a Novell Gold Authorized and Novell IBM Host Master Authorized Reseller. He has been working with computers for 18 years. He holds an MBA in Marketing and a BS in Information Science and was originally educated in Computer and Information Systems Analysis Engineering. He holds Novell CNI and ECNE ratings and all five of the original LANDA Certifications. Wayne can be reached for questions, comments, or topic suggestions at (708)298-7695 or via E-mail to wayne@mcs.com.

It's here! It's new! It's exiting! It's wild! No, I'm not talking about the latest pheromone-spiked fragrance, created to have members of the opposite sex buzzing around you like lovesick fruitflies. I'm talking about Netware 4.1, the latest offering from the 600-pound gorilla of the networking industry, Novell. This version is the real thing - finally. Netware 3.12 has been relegated to has-been status. If you are thinking of upgrading your network server, either hardware or software, this is the time to jump on the Netware 4.1 bandwagon. (Let's see now, how many bad marketing clichés can you find in this paragraph?)

If all of this sounds familiar, then you probably remember the release of Netware 3.11, the first stable release of the Novell's 3.x series and the subsequent controversy over Netware 2.2, the last of the 2.x series. While Netware 2.2 wasn't officially declared dead until mid-1994, it was released after Netware 3.11 and, therefore, doomed from birth.

Netware 3.12, in contrast, had a period of usefulness. Versions of Netware 4.x prior to 4.1 were incomplete and not nearly as reliable as 4.1. Many programs would not run properly under Netware 4.1 or under VLM's (Virtual Loadable Module's), the workstation software that is required if you want to access all of the features of Netware 4.1. This left some companies with no upgrade alternative other than Netware 3.12.

Even today, if you are stuck with an old application that is explicitly incompatible with Netware 4.1, your business needs this specific application to survive (you can't switch to another brand of software), and the software publisher isn't smart enough to modernize or is not skilled enough to convert the program, then you may be stuck with Netware 3.12. (This might be a good time to reevaluate potential replacement software.) Also, since Netware 3.12 is merely a "cleaned-up" 3.11, you might end up just staying (for the moment) with 3.11.

As a side issue, prior to the release of Netware 4.1, I only did 3.11 to 3.12 upgrades for those clients who were going to purchase a complete replacement server. Upgrading Netware at the same time requires little added labor, since we would be simultaneously upgrading all of the driver software anyway. (Note that these "drivers" are intermediaries that allow communications between the hardware - network cards, disk controllers, etc. - and the system software - in this case, Netware.) I have not suggested upgrades for other clients, since versions 3.11 and 3.12 are so similar.

This situation has inspired me to rewrite my original litmus test for choosing a reseller. The original test covered Netware 2.2, while this test covers Netware 3.12 (subsequent to the release of Netware 4.1).

If, (excepting the compatibility problem mentioned above) through advertising or personal contact, a reseller has proposed to sell Netware 3.12 (or worse, an earlier version), one of the following is true:

A. The reseller is so out of touch with the real world that he or she doesn't even realize that Netware 3.12 is a complete dead-end, long relegated to second class status by software and hardware manufacturers, and more recently, by Novell. If they are that out of touch with industry conditions, why deal with them?

B. The reseller knows of Netware 3.12's obsolescence, but wants to get rid of leftover stock, and has no ethical problems with tricking a client into buying an inappropriate product. If they are that unethical, why should you deal with them?

C. The reseller knows of Netware 3.12's obsolescence, but wants to make extra money by making the customer upgrade twice, first to 3.12, and later to 4.1, with all of the attendant duplication of labor. This reseller has no ethical problem with “churning” an account (at least that’s what they call it when a stock broker performs extra transactions to get more commissions - this concept is similar). If they are so unethical, why should you deal with them?

Novell has previously used a carrot and stick approach to accelerate the natural trend of market forces when trying to persuade customers to upgrade to the latest version of Netware. In the past, this has included limited-time upgrade discounts on the carrot side and implied or explicit promises of lack of future development on the stick side.

Novell is taking a similar tack in the push for Netware 4.1 acceptance. It announced an official policy of neglect for the Netware 3.x product line - and then backpedalled, the same pattern they originally followed with the 2.x line. Novell has also decreased the price on Netware 4.1 to match those of 3.12, eliminating the direct financial incentive for purchasing Netware 3.12.

For the more cynical among you, please note that there are patches available for those remaining reported errors within Netware 4.1 (including the infamous compression slowdown problem) on Compuserve (GO NETWIRE) and the Internet (http://www.netware.com). Of course, Netware 3.12, Netware 3.11, and even all of the Netware 2.x patches and fixes are also available from these sources.

The dictionary defines deja vu as the feeling of having lived through the same event previously. I think this situation qualifies.

LEMMINGS 95?

by Wayne M. Krakau - *Chicago Computer Guide*, July 1995

Wayne is a partner in Krakau Business Computer Systems, a systems integration firm that is a Novell Gold Authorized and Novell IBM Host Master Authorized Reseller. He has been working with computers for 18 years. He holds an MBA in Marketing and a B.S. in Information Science and was originally educated in Computer and Information Systems Analysis Engineering. He holds Novell CNI and ECNE ratings and all five of the original LANDA Certifications. Wayne can be reached for questions, comments, or topic suggestions at (708)298-7695 or via E-mail to wayne@mcs.com.

Lemmings - "Small animals of northern latitudes . . . noted for occasional migrations . . . many thousands of animals . . . always pushing on until they enter the sea and are drowned." (From *The Little & Ives Complete Book of Science*.)

Lemmings 95 - Large animals (two-legged) of multiple latitudes . . . noted for occasional mass operating system purchases . . . many thousands of animals . . . always buying until their systems crash. (From *Chairman Wayne's Little Red Book of Semi-Coherent Rave-Outs* - yeah, right.)

You don't have to read Nostradamus to predict this upcoming disaster. Millions of users purchasing a new operating system and blindly running the install program. Windows 95 has been the subject of what will probably become known as the greatest marketing campaign in history. It's the Hype-Monster that ate New Jersey! The people at that marketing juggernaut Procter & Gamble are probably weeping in envy.

Has everyone forgotten the old (in computer industry terms) adage about not buying version 1.0 (even if they name it 4.0, since it is a descendent of Windows 3.11) of any software - especially operating system software? Have they forgotten the history of the original versions of DOS, OS/2, and Windows NT? Lest you think that I am anti-Microsoft (or worse yet, a Luddite infiltrator within the computer industry), let me place NetWare 1.0, 2.0, 3.0, and 4.0 into evidence. (Hmm - perchance just a bit too much O.J. trial viewing?)

Remember, that this use of initial NetWare versions as bad examples is from a person known for being an enthusiastic (some would say rabid) supporter of Novell. As a matter of corporate survival, Novell was forced to publish NetWare 4.0 in spite of bugs, incompatibilities, and the last-minute removal of advertised features. NetWare 4.01 and 4.02 consisted mainly of patches and bug fixes. It wasn't until the issuance of NetWare 4.1 that a stable, full-featured version was available. (This brings back memories of the NetWare 2.0 to 2.0a and NetWare 3.0 to 3.11 transitions.) On ethical grounds, I discouraged clients from jumping on the NetWare 4.0 bandwagon without giving serious consideration to its aforementioned (Oops, another O.J.-ism.) limitations. NetWare 4.1 is, of course, another story. I have become an ardent proponent of that network operating system.

In Microsoft's case, the company's reputation and future enhanced profitability, rather than its survival, are at stake. (I am convinced that Microsoft will go out in a blaze of glory when Bill Gates buys several of the former U.S.S.R. republics, uses their remaining missiles to launch a nuclear first-strike against Orem and Provo, Utah [Novell's headquarters] and Armonk, New York [IBM's headquarters], and is killed by an accidental retaliatory strike from U.S. forces after a Washington bureaucrat misfiles a newly passed [strictly along party lines, of course] bill exempting nuclear warfare from antitrust prosecution - and retaliation.)

Along with the obvious, absolutely gushing, free publicity that Windows 95 is getting, much of the hype is driven by myths. The first myth is that good ol' Windows 3.1x is inherently unstable and Windows 95 is rock-solid stable. While I will admit that it is easier to crash 3.1x than 95, the only unstable Windows 3.1x systems that I encounter are due to the improper configuration of DOS, memory managers, add-in cards, or Windows itself. No operating system (or for that matter application or utility

software) will run optimally by just executing the install program. It's even worse for hardware. If you want "civilians" or computer professionals from other arenas to install and maintain their own systems, you just have to accept the risk that they won't do it right. It's not their fault - they just don't have the appropriate level of systems integration knowledge.

As to the alleged stability of Windows 95, that is based on the belief that Windows 95 is a 100% 32-bit fully protected operating system. Sorry, folks - it's not. The only way Windows 95 could provide a relatively high level of compatibility (though not complete) with old applications was to keep portions of our old 16-bit friend, MS-DOS. This has been repeatedly documented in the national press (despite denials by Microsoft), so I won't get into the technical details. Many specific methods of crashing Windows 95 have also been published, providing practical examples of this theoretical weakness. Please do note that even a real full 32-bit "protected" operating system can crash, anyway! There is no magic involved. If you live by the bug you die by the bug. (Alright, stop groaning!)

Now, for the tricky questions. Do I believe that Windows 95 will bulldoze the competition into the ground? Yes. Do I believe that it is the proper successor of Windows 3.1x? Yes. Do I believe Windows 95 has valuable new features heretofore (O.J.!) unavailable to Windows 3.1x users. Yes.

Will I sell Windows 95 to my clients? Whoa - that's a tough one. I would be happy to sell it to corporate types in a position that requires them to research such products for possible future use in a production environment. I would also be happy to sell it to developers who need to get a head start on the upcoming market for Windows 95 applications. For everybody else, I will recommend that they wait for a "mature" version, such as 1.1 (or at least 1.01) before putting Windows 95 on production computers. This way, they can stand by and watch the lemmings flounder in the sea while waiting for all of the bug fixes, hints, tips, and tricks to be published. I will also recommend that they use latex rather than leather gloves and then remember to dispose of them discreetly.

CLASS ACT

by Wayne M. Krakau - *Chicago Computer Guide*, August 1995

Wayne is a partner in Krakau Business Computer Systems, a systems integration firm that is a Novell Gold Authorized and Novell IBM Host Master Authorized Reseller. He has been working with computers for 18 years. He holds an M.B.A. in Marketing and a B.S. in Information Science and was originally educated in Computer and Information Systems Analysis Engineering. He holds Novell CNI and ECNE ratings and all five of the original LANDA Certifications. Wayne can be reached for questions, comments, or topic suggestions at (708)298-7695 or via E-mail to wayne@mcs.com.

You've been tinkering with NetWare LANs for years, and you are tired of having to either call for help or spend hours on research, every time a new problem or question arises. Or, you've suddenly been put in charge of your company's LAN, but you only have experience with other computing environments such as mainframes, minicomputers, or individual, unconnected PCs. Or, you could experience the worst case, being shanghaied into administering a LAN with no background at all. You could even have the most benign experience, a desire to further your career in working with LANs.

All these questions suggest formal training, with the most cost-effective option being classroom instruction. Individual training by a systems professional is just too expensive for any but the briefest of introductions to the technology. (Keep in mind that I could make some serious extra money by conning people into letting me give them extensive personalized network training beyond this level, so the previous statement goes against my own financial interest.)

Once you've decided to go the classroom training route, you can easily be overwhelmed by the available choices. First, you should determine what you want out of the training, your time constraints (both time to completion and your personal availability), and your budgetary limitations. These factors will guide your decision.

If you are concerned only with getting the information that you need to run a network, and getting it fast, you need to look at independent (that is not associated with Novell) classes that emphasize the rapid acquisition of practical knowledge without regards to any of Novell's tests. Since these courses are not subject to any review process (short of outright fraud enforceable via criminal statutes), you must evaluate them carefully.

Obviously, a word-of-mouth referral is best, assuming the person giving the referral is LAN-literate enough to give a valid opinion and works on a system similar enough to yours. Learning skills appropriate to someone else's LAN won't help much.

Another technique is to analyze the course curriculum carefully. Does it fit your needs? If you have access to more experienced colleagues or friends, ask them to help you sort out what may be unfamiliar terms. (Though, if the training organization is incapable of translating their curriculum into something resembling English, you might be justifiedly wary.) They can also help you categorize skills by how applicable they are to your particular situation.

Also be sure that the target audience concentrates on people with your current skill level. You don't want to be either bored out of your mind or left in the dust. For example, if you only need information on the day-to-day running of a network, you might not care about buying advice. If you are getting a new network or upgrading an existing one, product analysis skills and buying hints would be essential.

Talking to the instructor is another way to select a class. If the instructor is honest, you might be warned away from an inappropriate class. Either way, you can get some idea of the instructor's knowledge and, most important, people skills. Realistically, the individual instructor is a lot more important than the organization in training situations.

Visiting the training site can be very revealing. If you see a poorly-maintained group of mismatched computers, you can expect the class to be disrupted by hardware and software failures. If you notice unexplained evidence of piracy, then the organization's ethics may be suspect.

Speaking from the point of view of someone who designed and taught an independent LAN course for several years, most of the courses that I am aware of don't look very appealing. Based upon my examination of their advertising material, they habitually misallocate time by making decisions based more on academic or marketing opinions rather than real-world considerations. They also neglect to narrow their target market properly. They just want to get as many bodies into the classroom as possible. A perfectly valid class can be rendered near-useless by inappropriate targeting and a failure to reveal prerequisites.

If you have any interest at all in using training to advance your career, the Novell courses are the best for you. All other classes are simply cram courses designed to stuff your short term memory with just enough material to pass the Novell tests.

Even the cram courses given by Novell Authorized Education Centers that use Novell course material fall into this category. (I'm sure that the previous statement will do wonders for my future employability as an independent Certified NetWare Instructor.) Keep in mind that the industry is wising up. Sneaking your way through the tests by either using a cram course or just "winging it" based on the prior acquisition of ad hoc (and by definition, incomplete) knowledge isn't enough these days.

Employers and clients have become aware of the overwhelming advantages of the structured classroom experience, the standardized course material, and, especially, the presence of a CNI, in Novell's training method. The exposure to a CNI, and to classmates with differing experience is the key to this equation. I'm not even particularly thrilled with people who purchase the course materials and then take the tests.

Again, even when choosing from among the many Novell Approved Education Centers, the individual instructor is the most important factor in the equation. Based on feedback from former students, the quality of the instructors was the biggest variable, even within a single training organization.

Note that my preference for the full Novell classroom "experience" is based both on my own experiences as a LAN student in classes where I had overwhelmingly more practical experience than all but one of my instructors, as well as the post-class comments by my own students. I managed to acquire tons of useful knowledge from my instructors and classmates in every class in spite of my extensive experience. (Just last week, I learned at least a dozen new NetWare tricks from the administrator of a client's LAN who is just about to start his formal NetWare training!)

When I search for new employees, I strongly prefer people who have attitudes toward training similar to mine. That is, not only to attend the class, regardless of prior experience, but, while in the class, participate fully, to the point of sucking the instructor's brain dry of information in a vampire-like quest for knowledge.

TECH SUPPORT?

by Wayne M. Krakau - *Chicago Computer Guide*, September 1995

Wayne is a partner in Krakau Business Computer Systems, a systems integration firm that is a Novell Gold Authorized and Novell IBM Host Master Authorized Reseller. He has been working with computers for 18 years. He holds an M.B.A. in Marketing and a B.S. in Information Science and was originally educated in Computer and Information Systems Analysis Engineering. He holds Novell CNI and ECNE ratings and all five of the original LANDA Certifications. Wayne can be reached for questions, comments, or topic suggestions at (708)298-7695 or via E-mail to wayne@mcs.com.

It's like the old saying "With friends like these, who needs enemies?" With technical support like this, who needs bugs.

It started simply enough. We planned a network upgrade using a new, more powerful file server with a 32-bit bus. We would reuse the existing SCSI-2 (Small Computer Systems Interface, Version 2) duplexed Toshiba hard disks and Data Technology Corporation (DTC) controllers. (Duplexing is a variation of mirroring, employing both dual disk drives and dual controllers to provide redundancy while improving performance over the equivalent mirrored configuration, which uses a single controller.)

The client already had a copy of NetWare 4.1. He also had a Pacific Micro Data DAT (Digital Audio Tape) tape drive (really an older Hewlett-Packard model in disguise), an Adaptec AHA-1522 SCSI controller, and Palindrome's DOS-based backup software to implement a server-based backup system, replacing the workstation-based backup system they had previously used. As per to common practice, we were giving the tape drive its own controller to maximize performance.

We were supplying the file server, a DTK 486/DX4-100 with three bus-mastering VL-Bus slots. This machine was a little light for their needs, but it is a not-for-profit organization that has to make due on a shoestring budget. The same was true of the salvaging of old parts, especially the three 16-bit SCSI cards that were being installed in a machine with 32-bit slots.

The first bad news was from Data Technology. After repeated attempts to install the new tape drive on the old file server, the client was told by DTC that Adaptec's ASPI (Advanced SCSI Programming Interface) software would not coexist with their ASPI-emulator software in a NetWare server. Also, their ASPI-emulator wasn't close enough to real ASPI (designed by Adaptec) to run an Adaptec controller while their Adaptec-emulator chipset on their boards wasn't close enough to the real Adaptec chipset to be run via real ASPI software. Stalemate! We decided to switch to a single 32-bit Adaptec 2840 VL-Bus card (to stay at least close to the budget) to mirror the disks.

When we delivered the new file server, with NetWare already installed on the redundant hard disk taken from the original server, we couldn't get the tape drive to consistently respond to system level commands. Finally, after two extended sessions with Adaptec's very patient technical support team, the server itself failed. Using a hardware diagnostic board, I determined that the motherboard-based processor interface had failed. Oh well, back to the office.

DTK gave us a warranty replacement motherboard immediately. (I drove over to their Elk Grove Village, IL plant to pick it up right after I requested it. The location of their Midwest plant - one of six nationwide - is one reason that I deal with them. Immediate parts availability and in-person tech support can be very handy.) Even with the new motherboard, the system wouldn't work. I started swapping parts with some of our in-house computers and found that the motherboard wasn't the problem! (Oops. New personal rule - don't try sophisticated troubleshooting when spaced out from too many hours of continuous work.) The Adaptec AHA-1522 was bad, and when it finally completely failed, it took the IDE/I/O board with it. Replacing both boards was the solution. (Note that I am switching between "we" and "I" because most of this project was a client/KBCS staff team effort, not because I have a big enough ego to use the royal "we".)

Since there were two open 32-bit slots in the server, the client wanted to replace the 1522 with another 2840 VL-Bus card. He checked with Palindrome and was assured that was a valid combination, but when we tried to install the Palindrome software, we had trouble getting it to work. Sometimes it worked. Sometimes it locked up the workstation that ran the control module (the actual working modules run on the file server). Sometimes it locked up the server. After an extended session with Palindrome tech support during which I could tell that the person I was talking to was constantly consulting his colleagues, we brought the system back to our office again.

On the next day, the client called Palindrome before I had a chance to. He was informed in no uncertain terms that Palindrome had never approved the use of the 2840 card(odd, since it uses the same chipset as the 2740 EISA series which is approved by Palindrome), and that there was no way that anyone on their staff would ever even imply that was a valid controller to use with Palindrome's software! This was after two phone calls approving the use of the 2840 card, one of which included over an hour and a half of experimenting with different configurations of the board. The tech support person (as I stated, backed by his colleagues) even conversed in depth with me about the ins and outs of VL-Bus busmastering technology. There could be no mistake in either my conversation or my client's that the 2840 VL-Bus card was under discussion.

At this point I gave the client one of our in-house Adaptec AHA-1542 16-bit SCSI controllers. I was able to back up successfully with the 1542 controlling the tape drive while a 2840 controlled the disk drive, so I redelivered the system.

Since the client now had an extra 2840 32-bit controller, he decided to implement disk duplexing, but he kept having trouble with the tape software. Thinking that he would save me some trouble, he called Adaptec to find out how to get the three SCSI cards to coexist. This turned out to be a well-meaning mistake since I had talked to Adaptec and had instructions already prepared for the three-card configuration. He managed to get an inexperienced technician who didn't seem to know his own limits. The configuration advice that he got was so bad (though admittedly Adaptec's only error in a long stream of calls on this project) that it the system started generating random hard disk errors, one of which, by chance, trashed sector zero, track zero - the one vital spot on the drive. After repeated attempts to recover the drive using various disk utilities, we finally had to give up. We reformatted the drive, losing all of the data, and hoped that the one backup we had was good. Note that the old server was still in production and could be used if the tape was bad, but only with much redundant work.

After reinstalling NetWare, the client could not get the tape software to work. This time, when he called Palindrome, they told him that their software wouldn't work in a server that had even one, much less two, 2840 cards in it, even though those cards didn't have anything to do with the tape drive! They further stated that they had called both Adaptec and Novell and found that the 2840 was "NOT APPROVED FOR USE IN A NOVELL SERVER"!!!!

Right after this, the client pulled out the 2840 box and documentation. The box was plastered with "Novell Approved" and "Works With NetWare" stickers. The documentation had multiple chapters on how to install the board in a NetWare server. He called me regarding this shocking conversation, so I immediately called Adaptec to get the facts. They told me that not only was Palindrome wrong, but that Palindrome might be expecting a letter from Adaptec's legal department with the magic phrase, "Cease and Desist" on it, along with a not-to-subtle threat of legal action.

At this point, the client demanded that I take action to solve this dilemma, so I gave him a new AHA-1540 16-bit card to use with the disk drives in place of the now-forbidden 2840s. This combination at least partially worked, though using 16-bit controllers slowed the server appreciably. I use the term "partially" because the Palindrome software was unable to successfully restore NetWare 4.1 rights and NDS (NetWare Directory Services) information. Only the raw data was retrieved. The client had to reenter the NetWare-specific information, even though the software is, in theory, approved for use with NetWare 4.1.

I knew I should have been suspicious when the Palindrome installation routine demanded to be run in bindery mode using the old NETX Workstation Shell instead of the current VLM DOS Redirector.

This is of course, not to mention a user interface that would have embarrassed a 1970's era mainframe programmer (keeping in mind that I was one).

The final insult happened when the client called Palindrome to demand a refund, since, after even more tech support calls, the backup system still wasn't trustworthy. Palindrome management told him, in essence, "tough luck". They cast aspersions on my company's competence and ethics, and hinted, subtly, that the client had been stretching the truth for effect.

Oh, I thought that it might be interesting to point out that the client system administrator in question is an ordained cleric who uses his title in all introductions, and that organization that he works for is, naturally enough, a religious one. If you believe in that sort of thing, is there a special place in Hell for those who screw members of the clergy in business transactions?

COBBLED CABLING

by Wayne M. Krakau - *Chicago Computer Guide*, October 1995

Wayne is a partner in Krakau Business Computer Systems, a systems integration firm that is a Novell Gold Authorized and Novell IBM Host Master Authorized Reseller. He has been working with computers for 18 years. He holds an M.B.A. in Marketing and a B.S. in Information Science and was originally educated in Computer and Information Systems Analysis Engineering. He holds Novell CNI and ECNE ratings and all five of the original LANDA Certifications. Wayne can be reached for questions, comments, or topic suggestions at (708)298-7695 or via E-mail to wayne@mcs.com.

It's time again for what is becoming a regular event - my annual rave-out on LAN cabling! Why, you may ask, am I making this an annual event? That's an easy one. It's because my rate of discovering bad cable plants versus good ones has risen from two out of three to about five out of every six new networks that I encounter -- and the only reason the ratio isn't worse is because I am including cases where I was called in by a cabling company for non-cabling troubleshooting immediately after a complete rewiring job!

Since my last series on cabling, I have received a continuous stream of inquiries on violations of LAN cabling standards. Most calls involve requests for, in essence, a special dispensation from the standards (and by implication, the laws of physics), along with a demand for an exact enumeration of just how much degradation will occur for a given set of standards violations.

As to the first issue, I usually refer the questioner to his or her favorite deity. (With the stipulation that I get rights to the concession stand if there is a major transgression of the laws of physics. If needed, I can refer them to a team of lawyers in Los Angeles who have a proven record of a willingness to promote absolutely anything without fear of divine - or even judicial - retribution.)

As to the second issue, I refer the questioner to The Psychic Hotline. The people there are better equipped for that type of prognostication. They may also give advice as to an alternate career path for those who are in charge when the aforementioned standards violations cause a network to collapse. (Learn how to say "Do you want fries with that?")

These are case studies of cable plant problems, accompanied by my semi-coherent ramblings. What's scary about these cases is that they are all from one three-month period!

The first case started with a complaint about incomplete overnight tape backup jobs. After inspecting the workstation holding the tape drive, I found an invalid configuration. The CONFIG.SYS and AUTOEXEC.BAT, the tape drive setup, and even the network card setup were a mess. The settings had been changed repeatedly during multiple conversations between the administrator and different technical support representatives of the tape drive manufacturer, causing a gradual drift into this unstable configuration.

After initial testing, I spent some time completely reconfiguring the workstation. Then I tested the tape drive by backing up the C drive. That test was successful. I started another test, this time by backing up a portion of the F drive, the DOS representation of the volume SYS on the file server. This test worked, but it was alarmingly slow. It was so sluggish that I figured that the workstation could easily lose its network connection in midstream when backing up all of the network drive. My theory matched the description of the initial problem that was reported before all of the configuration fiddling.

This led me to examine the file server. It was running quite efficiently (luckily, since I'm the one who last reconfigured it). Finally, I checked the cable plant. I examined the individual cable segments of this Thin Ethernet network and found several that were plain RG-58 coaxial cable instead of the specified RG-58 A/U. Those extra couple of letters make a big difference in electrical characteristics, and the lack of them can cripple a network.

I also found several bad terminations. Some were poorly applied screw-on connectors (which seldom make a very high quality connection). Others were improperly installed crimped connectors. One

termination even had a two-inch strip of completely exposed wire braid! The outer sheathing was completely cut away.

A quick check with my LANCat (Datacom Technologies, Inc., 800-468-5557) showed the biggest problem. The LANCat is a LAN cable testing device that includes a Time Domain Reflectometry feature (TDR is a kind of radar). It showed a total length of nearly 1,000 feet! The Thin Ethernet standard requires a maximum trunk segment length of 185 meters - about 607 feet. The fact that the signal got through at all is a tribute to the resilience of the underlying specification.

We managed to get the network down to less than 800 feet after painstaking segment-by-segment testing (since nobody new where the cable was routed) by removing a completely redundant segment. We also discovered that most of the segments had a large coils of excess cable stashed in the ceiling. This means that the client company can clean up its cable plant by replacing nonstandard cable, cutting off the surplus cable, and reterminating with professionally crimped connectors. This will fix both the length and faulty termination problems. The notes from the tests will be used to create cable plant documentation. Interestingly, the cable was strung by the client's computer maintenance company.

The second case also involves Thin Ethernet. I was inspecting a LANtastic network for a possible upgrade and expansion to a NetWare 4.1 system. The client mentioned that the current system was not reliable. Workstations locked up in ways that I recognized as indicative of losing their connection to the file server, possibly due to cabling problems. The client also told me that the cable was strung by their previous computer vendor - a dead giveaway for cabling problems. (Just for good measure, the computers weren't even set up competently!)

I immediately noticed that the coaxial cable went up into the ceiling and was just randomly dragged across the suspended ceiling as needed, even though it wasn't the mandated plenum cable. Plenum cable is sheathed in Teflon, which doesn't burn as easily as the PVC (polyvinyl chloride) that sheaths regular cable, and doesn't give off poisonous gas - mostly cyanide - when it burns, like standard cable does. Using loose, non-plenum cable in suspended ceilings used for ventilation is a violation of building and fire codes.

I also noticed that the cable was a mix of a couple of segments of the specified RG-58 A/U coaxial cable along with multiple segments of RG-62 cable - the kind used for ARCnet LANs and IBM 3270 series mainframe terminals! To add spice to the mix, there were multiple "opens". This is a term used to describe a point where a cable has an open, exposed socket, not filled by another cable or a terminator. A terminator is a resistor of the appropriate resistance - 50 Ohms for Thin Ethernet - encased in a sort of cap that fits on the open end of a cable segment.

The most common open is usually a T-connector (a sort of three-way intersection for cable) left dangling after a workstation has been moved. The side of the T connected to the workstation's network card is left open. The correct procedure is to replace the T-connector with a barrel (inline) connector that has only two sides, thereby removing the offending open socket.

This site had multiple opens created as the company redecorated and renovated its office. They did, however, have nice, attractive crimped connectors on all of their cables. (That's one in a row!) Considering the mix of 50-Ohm (impedance) RG-58 A/U and 92-Ohm RG-62 cable, and adding in the numerous opens, I was quite surprised that they could get any work done between the workstation disconnections! This client has decided to scrap the existing cable plant and has drawn up plans to switch completely to 10Base-T (Ethernet over Unshielded Twisted Pair - UTP - cable).

Next month I will continue this theme with examples that show that even professional phone cablers - as opposed to the (alleged) computer people in this article - can't be expected to dabble successfully in high-speed (LAN) data cabling.

COBBLED CABLING - Part 2

by Wayne M. Krakau - *Chicago Computer Guide*, November 1995

Wayne is a partner in Krakau Business Computer Systems, a systems integration firm that is a Novell Gold Authorized and Novell IBM Host Master Authorized Reseller. He has been working with computers for 18 years. He holds an M.B.A. in Marketing and a B.S. in Information Science and was originally educated in Computer and Information Systems Analysis Engineering. He holds Novell CNI and ECNE ratings and all five of the original LANDA Certifications. Wayne can be reached for questions, comments, or topic suggestions at (708)298-7695 or via E-mail to wayne@mcs.com.

This is the second part of yet another series that I'm doing on cabling. (Considering my last name, perhaps I should have called it "Kobbled Kabling".) I shall continue with my litany of cable madness.

The next case involves another LANtastic network using Thin Ethernet. I was visiting the site ostensibly to regain access to printers and to reconfigure their security system. While cleaning up each combined server/workstation and plain workstation - a major task in itself - I kept having workstations randomly fall off the network. (That's technicalese for losing the connection to the file server.)

Since I had stabilized the configuration of the computers, I suspected the LAN cable. One obvious, and very common problem, was the lack of grounding. In addition to the 50-Ohm terminators on the ends of Thin Ethernet, the specification requires that one of those terminators - arbitrarily chosen - be grounded. This cable (as well as most others that I encounter) wasn't grounded. This can slow a LAN because packets (the smallest unit of data sent) are often sent multiple times before they get through, and it can compound other errors, but is almost never a serious enough standards violation to blow workstations off the LAN. More detective work was in order.

I examined each individual length of cable and the T-connectors and found that although the cable was the proper RG-58 A/U, the BNC connectors on the ends of each length were atrocious. I suspected that they were crimped by an amateur (noting that I am only a "talented" amateur in LAN cabling, and I defer to true professionals in that field) using the unreliable, inexpensive (All right, I'll admit it. I really mean CHEAP!) crimping tools sold by a certain national chain of electronics stores. (Notice how I deftly sidestepped mentioning the chain's name? I don't want to spend the rest of my life arguing legal niceties in front of one of Judge Ito's professional kin.) This company's crimping tools work well enough for cable TV, but not for LANs. Moreover, even their tools require at least some expertise to work correctly.

The cabling system had multiple occurrences of loose pins (the part that caps the center conductor of coaxial cable), some to the point of dangling freely, misshapen ferrules (the part that goes over the cable and gets crimped), misaligned bodies (the outer portion that actually makes the connection), and the ever popular random bits of braid wire (the crisscrossing mesh that forms the outer conductor of coaxial cable) sticking out all over the place. The terminations were so bad that I found one spot where I could slightly rotate a length of cable just short of its end and turn the network on and off like a faucet! I found out later that this spot could have been used as a sensor for their burglar alarm. If you even slightly bumped the table that supported the PC at the endpoint of this cable segment, the cable would be rotated just enough to shut down the network! That'll make a fairly good motion detector.

As I suspected, the client told me that their "computer guy" installed the cable. This is yet another case of someone dabbling outside their field of expertise and having the client suffer for it. Considering what a mess the overall network was in, perhaps the term "amateur" could be applied to the computing aspects of this "computer guy's" expertise, too.

In the next case, I was called in by Bruce Kahn the president of my favorite professional data cabling company, Telnet Communications Consultants, Inc. (Buffalo Grove, IL, 708-215-0003). His

client was designing a new network and he wanted to work jointly with my company (KBCS) in redesigning the cable plant in conjunction with redesigning the whole network.

Another company had already laid cable throughout the building. Their contract called for cable appropriate for future use with 100-megabit networking. A friend of one of the client's managers had informally reviewed the cable system and had noted some potential length standards violations. He was a computer professional in another area of computing, so he suggested that an outside company review and possibly even test the cable plant. That's why Telnet was called in. Bruce also noticed that the plan for the network, in general, seemed illogical and inefficient. That's why Bruce called me.

After the manager's friend cast doubts on the original system, the cabler partially rewired it to split it into smaller segments. Telnet's technicians tested the system and found that a third of the segments violated the 100-meter (about 328 feet) 10Base-T (Ethernet over Unshielded Twisted Pair - UTP - cable) standard. The lengths were more than 450 feet!

This did not even include an allowance for a 3-foot jumper cable in the wiring closet to go between the concentrator and the patch panel, or the 12-foot (average) jumper going from the wall socket to the workstation. It also didn't include an allowance for extra terminations. The termination at the concentrator socket and the one at the network card socket are a part of the standard, but additional intermediate terminations require a distance "penalty" or correction factor. We usually use a 10-foot penalty. There were three extra terminations, the "IN" port of the jumper block, the "OUT" port of the jumper block, and the wall socket adjacent to the workstation. That extra 30 feet of penalties added to the 15 feet allowance for jumpers means that even more segments were too long!

Telnet added secondary concentrators, placing them so they could halve the length of the over-spec segments. This brought them back into agreement with the rules.

The riser cables (between floors) were 25-pair cables that were broken out into 12 data links at the patch panels of each wiring close. That's a violation. They might have gotten a 10-megahertz signal down those cables, but a 100-megahertz signal would never make it.

Telnet reterminated the riser cables to use the 568B standard for terminations, using eight individual wires for each data link, for a total of four data links. They also cut the "leftover" wires so that nobody could inadvertently use them later, throwing the system back out of spec again.

The original cablers had strung multiple pairs of fiber optic cables to each wiring closet based on the manager's friend's mistaken belief that it was required for high-speed networking. They didn't bother warning him about this bit of misinformation. For good measure, they used older, potentially less reliable, and much more expensive methods of terminating the fiber optic cable. Telnet left this cable in place in case it is ever needed to overcome the distance limitations of UTP cable.

After the client's cable system was stabilized and retested, we looked up the details of the original system and calculated the segment lengths before the manager's friend spotted potential problems. There were multiple segments over 900 FEET long! (Remember. We are dealing with a 328-foot standard.)

The original cabler did much of the labor needed to bring the system into compliance, and they refunded a fraction of the client's fee. They never did reimburse the client for the outside resources needed (Telnet and KBCS) or for employee time lost during the project.

The really scary part is that this cabler is a huge multimillion-dollar telecommunications company, more known for its telephone services than its LAN expertise. (Notice how I weaseled out of an explicit name mention here, too! This company has more lawyers than Madonna has bedpost notches.) There is no excuse for them selling cabling services in a field in which they are ignorant of the standards. Also, why in the world don't they KNOW the standards? This is a case of major-league dabbling outside your expertise.

It is important to remember that cabling standards for high-speed data (like LANs, for instance) are much more strict than those for telephone cabling. This holds for material, tools, and even the techniques used in actually laying cable. A reliable data cabling company can double as a phone cabler (especially since it's very cost effective to string both at the same time), but the reverse is an invitation to disaster.

Please join me next month as we continue our exploration of the darkest unexplored regions of LAN cabling.

COBBLED CABLING - Part 3

by Wayne M. Krakau - *Chicago Computer Guide*, December 1995

Wayne is a partner in Krakau Business Computer Systems, a systems integration firm that is a Novell Gold Authorized and Novell IBM Host Master Authorized Reseller. He has been working with computers for 18 years. He holds an M.B.A. in Marketing and a B.S. in Information Science and was originally educated in Computer and Information Systems Analysis Engineering. He holds Novell CNI and ECNE ratings and all five of the original LANDA Certifications. Wayne can be reached for questions, comments, or topic suggestions at (708)298-7695 or via E-mail to wayne@mcs.com.

This is the third part of my series covering LAN cabling. Forget "Sick Building Syndrome". We're talking about "Sick Cabling Syndrome".

I'll start with a clarification of a potentially ambiguous paragraph in Part One of this series mentioning the use of plenum cable. The possibility of misinterpretation was pointed out to me by a reader via E-mail. (Yes, I do read and respond to E-mail!) Many localities require the use of plenum-rated cable if a suspended ceiling is used for air return. If the return air is ducted, you can often get away with non-plenum cable, though there may be rules detailing just how that cable is laid. Other localities such as the City of Chicago require the use of conduit (pipes) for any new cables going into a plenum ceiling. For them, even plenum-rated cable isn't good enough.

The next incident was discovered during a conversion of an Ethernet system from coaxial cable to UTP (Unshielded Twisted Pair). We had to keep the existing system stable while transitioning to 10Base-T in stages. Luckily, the network was already split into two based on location by having two network cards in the file server. This made our task easier since it automatically isolated half the network.

The network administrator got a head start by switching all workstations to combination network cards that could use either coaxial or UTP cable. This meant we could move them from coax to UTP at will.

Since we needed a stable network prior to starting the switch, I inspected the existing coax. I found numerous open T-connectors, substandard BNC connectors (including the ever popular dangling wires), and a lack of grounding. I also found many kinked segments, including some with multiple kinks.

It took a little exploring to find out why there were so many kinks. I finally realized that whoever originally wired the place thought it was a good idea to stash the cable out of the way for convenience. Since the cable is naturally springy, they had to find a way to keep it from popping out from under desks and the like and getting underfoot. Their solution was simple. They just pinned it down by putting it under desk legs, cabinets, and printer stands! Over time, due to furniture moving, multiple kinks and crushed spots were created. I have no idea how any signal was getting through.

I went back a week later to help with the transition and found that all of the cubicles in the main portion of the office had been removed, leaving an open area. When I asked the administrator if that portion of the cable plant had been rerouted to accommodate the missing section that used to support the cubicles, she told me about her encounter with one of the workers. He had asked her if he could cut the cable in the ceiling. She had, of course, forbidden him, since the proper method would have been to install two BNC connectors at the cut and then connect them with a barrel (inline) connector - while the LAN was down. Later, he told her not to worry because he had "fixed" everything. He cut the cable on an active LAN and had taped the two ends together!

Since this part of the office was scheduled to be converted to UTP in less than two weeks, and, amazingly enough, the remaining workstations on the taped coax still seemed to work, I decided that retreat was the better part of valor and kept my paws off the system. I was afraid that if I touched the cable while it was in such a precarious state, the network would collapse. It had survived other abuse and I wasn't about to fiddle with it in spite of the fact that I couldn't figure out how the LAN stayed alive!

The final case I will present is the ultimate in cabling adventures. I was called in to optimize a LAN and to help prepare the administrators for an upcoming upgrade. As a part of my inspection of the existing LAN, I examined the cable plant. First, I checked the coaxial cable that was used to connect the 10Base-T concentrators (hubs). It was 92-Ohm RG62 ARCnet cable instead of RG-58 A/U Thin Ethernet cable. Also, the coax was partially unterminated and was not grounded.

Then I checked the front of the hubs and the patch panel. Almost all of the RJ-45 connectors were incorrectly installed. They had lots of unsheathed individual wire hanging out and didn't look well crimped. The system also had Category 3 cable going into the wall and Category 3 patch panels, but used a random mix of Categories 4 and 5 jumpers. As I've stated in previous columns, you can't mix cables with different characteristics (which is what "Category" refers to) without causing a point of distortion where the different types meet. Think of it as the difference between clear glass and translucent glass. The light gets through both, but only the clear glass provides a clear, undistorted "signal".

Just to make things really interesting, several jumpers were completely unsheathed and untwisted wires! The specification calls for Unshielded Twisted Pair (UTP) cable, not Unsheathed, Untwisted, and Inherently Unpaired (UUIU?) cable!

As I was about to continue my inspection in other areas, I noticed a metallic glint among the jumper cables. I pushed aside the intervening cables and discovered that one jumper was made up of NAKED COPPER WIRES! It was with great relief that I tracked down the ends of that jumper and found out that even though this cable (maybe "bunch of wires" would be more accurate than "cable") was tangled within the LAN jumpers, it was really a part of the phone system. Whew! It was only a false alarm. I strongly suggested to the LAN administrator that a company with a live LAN but a dead phone system would not necessarily stay in business. He got the hint and promised to notify the phone system administrator about the wacky phone cabling.

Now, onto the workstation side of the system. I checked several offices and found that the LAN sockets were RJ-11 four-wire jacks, not the required RJ-45 eight-wire. Although Ethernet only uses four wires (two pairs), the specs are based on the internal spacing within the RJ-45, not the smaller wire spacing found in RJ-11 or even six wire RJ-12 jacks.

To test some sample cable runs with my LANCat (Datacom Technologies, Inc. 800-486-5557), I had to disassemble the socket and experiment with alligator clips to find which wire was which. Every run that I tested completely failed even at the lowest quality (Category 3) settings. This meant that even though it was Category 3 cable, it was strung in a nonstandard manner and couldn't carry an Ethernet signal reliably.

Because there were so many standards violations, I called my cabling authority Bruce Kahn of Telnet Communications Consultants, Inc. (Wheeling, Il, 708-215-0003) to discuss the results of my inspection. He confirmed my conclusion that the system was a disaster waiting to happen. It was so destabilized that the slightest change could cause it to collapse.

Because I had only tested a few cable runs, the administrator requested an estimate for Telnet technicians to come in and do a thorough test of the entire cable plant. He felt that my inspection was too limited for him to justify a major recabling job. He also felt that since I was, in essence, out of my major field of expertise (information that I had volunteered previously), that Telnet's results would carry more authority.

I discussed this client with Bruce a few weeks later and found that upper management had vetoed the request for a thorough cable plant test. The system wasn't broken (at least from their point of view), so why fix it? In retrospect, I actually agree with their decision to skip further testing, but not for the reason they would have expected. I am convinced that the system is so totally unsalvageable that any testing done in hope of saving portions of the existing cable plant would be a total waste of time and money! I think that they should just tear out the whole thing and start over from scratch - preferably before it collapses and heads start rolling.

The overall theme of this series harks back to one that I've repeatedly used in when covering other aspects of systems integration. Don't put up with people dabbling outside their field - and for

goodness sake, don’t dabble yourself! For LAN cabling, deal with specialists with heavy experience in high-speed data cabling for LANs. This excludes computer service companies, programmers, dealers, resellers, VARs, systems integrators (including my company), other assorted amateurs, and even telephone cabling companies. The job you save may be your own.

AND THE WINNER IS? - Part I

by Wayne M. Krakau Krakau - *Chicago Computer Guide*, February 1996 - *NewsWare*, September 1996

Wayne is a partner in Krakau Business Computer Systems, a systems integration firm that is a Novell Gold Authorized and Novell IBM Host Master Authorized Reseller. He has been working with computers for 19 years. He holds an M.B.A. in Marketing and a B.S. in Information Science and was originally educated in Computer and Information Systems Analysis Engineering. He holds Novell CNI and MCNE ratings and all five of the original LANDA Certifications. Wayne can be reached for questions, comments, or topic suggestions at (847)298-7695 or via E-mail to wayne@mcs.com.

Oscar, Emmy, Tony, Grammy, and countless other awards are presented each year in the entertainment field. I've decided to create my own award - the SI (pronounced like a nickname for "Simon") for the systems integration field.

In my scheme of things (Hey, it's my award, so I get to define the rules!), the SI awards can be both in positive and negative categories. Also, I make up the categories and choose the winner. This eliminates all those nasty levels of bureaucracy that usually interject politics into the award process. The results are stored on my hard disk until such time as I E-mail my column to my editor. Forget Price, Waterhouse and those other accounting firms. Finally, I can give out the awards any time I want.

To start on a positive note, the first SI Award is for **Honoring a Warranty Far Above and Beyond the Call of Duty**. It goes to Micropolis Corporation for work on their Radion LT 4200 model RAID (Redundant Array of Inexpensive Disks) systems. In two separate incidents Micropolis absolutely shocked me in their willingness to back their product.

In the first incident, the RAID unit had been placed in the far corner of a small storeroom that had been converted into a computer room. For its protection, it was surrounded by a tower-cased server, a large uninterruptible power supply, and an industrial-style heavy steel table. Access was further impeded by the monitor, keyboard, and tape drive atop the table. Surely, no one could inadvertently damage the RAID system - or so we thought.

A few weeks after installation, a technician from a security company was installing an entry system for the computer room door. He was using a huge rechargeable Mikita drill to punch through concrete and brick walls. When he wanted to recharge the appropriately large detachable battery, he plugged it into its equally weighty recharging unit and started looking for an open electrical socket. The first one he found was in a bank of open sockets located directly above - you guessed it - the RAID system!

He reached across all of the aforementioned obstacles (he was tall and had a long reach) to plug in his recharger. He then accidentally let go of the combined recharger-battery unit, but retained his grip on the cord. This caused the unit to swing, pendulum-style, into the side of the RAID system, immediately destroying two of the three disk modules, while subtly damaging the third.

Unbelievably, Micropolis shipped two replacement modules next day air at their expense! The boxes even included repacking instructions for the old modules and a prepaid return-shipment sticker! A few days later, when the damage to the third module was discovered, they replaced it, too.

In the second incident, we safeguarded the RAID system in a similar protected position. It was sheltered by a tower-cased server, an uninterruptible power supply, and a heavy wooden table. A large empty space was left in front of the RAID so nobody could trip over or otherwise bump it.

Some time after the installation, someone at the client site decided that the empty space was going to waste and would be better put to use as a storage area for software boxes. They completely filled the space with tightly packed boxes. You could have easily predicted the result. Someone tripped and accidentally kicked the first box in the row. Newton's Laws being what they are, the force of the blow was transmitted to the last box in the row and, hence, to the RAID system, crashing one of the three disk

modules. As in the previous case, Micropolis immediately cross-shipped a replacement module along with the prepaid return-shipment sticker.

I have reread the fine print in the Micropolis warranty several times since these incidents, and I can't find anything to indicate that either dropping heavy objects on or kicking of their products should be covered. Nevertheless, they covered both of these clients.

The second SI Award is for **Selfless Devotion to Technical Support**. This award goes to American Power Conversion (APC), the power protection device manufacturer, for their extraordinary treatment of my clients in multiple situations in which APC wasn't directly backing an existing product nor were they in a position to influence future sales of their products.

While I have encountered several examples of this outstanding attitude from APC, the most dramatic occurred with a client who was having mysterious crashes, lockups, and communications glitches. We had installed a new network and kept getting calls for help from this client. Workstations were locking up randomly, the fax/modem would lock up for hours at a time, and the most telltale symptom, the images on some monitors were "dancing".

The client had just had a complete electrical rewiring job, so we were initially denied permission even to explore the possibility of electrical problems. Any discussion of electrical problems was perceived by the client as a smokescreen for my company's incompetence or outright fraud. We completely retested the 10Base-T cable system, and replaced motherboards, video cards, network interface cards, and even disk controller cards in an effort to pin down the cause of the crashes.

Finally, the client let us borrow an additional Uninterruptible Power Supply (the server already had one) from APC's local office to try to solve the communications problems. The theory was that the UPS could cure a power irregularity. The client refused to believe that a UPS without a valid ground can't do its job. He even declined my offer to speak directly to an APC engineer for an explanation of the details.

Meanwhile, I spent hours on the phone with APC technicians, eventually working my way up the line until I was dealing directly with one of their top engineers. He patiently educated me on the intricacies of power protection, especially regarding interference. He emphasized that without a legitimate ground, all bets were off.

Using this information, I was able to alter the client's screen images at will by turning various fans and lights on and off. At last, I had evidence that electrical problems existed. The client brought back the electrician and spent several thousand dollars on rewiring the rewiring! Funny - I would have thought the electrician would fix what the APC engineer had described as below-standard wiring practices as part of his warranty work - for free!

Still, even after having extensive direct talks with APC, the client trusted the electrician and refused to take APC's suggestion to call in an engineer specializing in power to check the system. APC even offered to lend thousands of dollars worth of test equipment to the client and to assist in interpreting the results. The client declined.

Some months later I heard from the client that there were occasional crashes if all of their computers were on at once. The client again refused to believe that power problems could be the cause, since the electrician said the wiring was just fine and that the ground passed his test. His only test, however, was with a standard multimeter, not the specialized instruments offered by APC. That is how I left things. I couldn't force the client to hire a power specialist and he still doesn't trust my analysis.

Throughout this project APC was amazingly supportive. They expended dozens of hours trying to help in a situation that would not earn them money and was obviously not a part of any warranty. That was an award winning performance!

Next month, I'll present awards covering the dark side of the Force.

AND THE WINNER IS? - Part II

by Wayne M. Krakau - *Chicago Computer Guide*, March 1996 - *NewsWare*, September 1996

Wayne is a partner in Krakau Business Computer Systems, a systems integration firm that is a Novell Gold Authorized and Novell IBM Host Master Authorized Reseller. He has been working with computers for 19 years. He holds an M.B.A. in Marketing and a B.S. in Information Science and was originally educated in Computer and Information Systems Analysis Engineering. He holds Novell CNI and MCNE ratings and all five of the original LANDA Certifications. Wayne can be reached for questions, comments, or topic suggestions at (847)298-7695 or via E-mail to wayne@mcs.com.

This is the continuation of the presentation of my own series of awards, the SI (for System Integration) Awards. Last month I covered the positive categories. This month I'll cover the negative categories. As I previously stated, it's my award, so I make all the rules, create the categories, nominate the candidates, and choose the winners - and if you don't like it - tough!

The third SI Award is **The Joseph Stalin Award for Spreading Misinformation During a Technical Support Call**. It goes to Palindrome Corporation, for claiming in their final statement that certain Adaptec controllers were not only NOT officially "Novell Approved", but that they were never designed to be used in a NetWare environment. Since both Novell and Adaptec agree that the controller in question is approved for use with NetWare, the controller's documentation contains whole chapters devoted to NetWare installation, and the product's box is plastered with "Novell Approved" stickers, I consider Palindrome's assertion such a huge piece of misinformation that the old fact-bender, himself, Joe Stalin, would have been proud.

Did you notice that I used the term "final statement"? That's because of Palindrome's previous declarations, which were fairly outrageous, but not quite award-winning. This includes declaring their lack of support for a particular controller after repeated extended discussions of how to use their software with that controller. It also includes an after-the-fact assertion that the mere presence of this controller would cause their software to malfunction, even if that board had nothing to do with their software.

Since I have already documented these events in detail in a previous column, I won't go into the particulars of this incident, however, I will note that the people at Adaptec were definitely not pleased about this affair.

The fourth SI Award is for **The Automated Phone System From Hell**. This award had many candidates, but a late entry from Iomega Corporation simply outclassed the rest. It is a system seemingly taken straight out of Dante' *Inferno*. All it needs is a large rock to roll uphill while waiting on hold.

Over a period of months, I had been trying to get help for a problem with one of my own systems. Over time, the gradual deterioration of Iomega's phone system became apparent. I will admit, though, that I might have been a bit more tolerant of phone system shenanigans if I hadn't gotten consistently inaccurate answers to my questions.

As it stands now, each menu and sub-menu within their system is preceded by its own unique, interminably long message explaining Iomega's displeasure with your tying up of their valuable phone line and suggesting you get the hell off the line as soon as possible. (Note that this is my translation of their messages into clear English. The real messages aren't quite as direct - but they are close.) Another underlying theme is that if you are going to tough it out and stay on the line, they are going to make you suffer for it. The general subtext is, "We already have your money, so don't bother us!"

For good measure, the system is organized around multiple, confusing levels of menus, with no easy method of backing up if you go down the wrong path. Extra aggravation is guaranteed by delays within the menu system. When you make a selection, there are often annoying delays while you wait for the next level of response.

Throughout the system, there is no way to get to a human being either to ask which sub-menu to use to get to the right personnel, or to get a direct routing to a particular person on a follow-up call for a previously established problem.

Assuming you have a finite amount of time at your disposal, you might try searching for a way to leave a message. That would be a complete folly, since there is no messaging capability.

Also, while you are waiting, potentially for hours, keep in mind that you are paying for the call! This is what finally prompted me to give up on obtaining help. After repeated calls led only to incorrect info, I was eventually requested to purchase a new controller board for my system. The board didn't work. Before this, I had already spent dozens of hours experimenting, using the advice given by Iomega personnel. Even the updated software that they supplied via their bulletin board system (BBS) didn't help. The trick that finally worked was so obvious that I am still shocked by the fact that nobody at Iomega told me about it.

In an act I can only chalk up to pure stubbornness on my part, I tried to use the software that came with the new board, while still using my existing hardware. It worked! I had previously been supplied with the wrong software. After those long hours of experimentation and the frustrating sessions with their phone systems, it took only about five minutes of useful work to solve all of my problems.

Then I made the mistake of trying to contact Iomega to return the useless board. Based on my conversations with technical support people, I reasoned that the software that I ended up using was either available on their BBS for free or available on disk for a nominal fee. After repeated attempts to get through to a human, I finally had to give up after I realized that my phone bill was going to be so high that it would exceed the value of the returned board. (Hey, anybody want to buy a slightly used controller?)

The fifth SI Award is the **People Who Live In Glass Houses Shouldn't Throw Stones Award**. It goes to my company, or more specifically ME, since I'm the one who screwed up. While I am only human, and, therefore, make mistakes regularly, I usually don't do stuff this spectacular. I just found out that I generated an unrelated error while fixing a minor typo in a message line within the batch file that controls the nightly backup job for a longtime client.

This simple error, the accidental entry of a minus sign in a critical spot, caused the backup to look like it worked while, in fact, it didn't. The batch file was designed to first back up the network drive, designated F:, starting at the beginning of the tape, thereby deleting any existing data on that tape. Then it would back up the C: drive of that workstation, appending the data after the F: drive data.

While altering one of the messages that I designed to print out between the two backup commands, I accidentally altered the C: drive backup parameters so that it started at the beginning of the tape instead of appending the data. The revised job would back up drive F:, print a message and the log file indicating that the F: drive backup was successful, then overwrite the F: drive data that was just placed on the tape with C: drive data. Then, it would print a message and the log file indicating the C: drive backup was successful. There was no indication in the log files or messages that the F: drive data on the tape had been trashed!

Somewhere along the way, I got complacent and violated my own testing standards - those same standards I have repeatedly recommended in this column and in various speeches I have given over the years. I examined the printed output and the log files and decided that was enough to establish the validity of the revised batch file. Wrong! I should have used the manual commands available within the tape drive's software to actually view the data on the tapes. After publicly pontificating for years about the inadequate backup testing methods used within the personal computing industry, I did the very thing that I had condemned!

The scary part is that I can't remember when I last modified the batch file. My client could have been without a valid backup for months! I just don't know. I do, however, know that the current version of that file works. I manually checked the tape this time.

Now I can't wait to see if this admission of humanity on my part results in a decrease in the number of calls and E-mail messages that I receive every time I write about LAN cabling, asking me for permission to violate the laws of physics so that out-of-standard cabling systems will work.

COMPUTER TELEPHONY SAMPLER

by Wayne M. Krakau - *Chicago Computer Guide*, April 1996

Wayne is a partner in Krakau Business Computer Systems, a systems integration firm that is a Novell Gold Authorized and Novell IBM Host Master Authorized Reseller. He has been working with computers for 19 years. He holds an M.B.A. in Marketing and a B.S. in Information Science and was originally educated in Computer and Information Systems Analysis Engineering. He holds Novell CNI and MCNE ratings and all five of the original LANDA Certifications. Wayne can be reached for questions, comments, or topic suggestions at (847)298-7695 or via E-mail to wayne@mcs.com.

During my recent trade show expeditions, I have discovered several promising applications in the growing field of computer telephony integration. These are prime examples of the power and flexibility possible when you combine networks with telephone systems.

Computer telephony (pronounced with the primary accent on the second syllable and a secondary accent on the fourth) uses personal computers on a network to control and to augment the capabilities of a telephone switching system. Many products merely automate common phone system functions, while others "reinvent the wheel" by designing their own custom functions such as E-mail, contact management, and even network faxing. The three sample applications presented here excel in their ability to work with an organization's existing network applications, so the organization can use the best tool in each category and combine those tools with the telephony functions to provide enhanced capabilities.

I encountered the first two products at a private trade show put on by NEC America Corporate Networks Group (Irving, TX, 800-TEAM-NEC). During an interview with NEC officials, I was quite shocked to find an "old-guard" computing and telephone corporation with such an enlightened attitude. It was absolutely refreshing to hear them inadvertently paraphrase some of my own sentiments as expressed in past articles and speeches. Rather than go the route taken in the early years of the Desktop Publishing industry and currently being taken by most of the Document and Image Management Systems (DIMS) industry, NEC is pushing to be as non-proprietary and as flexible as possible!

In both Desktop Publishing and DIMS, the "old-guard" corporations, and even some newcomers, tried to make their systems as proprietary and "closed" as possible. In early Desktop Publishing systems, you often had to buy company-branded peripherals and stick with that company's limited word-processing and graphic-design capabilities. In current DIMS, not only are you often stuck buying special hardware, you may also be forced to buy proprietary E-mail, network fax, and even new network operating systems. In both cases, integration with other applications is forcibly and artificially limited by using nonstandard programming techniques within Microsoft Windows.

While NEC would obviously love to sell you a completely new phone switching system, with all of their own hardware and software, right down to the desktop, they have designed their computer telephony products to be integratable at both the hardware and software levels. For hardware, not only will their products work with NEC switches systems, but they will also work with any other brand that uses industry standard interfaces. For software, rather than reinventing the wheel or requiring proprietary systems, they have opened their system to outside software vendors, and have actually licensed several of them to market directly. I wish more companies had such a good attitude!

Two of these NEC-licensed programs caught my eye as being particularly powerful and flexible. Both Sixth Sense and Softphone are true client-server applications, with their server portions running as NLMs (NetWare Loadable Modules) on a NetWare Server, and their client portions running on the workstation as Windows programs. Each can be tied to other applications and control the passing of data among them.

Sixth Sense concentrates more on being the central application, with others revolving around it, while Softphone blends in more as just one of the gang. Sixth Sense is most often used as an enterprise database manager, gathering data from multiple databases and presenting it within its own screens.

Softphone acts more like a personal information manager, passing requests to other programs and letting them display their own data. Either style is workable, depending on how you run your organization.

Another product, CallWare from CallWare Technologies (Salt Lake City, UT, 801-481-8978) has taken a different tack. It can interface with and control existing phone switches, but it can also turn a NetWare server into a phone switch.

Like the others, it is an NLM-based client-server application. It can integrate with other applications in a style closer to Softphone than to Sixth Sense. You can run it either directly with your existing applications or with its own front-end software called Viewpoint. It also has a module available that integrates directly with Novell's GroupWise. CallWare's most distinct feature is its capability of using a multi-port phone-switching board to negate the need for an expensive standalone phone switch. The money and management savings can be tremendous.

I must admit that my initial attraction to CallWare, occurred when I came upon one of their representatives in mid-demo, just as he showed how the CallWare could be easily customized by the individual user, using easy-to-learn icons, to handle incoming calls differently, depending on the identity of the caller. I just love the idea of being able to categorize calls so that some of them ring through, others get appropriate different messages, and still others get blown off with a curt warning!

Since then, I have come to appreciate computer telephony for both the features explicitly mentioned here and for its customization potential. Computer telephony could be the next big "Killer App.".

MIGRATORY PATTERNS - Part One

by Wayne M. Krakau - *Chicago Computer Guide*, May 1996

Wayne is a partner in Krakau Business Computer Systems, a systems integration firm that is a Novell Gold Authorized and Novell IBM Host Master Authorized Reseller. He has been working with computers for 19 years. He holds an M.B.A. in Marketing and a B.S. in Information Science and was originally educated in Computer and Information Systems Analysis Engineering. He holds Novell CNI and MCNE ratings and all five of the original LANDA Certifications. Wayne can be reached for questions, comments, or topic suggestions at (847)298-7695 or via E-mail to wayne@mcs.com.

The gentle whoosh of the wings. The plaintive honk, honk, honk - splat. (Oh, well, I needed a car wash, anyway.) Yes, it's time for our feathered friend, the Canada Goose, to migrate. It's also time for our computerized friend, the LAN administrator, to migrate, too. Only this migration uses MIGRATE.EXE to upgrade to NetWare 4.1 and doesn't necessarily involve travel - unless, of course, the migration fails and your company asks you to leave!

The migration process is complicated, not always particularly logical, and is infrequently encountered, so my first bit of advice is the old standard, RTFM - Read The "Friendly" Manual! (You're correct if you guessed that I have replaced the real, unprintable word with the word "Friendly".) I know this is a trite suggestion, and that it isn't the "macho" way. (Manuals? Manuals! We don't need no stinking manuals!) Nevertheless, this may be the one time you should consider reading them. I've read the manuals, had formal training, and had lots of practice, but I am still discovering new wrinkles in the migration process.

While I will concentrate on using the MIGRATE command while upgrading a NetWare 3.x server to NetWare 4.1, you can also migrate from Banyan Vines, IBM PCLP 1.3 Extended Services, IBM LAN Server versions 1.0 through 3.0, Microsoft LAN Manager 2.0, and NetWare 2.xx, and NetWare 4.0.

Again, I have chosen to concentrate on a migration method that Novell calls "Across the Wire - Different Server", and will give short shrift to the "Across the Wire - Same Server" and the "In Place Upgrade" methods. I have been lucky in that clients have synchronized their NetWare upgrades with their purchase of replacement file servers, so the only one I regularly use is the "Across the Wire - Different Server" method. This is the easiest, safest, and, in terms of the performance of the server after upgrading, the most efficient.

For the "Across the Wire - Same Server" method, you need a workstation with enough storage space to hold all of the data from your server. This storage can be in the form of a large hard disk or a DOS-addressable (looks like a hard drive to DOS) tape drive. Since the data must be transferred twice, once to the workstation and once back to the file server, this method can take a long time. Add in the extra delay incurred by using the tape drive option (tape drives being way slower than disk drives) and you could set yourself up for a very long wait. This method is nearly an all or nothing proposition. If the upgrade fails or simply doesn't turn out as expected, you may have to completely restore the original server configuration and start over. If you are lucky, you may be able to jump-start the process in the middle - if you are sure that the "copy to workstation" portion of the migration completed perfectly.

The "In Place Upgrade" method is fast and relatively simple to implement, but it is the least safe method. It is a true all or nothing proposition, without any method for restarting in the middle. What's worse is the fact that you are stuck with the same disk block size that your old server had, usually 4K. NetWare 4.1 works much faster if you set the block size at 64K and let NetWare use its block suballocation feature to maximize efficiency. You lose that capability in this method. An extra complication is that old versions of NetWare used a different method of looking at the structure of IDE drives. If your DOS partition isn't large enough, you are out of luck. But, then again, this might be just the right time to take the opportunity tom switch to the more appropriate SCSI drives - in your new server, of course.

My favorite is the “Across the Wire - Different Server” method - and, no, it’s not just because I get to sell a chunk of hardware. It follows my long-held belief that during any budgeting for network expansion, always throw money at the file server first. If you need new workstations, for instance, first determine if it would be practical to replace the file server and retire the old one to workstation duties. This way, the server stays up to date technologically, thereby benefitting your entire network. If you extend this principle to include the evaluation of the timing of your NetWare upgrade, you can synchronize your hardware and software upgrades.

This method is the safest way to upgrade to NetWare 4.1, since your old server simply receives a long series of read requests. It is not altered. The required workstation only acts as a go-between to control the process and log progress and error messages. If part of the upgrade doesn’t work right, you can do it over. You can bring up both servers at once and compare them to double-check the results. You can even delay cannibalizing the old server until the new server has proven its stability, providing an extra measure of safety. On one occasion, I had to switch back to the old server on the day after the migration, so I know first hand how valuable this trick can be.

Next month, I’ll cover the process of migration along with my hints for making it easier, including some that apply to new NetWare 4.1 installations.

MIGRATORY PATTERNS - Part Two

by Wayne M. Krakau - *Chicago Computer Guide*, June 1996

Wayne is a partner in Krakau Business Computer Systems, a systems integration firm that is a Novell Gold Authorized and Novell IBM Host Master Authorized Reseller. He has been working with computers for 19 years. He holds an M.B.A. in Marketing and a B.S. in Information Science and was originally educated in Computer and Information Systems Analysis Engineering. He holds Novell CNI and MCNE ratings and all five of the original LANDA Certifications. Wayne can be reached for questions, comments, or topic suggestions at (847)298-7695 or via E-mail to wayne@mcs.com.

This article continues with coverage of NetWare's "Across the Wire - Different Server" method of migrating from a NetWare 3.x server to a NetWare 4.1 server. Since this process starts with a "clean" installation of NetWare 4.1, much of this information is also applicable to a completely new installation.

Before you install NetWare 4.1 on the new server, you must plan the installation. Version 4.1 is much more sensitive to filling the SYS volume than its predecessors, so a volume-creation strategy is needed. ("SYS" is always the name of the first disk volume defined in a NetWare server.)

First, decide what you want on the underlying DOS partition. (NetWare needs DOS to start.) You'll need enough space for the initialization files for NetWare, DOS, and whatever utilities you want. I've found that it is very handy to have diagnostic utilities on the DOS partition in case of an emergency. It is definitely not fun to fumble around with floppy disks or CDs in an attempt to fix a downed server while panicked users yell at you.

It's also handy to have a backup program loaded on the DOS partition. The single-user version of the network backup software you use is the best choice. That way, you can back up the DOS partition to the server's attached tape drive using the same software interface and hardware that you are used to. (This is assuming your tape drive is attached to the server - the most efficient arrangement.) In my initial visits to new clients I almost never encounter a server that has had its DOS partition backed up. Having that backup available can greatly reduce the time needed to reconstruct a badly crashed server.

You also have to keep in mind that if your server ever ABENDs (abnormally ends), and you need to dump the memory so it can be analyzed, you will need disk storage space equal to the size of your total RAM - and you may need the results of more than one ABEND for comparison. To avoid trying to get all of that data onto floppies, I leave as much extra on the DOS partition as I can spare.

Initially, for a 64MB RAM server, I used a 200MB DOS partition as a good compromise. Later, I realized that all of the systems that I sold had an IDE disk interface either built into the motherboard or as part of an Input/Output (serial, parallel, etc.) cards. Many RAID systems can't boot off the RAID drives. Even for those systems with bootable SCSI drives, it's a waste of space and somewhat more complicated to create a DOS partition. You have to split the disk manually, so you can't even use the preconfigured DOS partition that comes installed on most computers.

If you use the IDE disk that came bundled and preconfigured with many computers, you eliminate the extra labor. It also makes your system slightly more fault tolerant. If the SCSI drives malfunction, you have a head start on installation since your DOS partition might still be intact on the IDE drive. This gives you not only the base NetWare files, but also any updates and fixes that you have installed on that DOS partition. An added benefit is that if you choose to use a fault tolerance method that I have suggested for years, having a workstation/secondary file server (a workstation that is set up to be used as the file server if the primary file server fails), it makes it easier to make the primary and secondary servers similar.

If the IDE drive malfunctions, any available IDE drive could be cannibalized to use in the server. All you need is DOS and the single-user version of the tape backup software (along with a valid backup tape) to get your system up and running.

The NetWare partition can be split as needed. You need to set aside enough space for SYS to handle those programs that insist on putting files on that volume, plus a little extra for safety. In NetWare 4.1, print queues can be placed on any volume, so, if the migration is handled correctly, you won't have to worry about an errant print job running away and filling SYS.

After SYS, you can create other volumes as needed and name them what you want, preferably something logical. There are two schools of thought on whether to create one large volume or several smaller ones. I usually use one large volume since maintenance is easier and it eventually uses fewer drive letters (which affects both network efficiency and maintainability) when viewed from a workstation. (I'll save a debate on this issue for another article.)

Regardless of how many volumes you create after SYS, leave some extra space completely unallocated within your NetWare partition, just in case something happens and SYS gets filled. There is a copy of the INSTALL NLM (NetWare Loadable Module) on the DOS partition. With that program, you can add the extra, unallocated space to the SYS volume, immediately fixing your full SYS (at least once - you can't easily get that space back out of SYS). I usually leave at least 50MB unallocated.

During the installation process, make sure that you check off the option to copy the migration files. You can copy them manually, after the fact, but it's easier to let the INSTALL program do it for you.

The final preparation step is to use CompuServe, the Internet, Novell's Network Support Encyclopedia (NSE), or various manufacturers' bulletin board systems to get updated files. I've been known to tell people to keep SERVER.EXE and throw away the rest of NetWare. At times, that was not too far from the truth. These days, depending on the particular hardware and software you use, you may still have several megabytes worth of files to update. For the Novell files, I've found that using my NSE Professional Edition as the base and following up with a scan of Novell's NETWIRE (on CompuServe) as the easiest method. NSE Pro's files always lag slightly behind NETWIRE, so it's always best to cross-check for newer files. The most accurate way to make sure you've got the latest files is to download PATLST.TXT. It contains a complete list of the latest files.

Next month we'll get to the meat of the issue - that actual migration.

MIGRATORY PATTERNS - Part Three

by Wayne M. Krakau - *Chicago Computer Guide*, July 1996

Wayne is a partner in Krakau Business Computer Systems, a systems integration firm that is a Novell Gold Authorized and Novell IBM Host Master Authorized Reseller. He has been working with computers for 19 years. He holds an M.B.A. in Marketing and a B.S. in Information Science and was originally educated in Computer and Information Systems Analysis Engineering. He holds Novell CNI and MCNE ratings and all five of the original LANDA Certifications. Wayne can be reached for questions, comments, or topic suggestions at (847)298-7695 or via E-mail to wayne@mcs.com.

This is the actual migration portion of my series of articles on moving from NetWare 3.x to NetWare 4.1, concentrating on NetWare's "Across the Wire - Different Server" method. (In the spirit of a programmer I once worked with who, in the middle of a design meeting on a large commercial project said, "Enough of this planning nonsense - let's start writing code!" Luckily for him, we were in a windowless conference room, so he couldn't be forcibly hurled out the window.)

Before you use the MIGRATE command, you must be aware of one its most serious weaknesses. It can either move user objects without passwords, or move them with randomly generated passwords, but it can't retain their existing passwords. To overcome this, log everyone off your 3.x server and log in as the SUPERVISOR. Go to the SYSTEM directory on the SYS volume and run BINDFIX. Allow BINDFIX to remove outdated data and repair the three bindery files, NET$BIND.SYS, NET$BVAL.SYS, and NET$OBJ.SYS.

Rerun BINDFIX until it runs completely clean, and then run it once more. BINDFIX creates a backup of the existing bindery files with the suffix "OLD". By running it that one extra time, the resulting "OLD" files are copies of the clean files.

Copy these clean files, NET$BIND.OLD, NET$BVAL.OLD, and NET$OBJ.OLD, to the SYSTEM directory on the SYS volume of the new 4.1 server. Rename them back to their original names, using the "SYS" suffix. (It is not necessary to alter their attributes back to SYSTEM and HIDDEN.)

On the 4.1 file server console, load the INSTALL NLM (NetWare Loadable Module) by typing "LOAD INSTALL" at the colon prompt. Select "Directory options (install NetWare Directory Services)" from the menu. Select "Upgrade NetWare 3.x Bindery information to the directory" from the secondary menu. This option will open the bindery files in the SYSTEM subdirectory and use their information to add objects to your new server's NDS (NetWare Directory Services) tree. The USER objects will retain their old passwords!

The reason that this works is that any time MIGRATE tries to add an object to the NDS tree and finds a matching object already there, it merges the characteristics of the new object into the old one. If you instruct MIGRATE to skip migrating passwords, the password property of each migrated user will be empty and the password of the existing USER object will remain untouched. If you instruct MIGRATE to create random passwords, the new, random password would override the existing password, defeating the purpose of this bindery manipulation procedure.

To run MIGRATE, you need to log into both the old and new servers at once from the fastest workstation you can get your hands on. For the sake of speed, you might even want to isolate the two servers and the chosen workstation on their own cable segment so they won't be exposed to traffic from other devices on the network.

The connection to the NetWare 4.1 server must be a "bindery" connection. (The 3.x connection will be "bindery" by definition.) Although the documentation states that using VLMs (Virtual Loadable Modules) along with "LOGIN /B" is an acceptable method, don't believe it! (In Part One of this series I admonished you to Read The "Friendly" Manual - but I didn't tell you that I would guarantee its accuracy!) MIGRATE will not work properly unless the workstation uses the old ODI (Open Datalink Interface) shell programs, LSL, the MLID (the driver that the network card manufacturer supplies),

IPXODI, and NETX (or EMSNETX or XMSNETX). I found out about this documentation the hard way. I migrated a system and only afterwards found out that the users' rights didn't transfer to the new server. I have since located a Novell Tech Note mentioning this undocumented requirement for old drivers.

Also, don't forget that NetWare 4.x and NetWare 3.12 like to speak the Ethernet 802.2 protocol while earlier versions mostly speak 802.3 (keeping in mind that both terms are Novell's own nonstandard terminology). You must make allowances for this mix. If all of the devices on your network can speak 802.2, then this is the time to switch them. If some can't, you may need to load both protocols on the 4.1 server. Even if all of them can speak 802.2, you may want to temporarily load both protocols on the 4.1 server to allow for a transition period and to facilitate the simultaneous communication with both servers that is necessary for the migration.

Now, if necessary, create directories to hold the data transferred from the 3.x volumes to the 4.x volumes. If you don't have an exact match, volume for volume, between the servers, you must design a structure for volume placement. If, for example, you take my suggestion to create a limited SYS volume, and your old server had applications installed on its SYS, you need this type of strategy. After the migration is done, you can move the individual directories to their correct positions using either NetWare Administrator or FILER so you can retain their rights assignments.

At this point, copy the MIGRATE directory down to the workstation's local drive. This will keep the logging traffic off your server while MIGRATE runs. From this MIGRATE directory, run the MIGRATE command. Select the source (3.x) and target (4.1) servers. When you select what to migrate, you will be given the opportunity to select a specific target for each volume that you selected. That's why preparing the target subdirectories prior to starting MIGRATE is important.

Finally, unless you are willing to stare at the workstation's monitor for hours at a time, set the system NOT to pause on errors. MIGRATE considers even the smallest, most harmless warning message as an error, and will stop processing while awaiting manual intervention to get past these "errors".

After MIGRATE is finished, carefully review the seemingly endless log. This is the only way, short of stumbling into their aftereffects, that you will find potentially important error and warning messages. Then use NetWare Administrator (NWADMIN) to inspect the objects in your NDS tree. If you find any anomalies, such as missing rights, improperly assigned home directories, or unassigned group memberships, you can correct them prior to putting the server into production.

I have stopped using the MIGPRINT command for migrating the printing environment. By the time you figure out its bizarre syntax and manage to enter it without typos, you could have created all new print server, printer, and print queue definitions, with the added benefit of making sure that the new print queues don't end up on the delicate SYS volume. This is also a great opportunity to retire old cryptic device names and to reorganize the printing environment for ease of use.

With the migration and subsequent testing completed, you are ready to release the new 4.1 server into production. For safety's sake, leave the old server around, just in case there is a problem. That way, you can switch back to it or just manually copy over a critical file if necessary.

It's time for one last suggestion. It is very tempting to try to implement many other improvements and additions to the network at the same time as you upgrade your server. Please, seriously consider using a staged implementation schedule for these changes. Effecting too many changes at once can strain both your technical and training resources to the breaking point, with the end effect being an unreliable network and panicked users.

OS SOS

by Wayne M. Krakau - *Chicago Computer Guide*, August 1996

Wayne is a partner in Krakau Business Computer Systems, a systems integration firm that is a Novell Gold Authorized and Novell IBM Host Master Authorized Reseller. He has been working with computers for 19 years. He holds an M.B.A. in Marketing and a B.S. in Information Science and was originally educated in Computer and Information Systems Analysis Engineering. He holds Novell CNI and MCNE ratings and all five of the original LANDA Certifications. Wayne can be reached for questions, comments, or topic suggestions at (847)298-7695 or via E-mail to wayne@mcs.com.

It's time again for one of my semi-regular, semi-coherent, heavily opinionated rave-outs on what's happening in my little corner of the computer industry. This time I am covering some troubling aspects of operating systems (OSs). Remember, my motto is "Have Pulpit Will Preach!"

Windows 0.95

Yes, that's "Zero Point Ninety-five" - the traditional designation for the pre-release (also called "Beta") version that comes out before version 1.0 is ready. That pretty much sums up my opinion of this grandly overhyped product. While I am absolutely sure that the eventual "real", working version of the product will be spectacular, I can't justify asking my clients to risk their businesses and spend potentially huge sums of money simply to become Microsoft's guinea pigs.

My standard offering to those who ask about switching to Windows 95 is that I will charge them by the hour for all the time it takes to iron out any incompatibilities with their current software. So far I have seen many companies test Windows 95 and decide to hold off on implementation. I've also seen two companies implement it - and then erase their hard disks and revert to Windows 3.11 to avoid imminent bankruptcy. I wasn't involved in any of these "experiments", so don't think that they failed due some lack of expertise on my part. For the two complete disasters, I called Microsoft and referred these people to Microsoft Certified Professionals working for Microsoft Authorized Solution Providers. It didn't help.

A scary part of this is the amount of effort it took to get these referrals from Microsoft. After killing a couple of hours calling every Microsoft phone number that I could think of, I finally found the department that refers people to Solution Providers. The representative gave me the names of the three closest Providers. They were in Indianapolis, St. Louis, and Detroit!

While I often get sarcastic in my articles and speeches, I am usually able to remain reasonably civil in conversations. This time, I skirted the edges of civility. I absolutely badgered this poor woman for almost 45 minutes to keep retrying her search with different parameters until she finally gave me the names of the only five Providers she could find in Illinois. Luckily, two of them were near the companies in need of help.

Of course, I know that there are a lot more Providers in this area. I have run into more than five at various dealer association meetings. What I can't figure out is why nobody at Microsoft could give me a straight answer about which department handles referrals. Also, once I stumbled upon the right department, why couldn't they easily locate all of the Providers who have Windows 95 systems integration skills?

For comparison's sake, a few years ago, a client requested a referral to a NetWare specialist for his relative in another city. That city didn't have chapters of any of the associations that I belong to, so I called Novell at 1-800-NETWARE. Within three minutes, I was able to get a list of specialist companies to present to my client. I could have also used Novell's faxback system or their NetWire section within CompuServe. These days, I could also use their Web site. Even if I considered Windows 95 stable, it scares me to think about installing a new operating system with no easy access to help.

Windows NT

Microsoft has declared that Windows 95 is the ultimate desktop operating system for the home environment and Windows NT is the ultimate desktop operating system for the business environment. Sadly, the very reason that NT is superior to 95 is that it has the luxury of ignoring most issues of backwards compatibility (running existing software).

If you have the ability to upgrade all of your existing applications directly to NT, then you are ready for NT right now. If you have applications that don't have versions specifically made for NT, however, you may just want to keep the status quo - Windows 3.11. If you jump to Windows 95 (as only a tiny percentage of corporate computer users have) there is no official upgrade path. There is a standard way to move from 3.11 to 95 and from 3.11 to NT, but there is no easy way to transition from 95 to NT - and Microsoft has stated that there never will be!

The easiest decision for many is to stick with Windows 3.11 for now and, after they can test some future, more reliable version of Windows 95, decide which operating system to use. As a practical matter, even business users will eventually adopt 95 because they can't easily abandon their existing software and NT currently has only limited support.

While I believe in NT's future as a desktop operating system, I have never been enthusiastic about taking a desktop operating system, DOS, Windows 3.11, Windows 95, or Windows NT, sticking some extra features on top of it, and calling it a network operating system. (The same holds true for representing mainframe and midrange operating systems as network operating systems.)

The latest change in the justification arguments for NT Server have recently attracted my attention. The original argument eventually resulted in a tacit agreement that NetWare was great as a "mere" file and print server, but would never be as good as a strategic "enterprise" operating system. The answer for those, more important functions, was supposed to be Windows NT Server. The argument followed, that for the sake of compatibility and ease of management, you might as well standardize on NT Server for both types of file servers.

Since NetWare's NDS (NetWare Directory Services) has made it preeminent in enterprise networking, a new argument has appeared. Now NT enthusiasts are conceding (at least for the moment) that NetWare rules the enterprise, but are saying that you should use NT Server for all file and print services! Huh? Did I miss something here? I remember a section of *Philosophy 101 - Basic Logic*, dealing with circular logic. I think that idea applies here.

The really strange thing is that there are many reports of both arguments being used within a single company! NT supporters within a company have used the first argument to implement NT servers to control the enterprise activities of their companies.

When NT becomes unmaintainable in an enterprise situation (not an uncommon occurrence), they have to switch back to NetWare for their enterprise server. Then they use the second (and contradictory) argument to justify implementing NT Server for file and print server! It would probably be better for the companies involved if their NT enthusiasts quit and went to work for companies that originally used NT Server and have never had NetWare in the first place.

As an added bonus I read an article (obviously copied almost verbatim from a Microsoft press release) announcing additions that are being made to NT Server to provide many (but not all) of the features that are already in NDS. These new features are supposed to make NT the best choice for the enterprise. The articles went on and on about how Microsoft was going to bulldoze Novell into the ground with these amazing features. They suggested that administrators would be making a big mistake if they committed to NetWare. It was only at the very end that the proposed released date was mentioned - mid-1998! That means that, assuming Microsoft releases the software according to schedule (riiiiight), you get to wait two years to get features available in NDS right now. Gee, thanks, I needed that!

I know that some people will accuse me of being a rampant Novell fanatic on a mad spree of Microsoft bashing. Remember, however, that I never claimed to be unprejudiced. Also remember that I

regularly call Novell to task when they do something stupid - such as releasing NetWare 4.0 when it really wasn't ready. I'm not against Microsoft. I'm just against stupidity.

TIP SHEET

by Wayne M. Krakau - *Chicago Computer Guide*, September 1996

Wayne is a partner in Krakau Business Computer Systems, a systems integration firm that is a Novell Gold Authorized and Novell IBM Host Master Authorized Reseller. He has been working with computers for 19 years. He holds an M.B.A. in Marketing and a B.S. in Information Science and was originally educated in Computer and Information Systems Analysis Engineering. He holds Novell CNI and MCNE ratings and all five of the original LANDA Certifications. Wayne can be reached for questions, comments, or topic suggestions at (847)298-7695 or via E-mail to wayne@mcs.com.

This is a collection of tips for working with a NetWare server. They are the ones that I am most frequently asked about in my day-to-day encounters with NetWare LANs.

NETADMIN is the DOS-based administration tool for NetWare 4. It has a primitive and somewhat obtuse interface that follows only its own standards. Its capabilities to edit the System Login Script can, however, be very handy when testing, modifying, and retesting the script. It saves much time versus repeatedly going in and out of Windows to get to the Windows-based NWADMIN. The trick is, how to maneuver through NETADMIN to get to the script editing screen.

Here's how to get there. After entering "NETADMIN" at the DOS prompt (assuming you have the appropriate network rights), select "Manage" objects from the main menu and press "Enter". Then select the organization name and press "F10". (Warning! Don't press "Enter" - you must use "F10".) Now select "View or edit properties of this object" and press "Enter". This is the menu that gives you access to the Organization's properties. Finally, select "Login script" and press "Enter". At this point, you can edit the System LOGIN Script, press "Escape", and save or discard your changes similar to the way used within the old SYSCON utility on older versions of NetWare.

The Windows-based NetWare management utility, NWADMIN, has its own little quirks. For instance, if you double-click on the Country, Organization, an Organizational Unit, or a Volume, the tree structure under the object will alternately compress and then expand. To edit the properties of these objects, the shortcut is to highlight the object and then press "Enter".

While you are double-clicking on these container objects (Containers hold either other subsidiary containers or leaf objects. Leaf objects are the endpoints on the tree.), you may also note that the objects beneath them are resorted with each compression/expansion cycle. That's how newly added or renamed objects can be sorted.

To customize the sort order, click on "View" from the top menu and then click on "Sort by Object Class". This brings up a sorting menu listing object classes (a User Object, for example) in their current order. You can reorder them any way you like and optionally save the new sequence as the default. To get the new sequence to take effect, double-click on the appropriate container object as previously specified.

A quick way to give one object rights to another (make it a trustee) is to drag the prospective trustee and drop it on the object. This will bring up a "Trustees of" menu with the trustee name already filled in. Using the Shift and Alternate keys in the same way as they are used in Windows File Manager, you can select multiple objects to drag. Shift is for a contiguous series of objects while Control can handle noncontiguous groups.

Note that any discussion of rights and trustee assignments should include the warning that you shouldn't become so enamored of the incredibly flexible and powerful security system built into NetWare that you needlessly complicate your environment. Just because you can adjust rights at an incredibly detailed level doesn't mean you have to use every last option. Compare it to the maximum acceleration capability of a sporty car. It's nice to know that you have the acceleration capability when you really need it, but you don't have to go overboard by flooring it at every green light (unless, of course, you are male either under twenty-five or recently turned 40).

If you want to move an object, drag and drop it while holding the Control key. A “Move” menu will appear. If applicable, you can create an alias by checking off the appropriate option.

To set up a default user template, to be used by all newly created user objects, first click on “Object”
on the top menu bar. Select “User Template”. This will bring up a “User” properties window for a new user called “USER_TEMPLATE”. This object can subsequently have its properties modified at any time, just like any other user.

All of the standard user properties are available, including home directory (within “Environment”), password restrictions, login time restrictions, network address restrictions, mailbox, and group membership. After changing these settings, all newly created users will inherit them.

I have received many panicked calls for help from people who are in NWADMIN but who can’t see the tree. They just see a blank screen with the menu across the top. The way out of that leads to a secondary trick. First select “Tools” from the top menu. Then click on “Browse”. This makes a browse window appear.

The secondary trick is to select one of the container objects (such as organizations or volumes) and start another browse window. The new browse window will contain only those portions of the tree that are below the originally selected container object. You can make as many of these browse windows as you want, within the limits of screen acreage and visual acuity.

Vindication

Lest you think I went overboard in last month’s column on the circular logic frequently used to justify Windows NT Server (It’s a floor wax AND it’s a dessert topping, too!), the front page of the August 12 issue of *Info World* stated “Meanwhile, Internet tie-ins and bundled apps make NT Server 4.0 stronger then Version 3.51, but it can’t match NetWare in file-and-print or directory services.” Note that, for the sake of this discussion, directory services are synonymous with the enterprise capabilities that I mentioned. Hmmm. It seems that I am not alone in my opinions.

WE DON'T NEED NO EDUCATION

by Wayne M. Krakau - *Chicago Computer Guide*, October 1996

Wayne is a partner in Krakau Business Computer Systems, a systems integration firm that is a Novell Gold Authorized and Novell IBM Host Master Authorized Reseller. He has been working with computers for 19 years. He holds an M.B.A. in Marketing and a B.S. in Information Science and was originally educated in Computer and Information Systems Analysis Engineering. He holds Novell CNI and MCNE ratings and all five of the original LANDA Certifications. Wayne can be reached for questions, comments, or topic suggestions at (847)298-7695 or via E-mail to wayne@mcs.com.

No, the title doesn't indicate that I have completely given up on grammar. It's a line in a song (from Pink Floyd's *The Wall*) that always reminds me of an attitude that I first encountered when I was a programmer in the corporate/mainframe world. Note that the next line of the song is "We don't need no thought control." It's a scathing indictment of the British educational establishment, not the quest for ignorance that the line (and the attitude that I encountered) implies!

I had just revealed, to several more experienced colleagues during lunch, that I had begun taking classes to work toward my MBA (Masters Degree in Business Administration). They immediately laid into me with a host of derisive comments regarding the worthlessness - and even danger - in advancing one's education while trying to have a career in the computer industry.

The general argument was: "Who [sic] would you rather hire - someone with a lot of experience but little or no education or someone with lots of education (maybe with an advanced degree) and little or no experience?" They unanimously voted for the experienced applicant. They followed up this vote with additional nasty comments about those "overeducated" academic types.

When I speculated on the existence of an applicant who had both experience and an advanced education, they responded with blank stares. It was as if I told them that gravity occasionally reversed and the earth randomly spit people out into orbit. They couldn't conceive of such a person. They repeatedly fell back on their either/or logic.

The fact that we all worked in a department created to serve business users although none of us had any serious business knowledge did not impress them. Gathering such knowledge on an ad hoc basis, as needed, was perfectly satisfactory to them. They had no clue about why I would want to get formal training in business.

On the other side of the coin, at the same company, we were occasionally required to review incoming resumes. Usually, several of us (including the aforementioned colleagues) would each take a pile of resumes and evaluate them for possible interviews. Most of these resumes were from people with education but no experience.

These reviewing sessions were always lots of fun. An overwhelming number of applicants were convinced that real experience was only vaguely interesting, but the school environment was the peak of technological knowledge, and, therefore, was the only thing that mattered.

Their belief in the superiority of their educational backgrounds was evidenced by the goals stated in their resumes. They expected to start at a senior level and then rise to a computer management position within one or two years! In this company, that would be equivalent to getting at least six promotions. Just for comparison, the managers at that level each had seven or more years of experience - and a bachelor's degree. They had also accumulated multiple professional training courses over the years.

These resumes went over really well with my coworkers. The collision of these two diametrically opposed views was better than fireworks. I must admit that their remarks about the resumes often had me laughing so hard I nearly injured myself. If you have ever wondered whether programmers could be creative in a nontechnical sense - this proved it.

The moral behind this nostalgic junket is to watch out for extreme views and prejudices when planning your career in computing. The real key is a combination of education AND experience.

If you are starting out as a student, try to get any job you can that will get you hands-on experience with the type of computers on which you would eventually want to work. The proliferation of PCs, and especially LANs, provides a lot of opportunity. Many firms and departments within corporations simply do not need a full-time network administrator. Opportunities abound in traditional clerical school-supporting jobs for a person who can also help with computing needs.

This type of position is a great way to sneak some computer experience onto your resume, as long as you don't get a big head and think you are a big-time computer expert just because you've had a few computer classes or your professor thinks you're hot stuff. I have rescued many companies from employees with this attitude. Those employees often become ex-employees as soon as their employers realize just how much money they've lost due to screwed up systems. I have written before about the temptation in the PC and LAN arenas to overestimate your expertise. Avoid this or you'll lose your credibility - if not your job.

For those lucky enough to already be employed in the computer industry, continuing your education is the only way to stay on top. This could be the key issue when deciding who gets hired, who gets a raise, who gets a promotion, and, in this era of downsizing, who keeps their job.

The obvious channel for this continuing education is technical. Take computer classes. Get certified by the appropriate company or agency (CNI, MCSE, etc.). Self-study is mandatory just to keep up with this constantly changing business. Magazines, books, training software, trade shows, user and professional organizations, and seminars are all available to keep your skills updated.

The not-so-obvious channel is related to what type of company you work for. If you are in a general business environment like I was, an MBA, the ultimate business generalist degree, might be for you. At least take some business classes. Learning industry-specific skills is the ultimate ace in the hole. Most computer professionals are very concentrated in both their professional experience and their training. They have one very narrow technical field of expertise. Anyone with a wider view of computing as a tool to facilitate running a business has a distinct advantage.

Even the non-computer experience you may have picked up can be valuable. If, for instance, you worked in a clerical position for an insurance agency, you may have picked up enough insurance terminology and methodology to use as leverage to get a computer-related job in another area within the insurance industry. Industry-specific experience has been required for mainframe jobs for many years. That requirement is just now starting to filter down to PCs and LAN.

The most subtle ace to hold is people skills. The standard chicken versus egg joke in the computer business is "Are people with limited social skills attracted to computers, or does working only with computers cause people to lose their social skills?" This is the most polite way I can phrase this. It is usually expressed in a much cruder and insulting way. Sadly, it is not far from the truth. Many large computer installations have to assign selected employees to act as a liaison between end-users and the computing staff because they are genuinely afraid that the computer personnel will offend the users! (I was one of those liaison people.)

Anything you can do to maintain or enhance your people skills is a great asset in a job search. The ability to express yourself in English as opposed to "technicalese" is rare and valuable. Writing and speech classes are available and have direct application to many computer-related jobs. Organizations such as *Toastmasters* can give you serious practice in public speaking. Participating actively in a professional or social organization can polish up your social skills. I have even heard from people who took acting classes to get them out of their shell.

Some of the recommendations I have given are pretty obvious, but since I still find people getting the same kind of discouraging advice that I got early in my career, I felt that they were worth restating. Besides, remember that my motto is "Have pulpit, will preach."

VERTICAL TAKEOFF, Part One

by Wayne M. Krakau - *Chicago Computer Guide*, November 1996

Wayne is a partner in Krakau Business Computer Systems, a systems integration firm that is a Novell Gold Authorized and Novell IBM Host Master Authorized Reseller. He has been working with computers for 19 years. He holds an M.B.A. in Marketing and a B.S. in Information Science and was originally educated in Computer and Information Systems Analysis Engineering. He holds Novell CNI and MCNE ratings and all five of the original LANDA Certifications. Wayne can be reached for questions, comments, or topic suggestions at (847)298-7695 or via E-mail to wayne@mcs.com.

It's happened again. For the umpteenth time one of my clients had their business disrupted and possibly even endangered by their reliance on really shoddy vertical-market software. This trend is now at an epidemic level.

Vertical-market applications are written for specific industries or for narrow segments within an industry. The legal and medical fields are two of the most obvious examples of vertical markets. Different narrow segments within these disciplines could be divorce law versus product liability law, or individual medical office practice versus hospital supplies management.

There is no technical reason why vertical-market software products could not be created to be extremely high quality. They could be equipped with modern, user-friendly interfaces. They could be programmed using the latest productive languages and programming tools. They could be designed to use contemporary safe and efficient network-optimized non-proprietary databases. They could use a well-defined, properly documented file format. They could contain flexible report creation and modification capabilities. They could contain reasonable data import and export options. They could be delivered with enough user and technical documentation to avoid the necessity of making technical support calls. For that matter, they could be backed by some reasonably qualified and available technical support plan.

Sorry, Toto, we're not in the land of OZ, we're back in Kansas - or at least two states away in good ol' Illinois. Here in the real world, your vertical-market software wishes seldom come true. There is just too much opportunity for the ethically-impaired to take advantage of the situation.

Note that my use of the term "ethically-impaired" is specifically meant to include both those who plan to actively cheat people for fun or profit, and those who conveniently decline to recognize their actions as unethical - in this case, by refusing to acknowledge the limits of their training and experience (and perhaps the lack of limits of their egos). This refusal leads people to ignore the lack of weaknesses in their software, client training, support programs, upgrade practices, billing procedures, bug-fix policies, and customization charges.

Think of it in comparison to running over and killing someone with a car. Within the legal system, it matters whether you meant to run that person down or did it by accident as a result of, for instance, speeding. As a practical matter, it doesn't matter to the victim. He or she is still dead.

In the same way, business owners and managers don't really care about the motive behind the ethics lapses that are crippling businesses. The fact that the problems exist is the important thing.

Often, a software company gets the attention of some type of industry association. Since these associations are run by people in the appropriate industry, it is quite easy to impress them enough to get on their "approved" or "recommended" lists. Once you are on such a list, you are almost guaranteed a continual stream of customers. The nasty thing is that you don't have to update or improve your software to stay on these lists. If you got on the list ten or more years ago when there was little or no competition in a specific industry, your product can still be listed as "recommended" even though it is now archaic and no longer fits your customers' needs.

When a new client implements your product, they are normally comparing it to either totally non-computer methods or to some grossly inappropriate software that they have discarded. It is easy to impress them enough to get them to agree to supplying a written recommendation praising your product. Once they catch on and realize what they have gotten into, you don't care, because you already have their written endorsement and their company is already dependent upon your software.

If you have planned things carefully, you have made it nearly impossible to transition from your software to a competitive product. If your customer threatens legal action to force you to stop using their endorsement, or, when called by your prospects for their opinion, tell the brutal truth, you can threaten to withhold support, thereby maiming or destroying their business. That ought to shut them up. (This scenario is not fiction - I have seen it happen.)

Once you have established yourself as a "standard" within your narrow market, sales are a lot easier. Very few potential customers have access to serious research material regarding software. They will have heard of your product and perhaps just a few others (most likely at least as bad or even worse than yours). It will be easy to convince them that you are the only game in town.

I have lost track of how many times I have been introduced to a new client who has just chosen a vertical-market product after what is sometimes a multiyear search. Invariably, they are absolutely sure that their search was very comprehensive. It often takes me only a few minutes of research to find alternatives that they had not considered. These clients just didn't have the resources that a systems integrator has in terms of both qualified personnel and information availability, to do a thorough product search and subsequent product analysis.

In Part Two of this article, I will cover some of the other pitfalls of vertical-market software. In case you think there is no way to get industry specific features in software, make sure you read Part Two. "Vertical-market" doesn't have to mean disaster.

VERTICAL TAKEOFF, Part Two

by Wayne M. Krakau - *Chicago Computer Guide*, December 1996

Wayne is a partner in Krakau Business Computer Systems, a systems integration firm that is a Novell Gold Authorized and Novell IBM Host Master Authorized Reseller. He has been working with computers for 20 years. He holds an M.B.A. in Marketing and a B.S. in Information Science and was originally educated in Computer and Information Systems Analysis Engineering. He holds Novell CNI and MCNE ratings and all five of the original LANDA Certifications. Wayne can be reached for questions, comments, or topic suggestions at (847)298-7695 or via E-mail to wayne@mcs.com.

This is a continuation of my column on vertical-market software, that is, software written to (allegedly) meet the needs of a specific industry. Since the vertical-market software industry itself is heavily populated by what I refer to as "ethically-impaired", which includes both those who are actively unethical and those who simply decline to see the limits of their products or associated services, caution is advised.

I made the transition from "Big Iron" (mainframes) with the naive idea that I would see an end to the constant poor treatment suffered by users at the hands of computer people. After putting up with the lack of business design acumen, absolutely horrific user interfaces, unstable applications, and a paternalistic and, in many cases, downright dictatorial computer support staff, it's a wonder that mainframe users didn't switch back to quill pens and parchment to completely avoid technology.

I thought that the availability of easy to use programming tools along with market forces (survival of the fittest) would allow only the best products and companies to thrive. Why would users tolerate these characteristics in a free market? The reason is that this is not a free market. A "free market" requires that information is equally and easily available to everybody. (Hey, I just knew that those MBA classes were going to come in handy someday!) Only those who both have the technical and business background, as well as the appropriate research capabilities are really knowledgeable about the vertical-market software subsection within the computer industry. For everyone else, a computer is just another business tool - a mysterious "black box" designed by a quasi-religious sect of gurus.

My first encounter with bad vertical-market software really set me straight. A client's credit collection business depended on a particular vertical-market software package that was written using a then-popular database product and its associated language. After a recent upgrade of that software, the client's LAN began to fail during the nightly preparation for the next day's calls. (Think about the impact of a collection agency not having the information needed to make phone calls!) The server would crash and lock up completely. The software company claimed that the LAN was at fault.

Since my company designed and sold the LAN, I started looking into the problem. It is very difficult for "application" software (software that does a human-related function such as accounting) to crash a server unless it was designed to run within the server (and this product wasn't), so I thought that the LAN might really be at fault. After a couple of weeks of detailed analysis, I found some clues that led me to believe that is was the software's error. I managed to contact the head of development for the database company, and he agreed with my analysis and even filled me in on the details of his product.

A very old, and by then unsupported, version of his database would cause NetWare servers to crash if certain improper programming techniques were used. Correct programming methods would not cause this error. In subsequent versions of his database, the improper programming would be detected and automatically be replaced during the development processes. His solution was to switch to the latest version of the database. Doing that would simultaneously correct the problem, run faster, and make the software product eligible for support from the database company. There was even an automatic conversion routine included in the newest version of the database to accommodate upgrades from very old versions.

I told my client about the situation and he passed that on to the vertical-market software company. They denied everything (stonewall!) and gave him some nonsensical answers that seemed specifically designed to confuse a non-techie.

The next time I went to the client's office, I called the software company, using a trick that I have since found quite useful. I initially pretended to be a typical low-end computer geek from a computer store. I let the "programming manager" describe all the ways that the LAN was inadequate and all of the "stupid mistakes" that my client's staff made while using the software. I listened while he boasted of the advanced technology and outstanding programming techniques used in his product, and of the "many" satisfied customers they had. I then led him to a discussion of the details of his staff and their programming and network experience. (I can be very sneaky if I have to be to help a client.)

Then I dropped the bomb. I revealed my true background, including the fact that I had more years experience programming (prior to being a systems integrator) than the cumulative experience of his entire staff! I gave him a litany of programming errors that his company had committed on top of the one that caused the problem at hand. I also got him to admit that his company had been selling a network version of its software for over two years even though none of its employees had ever SEEN a local area network - not even at a trade show! They only tested their software on a multi-user system (one computer with multiple terminals attached) that claimed to be able to run software designed for a network. They felt that this was enough testing to make sure that their product was compatible with NetWare. Forget about ethically-impaired - they were ethically bankrupt!

In an effort seemingly designed to confirm my "ethically bankrupt" diagnosis, the company committed even more offenses. They refused to upgrade their product to fix the bug, claiming it would be "too hard" (contradicting what the database company manager told me). They fixed it by publishing a patch, for which they charged SEVEN HUNDRED DOLLARS! (I hope I'm not the only one who sees a problem with this.) To add insult to injury, they sent my client a bill for technical support charges! This was in spite of the fact that my client had a service contract with them (at $300 per month) which specified "unlimited" technical support. The software company representative explained that my client had used "too much" technical support time that month (gee, I wonder why) and, therefore, "deserved" the additional charge. He also stated that if they didn't pay the bill immediately, they would be refused all further support. Since my client depended on this software for survival, they paid the bill.

When my client later asked for an extra report, I was unable to help him since the software company encoded their data in a proprietary manner so only they could produce reports. Of course, they charged outrageous amounts for that service. A side effect of this encoding was that it was just about impossible to transition to a competitive product without manually typing in all of the data from this product. The software company made it clear that if they even heard a rumor that we were trying to crack their coding scheme, whether for reporting or transition purposes, they would cut off support. Since there were only a small number of experts in the underlying database, the client was unwilling to try anything. The risk of the software company finding out was just too great.

Two years later, when the software company finally upgraded their software to use the latest version of the database, they demanded that my client also purchase the underlying database update for double its suggested list price. When my client protested that they could purchase the database upgrade from my company (or, for that matter, just about any other reseller) at substantially less than list price, they were again threatened with the loss of all further support.

During these encounters with the vertical-market software company, that company didn't have the slightest notion that anything they did was in any way substandard or unethical. I hope you are seeing the trend. Next month I will analyze this trend and how to avoid it.

VERTICAL TAKEOFF, Part Three

by Wayne M. Krakau - *Chicago Computer Guide*, January 1997

Wayne is a partner in Krakau Business Computer Systems, a systems integration firm that is a Novell Gold Authorized and Novell IBM Host Master Authorized Reseller. He has been working with computers for 20 years. He holds an M.B.A. in Marketing and a B.S. in Information Science and was originally educated in Computer and Information Systems Analysis Engineering. He holds CNI, MCNE and CMTP ratings and all five of the original LANDA Certifications. Wayne can be reached for questions, comments, or topic suggestions at (847)298-7695 or via E-mail to wayne@mcs.com.

In the last column I described the circumstances of my forcible awakening to the potential hazards of vertical-market software. If you thought that politicians or used-car salespeople were the leaders in unethical behavior, you haven't had to work with the industry-specific programs, commonly known as vertical market software.

My most recent encounter with a vertical-market software company is a good example. A client already owned and depended on several products from this company and was in need of a payroll module. They were interested in this company's payroll module because the existing software from this company was totally proprietary, so that its payroll module was the only product which could communicate with my client's current system. Since I had interceded on behalf of my client with this software company on previous occasions where their existing software failed, I was not thrilled at the prospect of my client buying something as critical as payroll from them.

They had been unresponsive, and sometimes downright abusive toward my client. Their products had been written for a Unix environment and had been ported directly to a PC/LAN environment by the simple expedient of being run through a translating program. They had only Unix-oriented people on staff. Nobody in their company had any experience with LANs - they didn't even own one - and only a few had very light experience with standalone PCs. The software company also made it clear that their software would never be rewritten to run on that newfangled Microsoft Windows environment, since they seemed to think that Windows and its descendants are just passing fads. (Real men don't eat quiche - and obviously don't use graphical user interfaces, either.)

They repeatedly berated my client for being stupid enough to use "slow and inefficient" PCs and a LAN instead of their precious Unix. (Note that they didn't actually use the word "stupid", but the implication was very strong.) They made it very clear that it was my client's privilege to use their software and that they shouldn't complain about the lack of support, programming bugs, or poor performance because of their choice of a PC/LAN system over a Unix system.

In a call to this vertical-market software company I asked them about their payroll software, I asked them how they expected to support their products on LAN systems without any LAN experience or even an in-house LAN for testing. *They stated that they had no intention of providing any such support!* (Reader Alert: If you see nothing wrong with the previous sentence, please stop reading right now, go to the nearest college and take a course entitled "The Philosophy of Ethics" or the closest equivalent class that you can find. I got a lot out of it when I took it. After taking this course, you may resume reading.)

I was taken aback. It took me a few seconds to formulate a polite (that is, non-obscene) way of asking exactly how my client was going to obtain support. Their response was that they would refer PC or LAN related questions to the company that wrote the computer language that they used to write their software. Huh???

Since the beginning of the computer industry, when you buy a computer language and write software in it, the computer language company has no responsibility for directly supporting your customers. If, for example, I write a program in Microsoft's Visual Basic, sell it to someone, and that person calls Microsoft for support, Microsoft will refer them back to me, because I am the developer.

I asked the software company if they could provide evidence of a special contact that they had with the language company in which the language company agreed to override common practice and provide direct support to my customer. They told me that they didn't need one.

After they repeatedly stated that they saw no ethical compromise in selling a PC/LAN-based product that they had never tested (since they didn't own a LAN), I had to terminate the call. I just couldn't hold off telling them my opinion of their methods any longer (nor can I print that opinion in this publication), and I felt that doing so would jeopardize my client's ability to keep getting what little support was available from this software company.

In case you hadn't already guessed, the software company also charged outrageous amounts for additional reports and frequently charged for bug fixes. They use the local (that is workstation-based) version of Novell's Btrieve database, instead of the much faster and safer server version. Since they don't know anything about LANs, they obviously have no idea of the advantages of client-server databases. They merely ran the translation program from Unix to PC, printed some inadequate documentation, and shipped the product out the door. They wouldn't want to break their pattern of screwing their customers.

Next month, in the last of this series (finally!), I will cover the hows and whys (and in many cases, why nots) of selecting vertical-market software.

VERTICAL TAKEOFF, Part Four

by Wayne M. Krakau - *Chicago Computer Guide*, February 1997

Wayne is a partner in Krakau Business Computer Systems, a systems integration firm that is a Novell Gold Authorized and Novell IBM Host Master Authorized Reseller. He has been working with computers for 20 years. He holds an M.B.A. in Marketing and a B.S. in Information Science and was originally educated in Computer and Information Systems Analysis Engineering. He holds CNI, MCNE and CMTP ratings and all five of the original LANDA Certifications. Wayne can be reached for questions, comments, or topic suggestions at (847)298-7695 or via E-mail to wayne@mcs.com.

Yes, it's here - finally! I have finished beating you over the head with horror stories about vertical market software and am moving on to direct information (as opposed to being implied within the aforementioned horror stories) on what to do about it.

The main idea is to remember that vertical market software is a high-risk specialty tool. It is only to be used when absolutely necessary, and, even then, only to the degree required. It is only one step removed from the Never-Never Land of completely custom software. In fact, most vertical market software began as a single-client custom software project and, only later, evolved into a commercial product.

As a rough rule-of-thumb, a prioritized software search strategy that I originally published when covering custom software in 1992 still holds. For the sake of this list "packaged" software refers to standard commercial software that is commonly available, often in shrink-wrapped packages. The words "custom software" can be replaced by the words "vertical market software" with one adjustment. That is, you have to evaluate just how far a given product has transitioned on the path from custom to packaged software. If a company hasn't sold at least a few hundred copies of their software per year (not including upgrades), it is difficult to consider its software in the same reliability category as packaged software.

Here is my list in order of preference:

1. Single package.
2. Single package with application specific add-on packages.
3. Multiple packages linked with commercial packages.
4. Multiple packages linked with custom programs.
5. Modifiable single program with externally added custom modifications.
6. Modifiable single packages with internally added custom modifications.
7. Completely custom software.

The modifiable packages are most often accounting systems or contact managers (sometimes called sales or marketing systems). The best products in both categories include products that contain so many features that no single company or industry could use them all. They also include flexible and easily learned reporting. They have the additional advantage of huge customer bases. It is comforting to know that reductions in either software reliability or support quality will be protested not just by you, but also by thousands of other dissatisfied customers.

Even better, they are partially customizable right out of the box, and, through the purchase of programmers' toolkits, can be heavily customized. Because the products are based on Windows and its descendants, this customization uses the safer method of working through an external interface rather than the inherently riskier method of altering the base underlying programming code.

The companies that make these accounting and contact management products publish their data formats and actively encourage programming firms to create vertical market add-ons for their products. They provide sophisticated import and export procedures so it is easy to get data into and out of their programs.

By following these search suggestions, you can reduce your exposure to the risks of vertical market software. Only those components of your computer system that are truly unique to your specific type of business will be subject to this risk. A side-effect of this search technique it to restrict the vertical market software companies to the precise area in which (hopefully) they have the most expertise.

For instance, a programming team with lots of experience in the arcane aspects of medical insurance claims filing should not be forced (or probably even allowed) to write and support a complete, general-purpose accounting system, a pretty bizarre and arcane task in and of itself. The odds of such an accounting system actually following all IRS (Internal Revenue Service), AICPA (American Institute of Certified Public Accountants) and FASB (Financial Accounting Standards Board) guidelines while at the same time providing all the information needed to efficiently run a business are quite small. The odds of such a system matching the power, flexibility, and accuracy of a modern, full-featured packaged accounting system are absolutely abysmal.

On the other hand, a team of highly motivated specialists, with an appropriate ethical attitude, could make an accurate, elegant, supportable and, possibly most importantly, reliable claims filing system that works in conjunction with one or more packaged accounting systems.

Unfortunately, that "ethical attitude" phrase is too often a barrier. In a world without software quality standards, the buyer must determine just how much of a risk is being taken in the purchase of vertical market software. While the number of people consciously plotting to cheat customers is small, the ability of absolutely *anybody* to write and sell any software they want without justifying their qualifications makes choosing vertical market software very difficult.

Basically, nearly everyone who ever wrote a line of code, either professionally or as a hobbyist, thinks they are the ultimate programming expert. This phenomenon has been variously described in technical journals as "Ego-Based Programming" or "The Programmer as an Artist Syndrome". In treatises that I have personally traced back as far as 1969, this syndrome has been described as possibly the most severe problem in computing.

It is based on the difference between a "pure" artist and an artisan (or a craftsperson). An artist can create a work which creates an emotional reaction in the viewer and be considered incredibly successful. An artisan who creates a piece that is merely esthetically pleasing, but does not meet practical requirements, is considered a failure.

Put another way, if Norm of *The Yankee Workshop* (on PBS) created a chest of drawers that was merely beautiful and elegant, he has only met a part of his customer's needs. That customer would be justifiably dissatisfied if, after only six months of use, the drawers repeatedly jammed, a leg loosened, causing the chest to tilt, and the finish started coming off. If Norm refused to return that customer's calls for help, or if, after repeated calls for assistance, the chest was still not properly repaired, that customer would probably be downright angry. The artisan would not have lived up to his responsibility to provide a truly functional work.

In programming, the standards are much more slippery. What's worse, the customer is normally not equipped with the knowledge needed to realize that the program is substandard. It is common in computing, from mainframes through microcomputers, to believe that difficult-to-use, unreliable software is normal.

Even if the customer recognizes the problem as a general lack of ethics (in that the software company doesn't realize its own limitations), it is difficult to press home the complaint. In a war of dueling technical credentials, the customer loses.

As I've demonstrated in the horror stories in this series, once a customer has committed to a given product, there are practical, real-world limitations in what they can do even if they have someone with appropriate credentials running interference for them. Such are the hazards of vertical market software. May the buyer beware.

THE REAL WORLD

by Wayne M. Krakau - *Chicago Computer Guide*, March 1997

Wayne is a partner in Krakau Business Computer Systems, a systems integration firm that is a Novell Gold Authorized and Novell IBM Host Master Authorized Reseller. He has been working with computers for 20 years. He holds an M.B.A. in Marketing and a B.S. in Information Science and was originally educated in Computer and Information Systems Analysis Engineering. He holds CNI, MCNE and CMTP ratings and all five of the original LANDA Certifications. Wayne can be reached for questions, comments, or topic suggestions at (847)298-7695 or via E-mail to wayne@mcs.com.

No - this is not a review of a series of videos covering the lives of a bunch of whining, self-indulgent twenty-somethings (which is probably singlehandedly responsible for the bad reputation that this generation has acquired). This is about the real world conditions encountered in the computer industry.

In past columns I have criticized my own little niche (systems integrators) within the computer industry as well as many related fields including mail-order dealers, storefront dealers, value-adder resellers, cable installers, electrical contractors, various packaged software companies, operating systems software companies, custom software houses, and their close cousins vertical-market software companies.

Much of this criticism revolves around various obvious and subtle ethical violations. Obvious violations are those involving intent. If you use your superior knowledge about a specific aspect of computing to put one over on a customer, that's an obvious violation. If you hurt a customer's business by your own ignorance, or, more to the point, the inability to admit that you are beyond your depth, that's a more subtle ethical violation. Either way, my emphasis is on those activities and products which go beyond an arbitrary standard that exists only in my head.

In the real world of computing there are very few standards in the quality of either products or services. Everyone must make up their own standards - both the seller and the buyer. In an industry that moves this fast, nobody is capable of imposing standards.

Certainly the government is not able to keep up with the computer industry. Look at our air traffic control infrastructure. They are still using the same model computers that I first worked on in 1976, when they were already several years out of date.

Professional and trade associations can't do the job. Many of them actually encourage people to dabble in fields that, on paper, seem similar, but are not really related. Programmers are encouraged to work on computer hardware because they are supposedly "computer" experts as opposed to being programming experts. Copier dealers are encouraged to sell LANs because the sales cycles and prospective customer bases of the two industries are similar.

Even those competent within their chosen field can get away with providing substandard services and products. As long as the customer either does not realize the lack of quality, or, is convinced that the level of service they are getting is standard, they can survive and even prosper.

For example, I have repeatedly worked with customers who were used to having their mainframe systems crash frequently. To them, it was just a normal part of doing business. When their LAN crashed frequently, they just assumed it was business as usual. It wasn't until their LAN was redesigned and repaired that they realized that system crashes should only be a rare exception.

In the real world outside of computing, nothing works exactly as expected. When magazine articles are written about automobile reliability in terms of the number of defects found in new cars, have you ever noticed that number is never zero? On the better cars, it is usually somewhere around four to six defects per car. You expect that your new car will have an initial period of adjustment before it works just right. You hope that you don't catch a lemon, with difficult or impossible to fix recurring problems. You also hope you don't get one with a true design defect. (An exploding gas tank can really ruin your day!) After doing the appropriate research, you must eventually make your purchase decisions based on your

personal standards. You choose the risk-benefit equation for both the actual product (in this case a car), and the level of service you expect to get (from both the dealer and the manufacturer).

The same holds true within the computer industry. Even my favorite hardware and software products frequently contain bugs, manufacturing glitches, and, occasionally, even design defects. The best manufacturers don't maintain a zero DOA (Dead-On-Arrival) rate for their products. For all but the smallest programs, bug-free software is pretty much nonexistent. Indeed, companies who scrupulously avoid unethical actions are populated by mere humans who can potentially get in over their heads accidentally. (I'll plead No Contest to that one.)

Even reviews are subject to variation in quality. I always get a good laugh from reading reviews of computers, for example, in which the reviewer spells out in gory details how many different incompatibilities a system has, and then goes on to describe the many proprietary (read "even more incompatible and nearly impossible to fix or replace economically") parts within it. The reviewer will then give a glowing recommendation for that system! This is a prime example of having a very low set of quality-related personal standards.

I also get some jollies with software reviews. One example that I remember was billed as a review of "relational" databases, even though it included several products which are inherently non-relational or only followed a few of the rules that define what is relational. (Note that you don't need to know the definition of "relational" to understand this example. Just be aware that there are standards for this term.) To make up for the weaknesses in the non-relational and quasi-relational products, the main test procedures were reformulated to specifically exclude relational commands. This made the modern, relational databases look artificially bad in comparison with their older competition. The conclusion of the review was that a quasi-relational database and a non-relational database took top honors in a review of relational databases! This is an example of low ethical standards by the journalists involved.

The fact that nothing is perfect shouldn't cause you to avoid all purchases. It should, however get you to reevaluate your personal buying standards. What level of product quality, service, and support can you live with in a real world situation? Even more important, what ethical attitude can you live with? Careful compromises in the first category may have to be made, lest you end up using only pencils and paper to run your business. Compromises in the second category are not only unnecessary, they are deadly.

I am always surprised when I encounter people who continue dealing with a vendor after being terribly mistreated or being abandoned at a critical time. I'm not talking about occasional lapses in service due to simple human error from a well-intentioned vendor.

I'm talking about people being refused phone support because they declined to purchase a service contract - even for a comparatively trivial question about the disk drive switch settings on one of seventy-five newly purchased computers - and they still purchase from that vendor! Or, people who have had run-ins with bad custom or vertical-market software (as documented in previous columns), but have no plans to attempt to extricate themselves from these vendors' clutches. If you keep dealing with someone who has a proven lack of ethics (in either of the aforementioned categories), you are asking for trouble.

As I write this article, I am in the final stages of negotiating a vehicle lease for my business. The dealer who is closest to me also has the best service in the area. I know that because I have had the old car serviced there for several years. When I was given a quote by the dealer's salesman, it was missing all of the underlying numbers. I was only able to get those numbers by repeatedly demanding more information. Later, when I reviewed the information he gave me, I found out that the salesman gave incorrect information as to list price, invoice price, trade-in values, and residual value. His information contradicted not only every other quote I was given, but it contradicted information that had just been given to me by the manufacturer!

Rather than waste my time going back and haggling with this salesman, I will simply strike his dealership off my list and will continue negotiations with those dealers and leasing companies who dealt fairly with me - there are several. (Note that the deal will be completed before this article is published, so

don’t bother contacting me with leasing information!) The dealer has already proven uncooperative, at the least, and unethical at the worst. Even if I negotiate a fairer deal with him, how can I be sure he will act honorably in the future? I doubt that he will.

In the same way, if a vendor of computing products or services repeatedly acts inappropriately, or even once acts unethically, take your business elsewhere! I know it sounds obvious, but in spite of the fact that I frequently write about horror stories, there are plenty of vendors out there who are willing to act responsibly and ethically, and, within the limits of human capability, act within the scope of their experience and education.

Meanwhile, does anybody want to buy a really nice used car. It was driven only on Sundays by a little old systems integrator who never exceeded the speed limit or accelerated quickly. (Yeah, riiiiight.)

FOLLOW-UP, Part One

by Wayne M. Krakau - *Chicago Computer Guide*, April 1997

Wayne is a partner in Krakau Business Computer Systems, a systems integration firm that is a Novell Gold Authorized and Novell IBM Host Master Authorized Reseller. He has been working with computers for 20 years. He holds an M.B.A. in Marketing and a B.S. in Information Science and was originally educated in Computer and Information Systems Analysis Engineering. He holds CNI, MCNE and CMTP ratings and all five of the original LANDA Certifications. Wayne can be reached for questions, comments, or topic suggestions at (847)298-7695 or via E-mail to wayne@mcs.com.

Due to popular demand, this column is devoted to following up on topics I originally covered in previous columns. I am constantly bombarded with "whatever happened to" type of questions regarding my old articles, so I thought that I should reply to them in print.

The Shoemaker's Children Go Barefoot

After all of the carping I did in my series of columns cautioning people about the potential hazards of vertical market software, you would think I would know better than to get caught in the very trap that I described.

After years of creating quotes using what I refer to as "My Giant Spreadsheet From Hell," I purchased special quoting software specifically marketed as the salvation of companies that sell computer hardware, software, and related services. The software is Windows-based, and is allegedly (rampant O.J.-ism) capable of importing data supplied by many major computer product distributors. That data, along with manually entered information on products and services not already available in computer format, provides the base information for quoting.

My initial experience with this software was quite unpleasant. After numerous attempts, while using the documentation, I finally had to call tech support just to get it installed properly. After we finally got it installed, the technician admitted that the information I needed was not in the manual, in a README file, or on their Web site.

Once the product was installed, I assumed it would work as it did in its reviews - WRONG! It constantly blew up, taking the whole system with it. When it wasn't locking up my system, it was taking a ridiculously long time to import the distributors' data. After it took 22 hours to get 7% of one distributor's data imported, tech support finally told me to download a new, experimental patch for the program. That patch fixed the import problem. The program imported all of the data in just a couple of hours.

Now, I had a working program, but its usefulness was in doubt. It was written in FoxPro for Windows. While I have read in reviews that FoxPro can create true Windows programs, I haven't seen one. This quoting program, and all of the other FoxPro for Windows programs that I have encountered run in Windows, but use a unique and proprietary interface which only vaguely resembles a real Windows program. The mouse clicks, keystrokes, shortcuts, and even the scroll bars, work in their own bizarre non-Windows-like way. Even the fonts are proprietary and limited in their sizes. The program can't even use standard Windows fonts! (One of my customers purchased a FoxPro-based program which can't even be used with Microsoft's own video drivers!)

After wasting thousands of dollars of my time installing, debugging, and attempting to use this software, I finally gave up and went back to my spreadsheet. That's what I get for not following my own advice. (Note that mentioning this products name will not help any of my clients or prospective clients, but will help my competition. I've provided more than enough information for them to figure it out for themselves.)

Windows 0.95

I got a mixed reaction to my comments about this operating system. Some people told me about their successes with Windows 95 and were offended that I considered it an immature product (as in “0.95" being the product prior to version 1.00). Some even suggested that I might simply be a bit dim (a distinct possibility!), since I had so much trouble with Windows 95. Of course, they obviously forgot to read the part of the article that mentioned that in each case where I even thought I might be out of my depth, I immediately called in experienced, Microsoft-certified personnel - who subsequently also had lots of problems.

Many more told me of problems just as bad as mine. Since I frequently write about computer problems, I suspect that I attract horror stories, so I wouldn’t apply much statistical validity to these results. Just be aware that I am not the only one having problems with Windows 95.

Some of my clients are doing limited installations of Windows 95. Where possible, they are limiting their efforts (as per my advice) to new computers loaded with all Windows 95 specific software to reduce their risk. Even within those limitations, it can still be a challenge.

I have discovered that the same software companies that heavily advertise the new features added to the Windows 95 versions of their products, neglect to be so forthcoming with the news that they have removed old and trusted features from these same products, ostensibly due to the difficulty in implementing these features in Windows 95. You only get the nasty surprise that the features are missing after looking in vain for them in the product and its associated documentation, and finally calling tech support for help.

For example, STAC electronics has removed the outbound calling capability from the Host-Only version of their Reachout remote control software. Also, Artisoft removed the UPS (Uninterruptible Power Supply) interface software from the Windows 95 version of LANtastic.

One of the problems with DOS and Windows 3.x was that, due to limitations in the products and their documentation, software companies regularly wrote programs that violated Microsoft’s rules (including Microsoft’s own applications programs). A computer system would then act strange or even lock up because of these transgressions. Windows 95 was supposed to fix that. It includes many features which previously had to be manually created by software companies, has much more extensive and higher quality documentation, and Microsoft has stated that they are willing to withhold certification from software that doesn’t follow the rules.

Despite these efforts, nonstandard, rule-breaking programming techniques have become common. While these won’t necessarily be unsafe, in that they will disable your system, there are potential complications. For instance, Artisoft added a modem sharing feature to their LANtastic for Windows 95 product. To do this, they added multiple, nonstandard communications port definitions to the Windows 95 registry. This causes warning messages that claim that the system timer and the DMA (Direct Memory Access) controller of your motherboard are malfunctioning due to conflicts with some of these bogus COM ports.

If you try to remedy the situation by deleting the fake COM ports, everything goes back to normal - or so it seems. In reality, deleting these fake devices corrupts the registry so badly that Windows 95 must be reloaded. Of course, the fact that the conflict warning messages are really false alarms is not documented. You have to call tech support to find out what is really happening.

An additional complication is the existence of OEM Release B of Windows 95. Every new computer bundled with Windows 95 since around late 1996 has this version. The catch is that none of the software companies I have talked to have tested their products with this version. Even STAC, which is partially owned by Microsoft, hasn’t tested with OEM Release B. Luckily, most software that works with earlier versions of Windows 95 work with this version, but I have already started finding exceptions.

These exceptions are gradually being fixed, but it means that you could be a beta tester without even knowing it.

Next month I will continue following up with past topics. I need to mention that I didn't want to pick on either STAC or Artisoft in this article. They are just the most prominent and recognizable examples that I could find for the type of difficulties mentioned. Their respective technical support departments were helpful, cooperative, and, considering the type of problems encountered, appropriately apologetic.

FOLLOW-UP, Part Two

by Wayne M. Krakau - *Chicago Computer Guide*, May 1997

Wayne is a partner in Krakau Business Computer Systems, a systems integration firm that is a Novell Gold Authorized and Novell IBM Host Master Authorized Reseller. He has been working with computers for 20 years. He holds an M.B.A. in Marketing and a B.S. in Information Science and was originally educated in Computer and Information Systems Analysis Engineering. He holds CNI, MCNE and CMTP ratings and all five of the original LANDA Certifications. Wayne can be reached for questions, comments, or topic suggestions at (847)298-7695 or via E-mail to wayne@mcs.com.

This is a continuation of my article following up on the subjects of past articles.

Migratory Verticals

I have encountered a trend lately which overlaps two of my previous subjects, the continued sales of NetWare 3.12 in the face of overwhelming evidence of its obsolescence, and the difficulties in getting good support from many vertical market software vendors. Companies are installing new 3.12 systems not because of poor advice from resellers (the most common reason), but because they are dependent on vertical market software that has not yet been tested and is not fully supported in a NetWare 4.x environment!

I can understand a software vendor's reluctance to work with the old 4.0, 4.01, and 4.02 products. They were definitely "not ready for prime time." NetWare 4.1, however has been available for quite some time and is a proven, reliable product. Its descendant, IntraNetWare, is an even more capable derivative product. There is no excuse for neglecting this industry-standard product family for such a long time.

The client companies are naturally reluctant to become unsupported beta testers, so they buy decrepit old NetWare 3.12. Who can blame them? I can only hope that customer demand will eventually force these software vendors to update their products (if necessary) so they can test and approve them for use with NetWare 4.x.

The funny thing is that I found that the reluctance of some of these vendors to do this testing is based on horror stories about the original version of the workstation interface software, called VLMs (Virtual Loadable Modules), released with NetWare 4.0. VLMs long ago were stabilized, and are now so old that they are in the process of being replaced! If you are going to sell a NetWare-based program, at least you should keep up with the latest information.

ET Phone Home

After examining many computer telephony systems, and writing about several, I finally jumped on the bandwagon and got certified in CallWare (CallWare Technologies, Inc., 801-486-9922). After an intense week-long course, most of which was on computer telephony as a concept, not just CallWare, I have added some more initials after my name - CMTP, for Certified Multimedia Telephony Professional. Yes, I know, that acronym sounds overinflated and a bit too self-important. I liked it better when it was just called CNTE, Certified Network Telephony Engineer, since that sounds more to the point.

Naturally, I think it's a great product, or I wouldn't have wasted my time and money on the training. Since this is obviously no longer an unbiased opinion, I won't go into the details of this product in this article. I just wanted to mention the result of the research I've previously written about and didn't want to censor information about a quality product just because I've decided to sell it.

Windows 97

Windows 97 (or, as I like to call it, New Windows Version 1.00, as opposed to Windows 95, which I call New Windows Version 0.95), it has been delayed again. It is now being called Windows 9x, and, since its new due date is the first quarter of 1998, will presumably be renamed Windows 98. (If it ends up named Windows 99, I suspect the software development industry as a whole will rise up and lynch Bill Gates.)

The previous due date for this product was August 1997. This is a major disappointment for me since I expect it to have the same relationship to Windows 95 as NetWare 4.1 did to NetWare 4.0 - that is the real, working version of what amounts to a beta release to the public.

I would refer Mr. Gates (and anyone else who wonders why software projects are seldom delivered on time or bug-free) to the research done by IBM in the 1960s covering the original automation of the *New York Times*, along with associated research from that period regarding NASA's Agena project.

In its research IBM found that at some point, throwing more people (and presumably money) at a software development project becomes counterproductive, due to the intricacies of multiple person-to-person communications procedures. As you add more people past that point, the project actually slows down and will eventually grind to a halt!

The real limiting factors turn out to be organization and planning. Software development teams that suffer from an autocratic leadership and bureaucratic interference cease to function effectively.

Translated into real terms, it means that if management changes plans every few days, don't expect to make your expected release dates or to deliver reliable software. Those who best balance the competitive demands of the computer industry with the practical necessities of the development process create reliable products that are delivered on time. Those who base their sales mostly on overwhelming marketing spending are likely to deliver late, buggy products.

BACKUP, OR ELSE

by Wayne M. Krakau - *Chicago Computer Guide*, June 1997

Wayne is a partner in Krakau Business Computer Systems, a systems integration firm that is a Novell Gold Authorized and Novell IBM Host Master Authorized Reseller. He has been working with computers for 20 years. He holds an M.B.A. in Marketing and a B.S. in Information Science and was originally educated in Computer and Information Systems Analysis Engineering. He holds CNI, MCNE and CMTP ratings and all five of the original LANDA Certifications. Wayne can be reached for questions, comments, or topic suggestions at (847)298-7695 or via E-mail to wayne@mcs.com.

As I write this column I am in the middle of a project to attempt to recover data for a client who's hard drive on a specialty workstation drive crashed, but who had no recent backup tapes. Actually, their hard drive didn't just crash. It MELTED! The plastic gasket that helped seal the insides of the drive oozed out and dripped sticky goo all over my fingers.

The people at Ontrack Data Recovery (800-872-2599) later told me that this was the first melted drive they had ever seen. They suspect that a broken part rubbed against the spinning platters of the disk and the ensuing friction generated enough heat to turn most of the drive into mush. Naturally, no data was salvaged. Remember this drive the next time someone mentions the reliability of modern computer systems as an argument to lower their standards on backing up data.

This incident illustrates the difficulty of convincing nontechnical people who have never experienced a major computer disaster that they should take extraordinary steps, if necessary, to protect their data. In this situation, I feel partly responsible (and have made a major downward adjustment in my billable hours accordingly), since I didn't get tough and offer an ultimatum to the client. (Something like "Back up or I won't deal with you again.") I guess I'm just too soft-hearted to be that nasty.

I was lulled into a false sense of security by the fact that this client was quite cooperative in backing up their file server. When they replied positively to questions about regularly backing up their specialty system, I ASSUMED that they used the same definition of "regularly" that I did. A lesson learned: Never ASSUME that a techie talking to a non-techie is speaking the same language! (Obviously, assuming anything is dangerous.)

Since the backup of their file server was mostly automatable, and the backup of their specialty system was not (due to the space limitations of their backup tapes), the specialty backup was given a lower priority. It tied up a valuable computer for several hours during the day and was somewhat complicated to as well.

Just as I was about to sell them a new tape system with much more capacity and more sophisticated, easy-to-use software, the meltdown occurred. Luckily, all of the data will be recoverable from an auxiliary database that is on removable optical disks, but only after a major, time-consuming and expensive process.

If this was the only incidence of problems with backup plans, I probably wouldn't have been inspired to write on this subject. Sadly, at least one-third of all the prospects I meet have major inadequacies in their backup systems. Some of these problems are procedural, some are software-related, and some are hardware related. Since software and hardware difficulties should be detected by a proper methodology, I suppose that you could chalk them all up as procedural if you wanted to get picky about such things.

These days, there are many options for backup systems. Tapes are inexpensive when compared to data-recovery fees. More tapes mean more safety IF and only IF they are used in an organized manner. (Notice that I didn't say "planned", I said "used"!)

For hardware, there are a variety of choices. Travan-4 tape drives can hold up to 8GB (gigabytes) very inexpensively. DDS-3 DAT drives now hold 24GB. 8mm tape drives hold up to 40GB. Finally, for those who want serious data storage, DLT drives hold an amazing 70GB!

If that's not enough for you, try getting an auto-changer (a high-tech juke box). While not all of the technologies I've mentioned have auto-changers available for their highest capacities, manufacturers usually catch up over time. For instance, 24-tape auto-changers have been available for some time for the older 8GB DDS-2 DAT drives. I'm sure that some bright manufacturer is working, even now, on putting a 24GB DDS-3 drive in those changers.

Because of the problems inherent in backing up databases while they are in use, you may feel that you have to compromise your backup process. If your backup process takes so long it overlaps into normal working hours, or your staff works around the clock (hopefully with multiple shifts), you have a problem that used to be critical. Now, all you have to do is get Open File Manager from St. Bernard Software (800-ST-BERNARD). It is an NLM (NetWare Loadable Module) that runs on your file server. It allows live files to be backed up accurately by freezing the file and saving ongoing updates while the backup takes place. After the backup software goes on to the next file, the stored updates are applied. Once the software is set up properly, it is all automatic and transparent to the users of the network.

Finally, there is the matter of the backup software. Server-based software is the only way to go with NetWare 4.x. I prefer to use Backup Exec from Seagate Software (800-234-3793) and have had much more success with their products than any others my clients have used, but I know that others have different preferences.

A key fact is that workstation-based software can't back up the NDS (Novell Directory Services). Several manufacturers profess to, but, so far, their tech support personnel could not answer my questions well enough to give me confidence in their claims. Senior Novell support staff members have confirmed the inability to back up the NDS from a workstation.

Even based on the chance that some company figures out a proprietary method for backing up the NDS from a workstation, you would be stuck using a method that was not officially approved by Novell. In that case, just try to get serious technical support from Novell while you are rebuilding a server after a major crash!

If you are stuck using an existing workstation-based backup solution, then you can download a program called DSMAINT from the Web. This is an NLM that does Directory Service Maintenance. One trick it can do is to back up and restore the NDS onto a floppy or some other disk drive.

With all users logged out, you execute a procedure which shuts down the NDS. Then you do the actual backup. This only takes a couple of minutes. Then you reawaken the NDS. Finally you bring the server down and bring it up again.

If, like me, you don't read the directions as thoroughly as you should before trying DSMAINT the first time, you can set yourself up for quite a scare. I thought that I could just do the down and up sequence after the backup. I believed that this alone would reenable the NDS. WRONG! All that does is generate dozens of nasty messages which seem to indicate that you have permanently lost or at least critically damaged the NDS.

I think I gained multiple gray hairs from that stunt. Loading DSMAINT, reenabling the NDS, and doing another down and up sequence fixed the self-generated problem. (So much for the ace systems integrator!)

It is important that you remember to do another DSMAINT NDS backup any time you alter the NDS. If you change anything in NWADMIN or NETADMIN (the Windows and DOS administration tools for NetWare 4.x), you have altered the NDS. This is also true if you add, delete, or modify any program that directly affects the NDS. If in doubt, then run another backup. Scheduling regular DSMAINT NDS backups is an even better idea. As a manual task that can only be done when nobody is logged in, this might seem inconvenient, but it is not as convenient as being without a current copy of your NDS tree after a crash.

As I write this article, new, lower unemployment figures are being reported. If you don't want to make your own personal dent in those statistics, then I suggest you make sure your backup system - hardware, software, plans, and especially the execution of those plans - really works.

NEW SCHOOL DAZE

by Wayne M. Krakau - *Chicago Computer Guide*, July 1997

Wayne is a partner in Krakau Business Computer Systems, a systems integration firm that is a Novell Gold Authorized and Novell IBM Host Master Authorized Reseller. He has been working with computers for 20 years. He holds an M.B.A. in Marketing and a B.S. in Information Science and was originally educated in Computer and Information Systems Analysis Engineering. He holds CNI, MCNE and CMTP ratings and all five of the original LANDA Certifications. Wayne can be reached for questions, comments, or topic suggestions at (847)298-7695 or via E-mail to wayne@mcs.com.

I am absolutely thrilled to announce that the educational use of computers is finally maturing. Though the old adage about taking three steps forward and two steps back frequently applies, progress is being made.

Originally, school administrators (and probably most of the general public) thought that the proper way to teach children how to use computers was to teach programming languages. Students from high school (including me) all the way down to kindergarten were taught how to write programs. This (with the addition of a few games) was the working definition of computer literacy.

This parallels the earliest use of IBM PCs and their predecessors in business. Business users (or at least most business users), however, rapidly caught on to the fact that programming is a specialty best left to highly-trained and experienced (and one hopes, competent) professionals. Because of the growing availability of reliable (okay, make that "reasonably" reliable) commercial software packages, most business owners and managers didn't even need to employ programmers.

They recognized the computer as a business tool. They realized that the only required expertise was in application programs like word processors, spreadsheets, accounting systems, and specialty software needed for a particular type of business (vertical market software).

The only programming skill with a guaranteed payoff in day-to-day use was being able to write simple macros for word processors, spreadsheets, and the like. Additional expertise in programming may occasionally be handy for those with a particular aptitude for it, but there is the obvious risk of going beyond one's level of expertise when amateurs dabble in areas outside their chosen field. There is an even greater risk of being distracted away from your main duties while trying to fiddle with this fascinating technology. (Just ask someone who has dabbled in desktop publishing and found themselves tweaking their publication for hours at a time to get it "just right" while neglecting their normal job. Or, these days, someone who spends all of their time browsing the Internet without noticing that they've blown the entire day without getting any serious work done.)

Schools are finally realizing that computers are tools. They are reserving programming classes for those with a particular aptitude and interest in a career in computing or a related field. The computers are being used as tools to enhance the learning experience in various subjects.

You can see a parallel if you compare driver education classes with auto shop classes. The driver education classes are universally applicable to any student who wants to use a car. They don't have to know the gory details of how to build or repair one. Auto shop class, however, is reserved for students who want or need to learn these details.

Drill-oriented games are used for basic math, spelling, reading for interpretation, and typing. For these repetitive tasks, no human teacher can be as patient (or as entertaining) as a computer. Children can now enjoy what used to be a painful rote-learning task.

More sophisticated mathematical concepts can be illustrated by various graphing tools in addition to simple question and answer games. Seeing a graphic representation of a mathematical function is far more intuitive than just dealing with the raw formulas, especially with the computer doing the busywork of actually drawing the graphs.

Various publishing and drawing programs are used to teach artistic and design concepts. My own favorites, however, are the music programs. They range from simplistic little note-players to sophisticated musical notation trainers. What human being can instantly transpose keys - and print the results - in a matter of seconds?

There is a huge variety of educational reference material available. Besides the obvious encyclopedias, there are products covering just about any scientific or artistic subject you can think of. Chemistry; biology (without killing animals, no less); geology; geography; history; great works of art - it's all there.

On top of the educational requirements within schools, you also have to consider the business of running the schools. Word processing, general accounting, payroll, class scheduling, and document and image management systems are examples of the applications in use.

Because of this combination of a businesslike attitude toward educational needs and the increasing dependence on computers to run the school, school administrators now have to consider many of the reliability and performance issues that are common in the business world.

They are changing over to external devices on the server so they can easily switch to another computer that I call a workstation/secondary server if the main server fails. They are using mirroring (duplicate disks), duplexing (duplicate disks and controllers), and even RAID (Redundant Array of Independent Disks) systems to preserve their disk data. Duplexing and RAID also increase the speed of the system.

To handle the extra traffic generated by whole rooms full of multimedia-equipped computers accessing shared CD systems, they are splitting the network using multi-port network cards in the server to spread network traffic among multiple segments. Segments dedicated to the standard low-demand tasks involved with running the school can use standard 10Base-T concentrators (Token Ring is going out of favor), but those segments leading to classrooms full of multimedia workstations are set up with switched Ethernet concentrators with duplexed (that is two-way simultaneous signaling) 100Base-T uplinks to the file server.

Shared CD towers and occasionally even CD changers are used to avoid both the financial and management burden of purchasing dozens of copies of software. License-management software is used to keep things legal.

In larger systems, management software and even cable testing instruments are being used to keep the network running.

Internet access is becoming common. The standard technique of simply blocking access to inappropriate sites or whole categories of sites (such as "alt.sex") has proven inadequate due the ever-changing site names found on the Internet. (I have seen it reported that a Web site design firm temporarily turned a business Web site into a free X-rated picture repository as a way to test the ability of the system to tolerate a high number of hits!) A more thorough method is to add the ability to report on every location visited and the duration of each visit (to avoid penalizing for accidentally stumbling onto a site) by each student (and maybe even the teachers). If students know that Big Brother is watching, the threat of loss of computer privileges followed up by appropriate disciplinary action is usually enough to keep them away from inappropriate Web sites.

Between all of this Internet access and the more immediate access to floppy disk drives, industrial-strength virus protection at both the server and the individual workstations is a necessity.

Finally, bizarre, proprietary computers and no-name clones are getting rarer in the educational environments as administrators realize the cost of installing and maintaining those machines. An organization that is at least partially dependent on random contributions and/or special circumstance discounts can't afford to buy computers that have a low standard of compatibility.

School computer systems are resembling business computer systems more and more over time. While these complicated systems require additional planning and a structured outlook on design, schools are benefitting from these faster, increasingly reliable, and eminently more useful systems.

POWER PLAY

by Wayne M. Krakau - *Chicago Computer Guide*, August 1997

Wayne is a partner in Krakau Business Computer Systems, a systems integration firm that is a Novell Gold Authorized and Novell IBM Host Master Authorized Reseller. He has been working with computers for 20 years. He holds an M.B.A. in Marketing and a B.S. in Information Science and was originally educated in Computer and Information Systems Analysis Engineering. He holds CNI, MCNE and CMTP ratings and all five of the original LANDA Certifications. Wayne can be reached for questions, comments, or topic suggestions at (847)298-7695 or via E-mail to wayne@mcs.com.

There is an insidious enemy striking at our computer systems. It is a lot less obvious than the inadequate network cabling that I've previously documented. I'm talking about our inadequate power infrastructure. The ubiquitous power plug just isn't as reliable as most people assume.

The basic rules of computer power protection are simple and perhaps even obvious (though often ignored). File servers and other critical machines get high-end Uninterruptible Power Supplies (UPSs). Simple workstations that suffer from regular power losses get less sophisticated UPSs. Workstations that don't suffer from power outages get high-quality (as opposed to hardware-store specials) surge suppressors. Power conditioners fill the gaps between UPSs and surge suppressors for those computers that encounter irregular power without blackouts, though they often lose out to the price-performance ratio of modern low-end UPSs.

All phone lines going into any component of the network should be protected by phone-line surge suppressors. Network lines leading into file servers or other critical machines should have network surge suppressors. In all of this, I am specifying the minimums. If you can afford it, or you can't survive without it, you should provide even more protection to your precious network. There are plenty of statistics available to show that having a dead network is a great way to drive a company into bankruptcy.

Even with all of this protection, circumstances are conspiring against you. It starts with the increasing strain on the capacity of the U.S. power grid. Power is getting more variable than ever. Everything from toothbrushes to massaging lounge chairs is drawing more electricity than ever from a system capable of generating only a finite amount of electricity. Plus, all these new toys are dumping major league interference back into the power system. (Remember, AC means alternating current.)

During peak periods (which are getting more frequent) brownouts and even blackouts are common. Even if you don't get a true brownout, the power during these periods is more irregular than during non-stressful periods. Just be thankful you don't live in Florida. It has a terrible history of almost constant bad power.

Then the hardware and software makers conspire to force you into buying ever faster and more power-sensitive computers. Actually, these new machines are often downright twitchy when not babied with perfect power.

Just to make things even more interesting, the overall quality of the power supplies inside PCs has fallen over time. This affects their ability to keep running with imperfect AC power.

I found out about power supply problems indirectly. I purchased a power diagnostic board that fits in an ISA slot and gives a visual indication via LEDs of the quality of power within a PC. Less than a month after I bought this board, I received a surprise replacement from the company that made it. A letter of explanation accompanied the replacement board. It was a bit vague, so I called the manufacturer to get the full story.

The manufacturer acknowledged that the original board was set to reject any computer power supply that did not follow industry standards for power quality within a PC, and that the board itself was not faulty in any way. The real problem was customer satisfaction. Most of the purchasers of this product were computer resellers who made their own computers to sell. They used the boards to test and debug

those computers both prior to and after sales. These resellers were upset because the test board flunked all of their internal power supplies!

Another large contingent of purchasers of this power test board resold mail order computers. Again, they were annoyed that their computers' power supplies consistently failed the power tests.

Finally, resellers who either sold or recommended the low-end home and SOHO (Small Office/Home Office) lines of many prominent manufacturers for business and especially network use were similarly upset. Those machines didn't universally fail, but they didn't come anywhere near universally passing, either.

To satisfy its customers, the test board manufacturer made a new version of the board which had two settings. The first setting was essentially the same as the only setting available on the original board. It followed industry standards. The other setting tested at far below industry standards. It was explicitly designed to allow substandard power supplies to pass. So, the solution to using cheap junk power supplies is to lower the standards! (Don't you just love the computer industry?)

As a matter of interest, if you need to replace a dead or defective power supply, I have had very good luck using replacement supplies sold by PC Power & Cooling (Carlsbad, CA, 800-722-6555). They carry a variety of power supplies specifically designed to provide better cooling with less vibration than standard supplies. They even sell the tiny fans used to cool the main processor chip. Best of all, they are quite patient with people (like me) struggling to accurately describe an existing unlabeled power supply so that the replacement will fit.

Next month I will continue expounding on my various conspiracy theories, all of course, dedicated to that ultimate of conspiracy theorists, Oliver Stone.

POWER PLAY - PART 2

by Wayne M. Krakau - *Chicago Computer Guide*, September 1997

Wayne is a partner in Krakau Business Computer Systems, a systems integration firm that is a Novell Gold Authorized and Novell IBM Host Master Authorized Reseller. He has been working with computers for 20 years. He holds an M.B.A. in Marketing and a B.S. in Information Science and was originally educated in Computer and Information Systems Analysis Engineering. He holds CNI, MCNE and CMTP ratings and all five of the original LANDA Certifications. Wayne can be reached for questions, comments, or topic suggestions at (847)298-7695 or via E-mail to wayne@mcs.com.

This is a continuation of my column on the factors conspiring against you to damage your computer systems through bad electrical power. I've already covered the attack on your system from the power company and from the internal power supply. This article covers the part in between - your building.

Inside electrical wiring and its associated circuitry is generally designed (or in older building, redesigned) based on the National Electric Code (NEC), developed in the 1930's. This code was established to provide safety and standardization for electrical systems. The emphasis was decidedly on safety, proving the old adage that if enough voters die (here from preventable fires and electrocutions), eventually, even the most jaded, self-serving politician is moved to act. (Though being dead in Chicago has not always been an absolute disqualifier for voting.)

Microcomputers, however, are designed based on the much more exacting Institute of Electrical and Electronics Engineers (IEEE) specification for electric power. These standards have been a work in progress since before there were microcomputers, but, due to the time lag inherent in standards creation, have only recently been fully approved. Luckily for manufacturers of computers and related equipment, the proposed standard remained steady enough over time to use as the basis for their designs.

Now, here's the fun part. The electricians who install and maintain building electrical wiring and equipment, and the inspectors who approve them are trained based on the NEC, not the IEEE standards! To test for good power, they use the AC voltage function of a simple volt-ohm meter worth $20 to $120 depending on how fancy it is, which only provides an average (actually the RMS, or root-mean-square) of the actual voltage. To test for a good ground, they use the continuity function or, if they are really conscientious, the resistance function of a volt-ohm meter. This gives only the broadest indication of the true conditions. An appropriately trained, experienced engineer would use something like $15,000 to $20,000 worth of sophisticated equipment, along with a good deal of professional interpretation, to determine the reliability of an electrical system.

(A note to anyone who thinks I don't like electricians: My sister is one! She has worked for years in the dangerous environment of an international airport. Initially, when she wasn't diving for cover to avoid being sucked into or exhausted at by jets, sliced, diced, or propwashed at by propellers, or outright run over by planes, she was ducking passes or dodging derogatory remarks from her coworkers. After she proved herself, her treatment improved. Luckily, she is an ex-Marine, so her coworkers rapidly learned not to try physical intimidation, lest they test the reliability of their group health insurance.)

In older buildings, a simple lack of capacity is the most common problem. These buildings were simply not designed to handle mass quantities of AC-powered devices. They also can't handle the huge amount of interference generated and dumped back into the electrical system by the switching power supplies in PCs.

Surprisingly, even fairly new buildings sometimes run into capacity limitations. I have encountered situations in which, due to shortsightedness by the designers, some buildings built as late as 1993 aren't equipped to handle a modern LAN environment. I can't tell whether it is true ignorance, or simply a cost-cutting measure, but in the long run, it doesn't matter.

For years, I have run into various power problems, often due to these capacity limitations, and had cured them with UPSs and, occasionally, power conditioners. In fact, I still remember the sound of the first wheezing PC I encountered. Due to massive voltage variations in a very old building, the fans of the PC and its companion external Bernoulli Box (removable disk system), would speed up and slow down in tune to the voltage fluctuations. It sounded like I needed to send it to a quit smoking clinic. Another dead giveaway of bad power was the monitor. It danced a jig that would have made Michael Flately proud.

Once I added a line conditioner (the client was in poor financial condition and couldn't afford the extremely expensive UPSs of the day), the PC, B-Box, and monitor ran normally. The only minor problem left was the odd noise produced by the line conditioner. Every time it absorbed a large power surge, it would emit a loud, raspy growl. The client put this to good comedic effect when giving visitors a tour of the office by making sure that the guest stood only inches away from the line conditioner during the explanation of the computer system. When the line conditioner inevitably growled, the guest would invariably jump back as if attacked by an errant beast! In fact, when that client finally disappeared, I wondered whether it was due to a lawsuit instituted by the next of kin of a now-deceased guest with a heart condition.

My first encounter with totally intractable power problems occurred with a client who had solved the problem before my fist visit by installing a completely separate power system reserved for computers. One weekend, I was working alone on various upgrades and fixes. To make it more convenient to work on one PC, I moved it off a desk and onto the floor in the middle of the room. I had to move its surge suppressor to another socket to reach. When I turned the computer on, I found that it would only complete the boot-up sequence about one out of every three tries. Even when it did boot, it would randomly lock up without notice. It took a while for me to remember that there was a separate power system for computers. I had inadvertently plugged into a plain beige AC socket. When I switched to a red "clean power" socket, the computer functioned normally. (In case you're wondering, I didn't charge the client for my half-hour episode of stupidity!)

Next month I will continue with a spectacular example of power problems, along with various ways around them. Meanwhile, considering my comments about politicians, I'll prepare my records for the arrival of the government auditor.

POWER PLAY - PART 3

by Wayne M. Krakau - *Chicago Computer Guide*, October 1997

Wayne is a partner in Krakau Business Computer Systems, a systems integration firm that is a Novell Gold Authorized and Novell IBM Host Master Authorized Reseller. He has been working with computers for 20 years. He holds an M.B.A. in Marketing and a B.S. in Information Science and was originally educated in Computer and Information Systems Analysis Engineering. He holds CNI, MCNE and CMTP ratings and all five of the original LANDA Certifications. Wayne can be reached for questions, comments, or topic suggestions at (847)298-7695 or via E-mail to wayne@mcs.com.

In this final installment of the trilogy, having already covered the basics of bad electric power from both outside and inside, I'll conclude with one last real-world example and some possible defenses.

A couple of years ago, I was lucky enough to be referred to a client without a network. (It's rare for me to find one. Usually I meet new clients who need an existing LAN fixed.) I sold the client a 10-user NetWare LAN, including all of the hardware (server, workstations, etc.), all of the system software (NetWare, DOS, Windows, shared fax system, various utilities), and most of the applications software (contact manager, database, word processor).

In addition, since the client's system administrator had little computer expertise, my company became, in effect, their computer department. This situation is common for my company and imparts a great sense of responsibility for the client's welfare.

The first problem was a report of defective monitors. Some monitors were frequently blurred to the point of unreadability. We replaced those monitors. That didn't help. The new ones were also blurred. That led me to suspect some underlying power problem. The system administrator vehemently denied any possibility of a power problem. The building was only about two years old and he completely trusted the very experienced electrician who wired it.

Then the administrator reported malfunctions of the internal modem within their shared-fax machine (Essentially it was a nondedicated fax server - a configuration which I didn't like, but the client demanded). The mode would completely lock up and even turning the computer off and on didn't always revive it. It would remain locked up for random lengths of time ranging from a few minutes to several hours. Again I suspected power issues.

At the administrators request, I replaced the fax computer's modem, then the motherboard, then the power supply, and finally the I/O card (which was also the disk controller). Nothing helped. Simultaneously, I began consulting with the top troubleshooters at APC (American Power Conversion, West Kingston, RI, 800-800-4APC). The file server (which had no problems) was protected by an APC UPS (Uninterruptible Power Supply) and the workstations, including the modem's phone line, were protected by APC surge suppressors.

The APC engineer (and I do mean "engineer" as opposed to "technician") agreed that there were power problems. Even as I described the situation to him over the phone, I was receiving reports of random boot-up problems with other workstations on the LAN. He suggested that I try putting a UPS on the fax computer. The district manager for APC lent me a small UPS for this test. Even that didn't help.

Meanwhile, I was examining the system under the direction of the APC engineer. Through trial and error experimentation with electrical devices, I found that I could upset the display of several monitors by turning on a freestanding fan. I also found that the desk light in one cubicle prevented another computer two cubicles away from rebooting!

Based on my description of the wiring, the APC engineer noted several NEC (National Electric Code) violations were apparent, and that an improper ground was preventing the UPSs and surge suppressors from doing their jobs. After talking to the engineer, the administrator finally agreed that there were electrical problems. He declined the engineer's offer of a Power Audit, in which APC would send

someone to this site with the appropriate test equipment to fully analyze the power problems and work with the client to see that they were fixed. The client said the Power Audit would cost too much and stated that he trusted his electrician to fix things.

A month later, I was called upon by the same client to inspect more allegedly defective PCs. The client had spent $3,500 on rewiring the building and was sure that the power system, including the ground, was now perfect. He took his electrician's word for this. Now, all of the workstations worked perfectly - if they were turned on individually. If all of the workstations were on, the last couple that were turned on wouldn't boot. The administrator couldn't run all of them at once. At least one, and sometimes two, had to remain off.

The physical position of the malfunctioning PCs wasn't an issue. The order in which they were turned on was the only deciding factor. The last one or two turned on invariably malfunctioned.

Although I replaced several PCs, retested all of their LAN cabling, and spent several days (and many hours on the phone with the APC engineer) trying to find a way around these problems, the administrator dropped my company as a vendor, accusing me of having cheated him. Through all of this, I had spent over 100 hours of my time on debugging their electrical problems, all without charging them. Oh, well - I guess there are some situations in which you simply can't prove you are ethical. (By the way, the last I heard, they still can't run all of their PCs at once.)

This episode leads me to my primary suggestion for attacking mysterious power problems. Arrange for APC to do one of their Power Audits. I have heard rumors that other companies may start offering similar services in the future, but, so far, I am only familiar with APC's offering. Also, I have had several years experience using their products and accessing their technical support, so I know that they can get the job done.

In a Power Audit, APC personnel are sent to inspect the client's site and test the power system based on the NEC, the IEEE and other industry and government standards for wiring and cabling. They use
equipment like harmonic analyzers, branch circuit analyzers, and various power monitoring devices to test the power system. Best of all, they use their experience to visually inspect each component of the power system, looking for standards violations.

Most of the time, they find multiple standards violations, including violations of the NEC, which, as I stated in Part 2 of this series, was issued in the 1930's and is the main code on which electricians' base their work. I have no idea if these violations are caused by extreme cost-cutting in building and remodeling projects, by a lack of skill on the part of electricians, or if the same type of uncaring attitude that is the curse of the computer industry has affected electricians, or if there's some other reason.

I just don't have the background to comment in any serious way on the causes of these code violations. All I know is that some of these violations are safety related. It would be interesting to know the legal ramifications of an electrocution or death by fire traced back to this type of violation. I'm sure the personal injury lawyers would have a field day. (This is all, of course, in addition to any loss due to either downtime or inaccurate data from a malfunctioning power system.)

The main objection that I've heard against an APC Power Audit (or the equivalent service provided by some other company) is the expense. The cost varies so much depending upon the exact situation that I can't quote exact numbers, but it is expensive to send a professional out into the field with a pile of fancy equipment. I have heard a rough estimate in the $3000-$4000 range for one 25-user LAN, but that was an off-the-cuff estimate specific to one particular client, so I wouldn't be willing to budget based on that estimate.

Another, admittedly limited, option is to purchase a Perfector UPS from International Power Technologies (Orem, UT, 800-944-0356). It is a 1500 Volt-Amp digital on-line UPS. It runs completely without reference to a ground. It can take any voltage from 60 through 270 volts and can output your choice of any voltage from 90 to 270 volts at either 50 or 60 Hertz. It is the perfect choice for systems that are out in the field running off portable or otherwise nonstandard power systems.

At a $2,995 list price, it's a bit expensive for day-to-day office use, but for situations where the underlying power infrastructure is not adequate, it could be just what you need. There are limits to its usefulness. First, you usually don't find out about bad power until after you have already purchased a "normal" UPS. You would have to discard the old UPS or put it to use elsewhere.

Also, it can only protect up to 1500 Volt-Amps worth of equipment. If your whole building's power system is inadequate, you might have to buy a lot of these Perfectors to run your business. I still haven't figured out how to get around the network communications problems inherent in having multiple independent ground references from all these Perfectors. In the long run, it may be cheaper (and safer!) to just fix your power system.

Now, pardon me, while I check out this loose wire behind my computer. I'll have to crawl under my desk and stretch out to reach it. Hmm, I wonder if it's carrying any current. I'll be right back to finish this article. Zzzzzzzzzzzzzzzzzzzzzzzzzzzzzzzzzzzzzzzzzzzzzzzzzzzzzzzzzzzzzzzt!

SERVICE CONTRACTS & EXTENDED WARRANTEES - Part One

by Wayne M. Krakau - *Chicago Computer Guide*, November 1997

Wayne is a partner in Krakau Business Computer Systems, a systems integration firm that is a Novell Gold Authorized and Novell IBM Host Master Authorized Reseller. He has been working with computers for 20 years. He holds an M.B.A. in Marketing and a B.S. in Information Science and was originally educated in Computer and Information Systems Analysis Engineering. He holds CNI, MCNE and CMTP ratings and all five of the original LANDA Certifications. Wayne can be reached for questions, comments, or topic suggestions at (847)298-7695 or via E-mail to wayne@mcs.com.

A service contract for hardware, software, or both, is a valuable option, but one with more limits and potential real-world pitfalls than most would surmise. For many years, I have sold service contracts, referred my clients to third-party providers of service contracts, and, most frequently, worked on my clients' behalf in obtaining contracted services from vendors. I have found that they must be evaluated as carefully as any other type of insurance product, since the life and health of your computer system are at stake.

Hardware-only service contracts are the most common. They come in two major flavors. The first is the base contract offered by the manufacturer of the hardware while the second is offered by either dealers or third party firms. There is usually a certain amount of service bundled directly with the hardware, and often, an optional "enhanced" or "extended" warranty. These can include various combinations of on-site and depot (carry-in) service.

An example is one manufacturer that bundles first-year on-site and second-year depot service with its premium line of computers. (The low-end line gets a two-year depot warranty.) For a little extra, you can upgrade to a two-year on-site agreement. For more money, you can upgrade to a maximum of a three-year on-site agreement. The dealer has the option of bundling one of the optional contracts with the computer. The extended warranty options are so inexpensive, that, depending on the needs of my clients, I often bundle the three-year on-site warranty with the computers. At other times, I simply offer the extended warrantees as a separately listed option.

The major limit of this offer is that only parts supplied by the original manufacturer are covered. If I add options to the computer, for instance, they aren't covered. Since the reign of the single-brand LAN ended (at least according to the national press) somewhere around 1987, you have to either trust the individual manufacturers of any added components, or you have to purchase an all-encompassing third party service contract. Note that any third party contract should take into account any overlapping manufacturers' warranties when calculating fees.

Other manufacturers offer more extensive warranties, including some with maximum response times. Examples could be twenty-four-hour, eight-hour, or even four-hour maximum response times. Obviously, the shorter the guaranteed response time, the more you pay. Also, very few manufacturers will cover more than just their own products.

Third party providers are much more flexible on brand coverage, but they still have their limits. I have found them very reluctant to provide anything more than extremely restricted coverage on RAID (Redundant Array of Independent Disks) systems and high-capacity tape drives. Since most of the servers that I sell have both, and they are mechanical devices (as opposed to purely electronic items like circuit boards), the kind most likely to fail, the value of the service contract is limited.

Another limitation of time-sensitive service contracts is the difference in interpretation of the average individual versus that of the contract provider. For a four-hour on-site contract on a computer, most people think that a computer expert with a wide array of diagnostic hardware and software and enough spare parts to rebuild the computer from scratch will arrive on site four hours after the call for help is begun. WRONG!!!!! Something like this might happen. You notice a problem and call for help. After waiting on hold anywhere from a couple of minutes to over thirty, you talk to a gatekeeper (a term

I've appropriated from multiple articles that I've read on abuses by Health Maintenance Organizations.). This person is essentially a nontechnical or semi-technical call-screener whose job it is to confirm that you have a valid contract, that the general description of the problem fits those specified in the contract, that an authorized person (as defined by the contract - a common restriction) is making the call, and that the person calling can recite all of the appropriate contract and/or code numbers.

At this point, you may be asked who played Tom Cruise's second back-seater in *Top Gun*. (Answer: Tim Robbins of *Shawshank Redemption* fame.) You may also have to come up with the names of the Seven Dwarfs (I'll let you figure out this one.) - anything to prevent you from getting through to actual technicians. That's the gatekeeper's real job. The gatekeeper may also try to convince you that the specific problem is not covered under the contract because of mishandling on your part, incorrectly configured or buggy software, or the malfunction or incompatibility of non-covered hardware. A simplistic answer to your problems might also be offered. If you tell the gatekeeper your computer is smoking, then you might be instructed to get your computer to either use a patch or, at minimum, to switch to a filtered, low tar brand.

Once you've made it past the gatekeeper (after having twisted your phone cord into knots in frustration), you are given an incident number and placed in a queue waiting for a technician. (Heaven forbid you should forget the incident number! Then you might have to start the whole process over, or at least wait fifteen minutes while someone tries to cross-reference your company name in a poorly designed database.) Once you've waited an additional two to thirty minutes (or more), you get to talk to the technician.

The technician will usually have an encapsulated description of your problem as interpreted by the gatekeeper. It may or may not be accurate - flip a coin. The technician will then use his or her greater skill and experience to come up with new and interesting ways to disqualify your service request on technical grounds - ones that the gatekeeper (and possibly you) could never imagine.

You will be asked difficult technical questions about your system's configuration. You will also be asked for intimate details about your computer, such as the "rev" (revision number) of your motherboard. If you are not a computer techie yourself, and have no computer specialist on staff, you might not be able to answer these questions, so you will be ridiculed for your lack of technical expertise, and possibly disqualified from receiving service for not supplying enough details to the technician. At that point you may mention to the technician that you purchased the contract specifically because of your own company's lack of in-house technical staff. This response will usually be an accusation of having an uncooperative attitude.

If you are technically skilled in computers, you will be asked to do much of the diagnosis yourself. You will be expected to open the computer to find the motherboard "rev" and other details. You will be asked to experiment with swapping out various boards or changing switch and jumper settings - after hanging up, of course. If you tell the technician that you don't have spare parts and cannot, as a practical matter, shut down one of your working computers to get spare parts for board-swapping experiments, you will be seen as uncooperative. If you explain that you have the skills, don't have the time to experiment, and specifically purchased the contract to avoid spending time fixing computers, you are again labeled uncooperative.

Once the technician has chosen the diagnosis, he or she will try to just send the replacement parts to you and expect you to install them. After you've wasted more time convincing the technician to honor the full service agreement by sending a person to do this installation, the technician will actually start the process of dispatching a person with the appropriate (you hope) parts and tools to your company.

Next month I will continue presenting this admittedly worst case scenario and go on to cover software service contracts. Once I have covered the potential pitfalls of service contracts, you will know what to avoid.

SERVICE CONTRACTS & EXTENDED WARRANTEES - Part Two

by Wayne M. Krakau - *Chicago Computer Guide*, December 1997

Wayne is a partner in Krakau Business Computer Systems, a systems integration firm that is a Novell Gold Authorized and Novell IBM Host Master Authorized Reseller. He has been working with computers for 21 years. He holds an M.B.A. in Marketing and a B.S. in Information Science and was originally educated in Computer and Information Systems Analysis Engineering. He holds CNI, MCNE and CNTE ratings and all five of the original LANDA Certifications. Wayne can be reached for questions, comments, or topic suggestions at (847)298-7695 or via E-mail to wayne@mcs.com.

This is the second part of my coverage of service contracts and extended warranties. Realistically, much of this admittedly worst case scenario also applies to original warranties offered with new hardware and software.

In our last episode (organ music starts quietly in background and rapidly rises to a crescendo), our hero, the system administrator, had braved the tortuous path to getting authorized to actually receive an on-site service call. The unsuspecting administrator thinks that the system will quickly and efficiently be repaired, reconfigured, and back up in just a short time. (Riiiiight! And, I've got some riverfront property in Des Plaines to sell you to go along with that swampland in Florida I sold you last year! I'll even throw in some sandbags with it.)

The service person might arrive on time - or not. The phase of the moon is as good a predictor as any. I hesitate to use the word "technician" since that word implies at least a minimum amount of training and experience. The most common (and derogatory) designation is "board-swapper". A board-swapper is one who can open a computer, replace a circuit board or a drive, and run a simple software-based test to see if the swap worked. Getting a true technician is much less common.

The service person will attempt to fix the computer in question. If swapping out the particular item brought is not successful, the service person will call for help to get the opinion of a technician. If they find that the wrong replacement part was specified (not uncommon with a purely phone diagnosis, especially by someone who has probably never seen your particular configuration), then you will have to wait for another visit.

All of this assumes, of course, that the phone technician was competent in the first place - not necessarily a valid assumption. I have lost count of the times that I have had to explain things like basic PC technical concepts (like CMOS or serial ports) to an allegedly experienced support technician.

In one example that I witnessed in which lack of expertise was not an issue, the phone-based technician thought that a bad power supply was causing boot-up problems. He also acknowledged that a malfunctioning motherboard could be the culprit. Under pressure from his management to specify an exact solution and to authorize only the minimum number of replacement parts, he mentally flipped a coin and sent a service person out with only a power supply. You guessed it - replacing the power supply didn't fix the problem. The customer had to wait for another visit to get the new motherboard that actually fixed the problem. So much for time-guaranteed warranty service.

If the on-site person is a representative of a computer manufacturer that uses a highly proprietary design for its computer (as several top computer manufacturers do), and you have another brand component inside, or, occasionally, even attached externally to the malfunctioning computer, you might be directly or indirectly denied service. Sometimes unsupported or competing software can also cause a denial.

A direct denial of service occurs when the service person refuses to even start any serious work because of the presence of the offending product. More frequently, there is a more subtle action involving vendor-to-vendor finger-pointing. "If you hadn't used that part or software, our product would not have failed." "Our product is not the problem. That nasty unapproved hardware or software is to blame. Call its

manufacturer and stop bothering us." "We are not responsible for the configuration of the unapproved product. You configured it wrong and that is causing the problem."

Another variation on this trick is to get caught between the manufacturer who issues a warranty and the authorized service provider who actually carries it out. The service provider must follow very strict rules to get reimbursed, and has no power to adjust to the situation. You can get caught in the middle. Either the manufacturer authorizes service that is subject to interpretation and the service provider won't do the work for fear of lack of reimbursement, or the service provider has enough expertise to catch a mis-diagnose by the manufacturer and can't get authorization to repair or replace the right parts. (I've seen people replace a disk drive and controller when it was obvious that a damaged cable was the problem, just because that's what the service authorization paperwork ordered them to do!)

There are multiple variations on this theme. No matter which one you get, the final result is that you don't get the service you need (and presumable paid for either as part of the product price or as a separate purchase).

If you are lucky enough to be provided with the appropriate part, and it is installed correctly, then the real fun starts. The service person will wake the computer up only enough to test it using simple diagnostic software. Any custom configuration that you had, such as modification of the CMOS setup menu, for instance is lost. You have to put it back. If you had software installed on your hard disk, you have to restore it. If the repaired device was a server, you are also responsible for reloading the operating system prior to restoring the contents of the drive.

Most on-site service people sent under warrantees and contracts barely have enough expertise to put DOS and Window 3.x or Windows 95 on a machine using the standard, menu-driven install programs. Any modifications you've made to these operating systems to get your system to run properly are beyond them. Heaven forbid you should ask them to put NetWare or another network operating system back on your system.

Next month I will complete this worst case scenario and provide suggestions on how to evaluate service contracts and extended warranties. I will end this month's column with a special note to my readers who are involved in some aspect of computer hardware or software service. REMINDER: This is a WORST CASE SCENARIO that includes encapsulations of actual events. Please don't get bent out of shape and feel the need to write nasty letters because you are being criticized. If you provide competent, ethical service, I'm not talking about you! It's the bad apples that I'm after, and I think that you must admit, there are far too many of them out there. (He said as he stepped down from his soapbox.)

SERVICE CONTRACTS & EXTENDED WARRANTEES - Part Three

by Wayne M. Krakau - *Chicago Computer Guide*, January 1998

Wayne is a partner in Krakau Business Computer Systems, a systems integration firm that is a Novell Gold Authorized and Novell IBM Host Master Authorized Reseller. He has been working with computers for 21 years. He holds an M.B.A. in Marketing and a B.S. in Information Science and was originally educated in Computer and Information Systems Analysis Engineering. He holds CNI, MCNE and CNTE ratings and all five of the original LANDA Certifications. Wayne can be reached for questions, comments, or topic suggestions at (847)298-7695 or via E-mail to wayne@mcs.com.

This is the third part of my coverage of service contracts and extended warranties. This article covers the practical evaluation of these services.

In the example I have given in Parts One and Two, all of the problems are real, and are based on actual events that my clients have experienced. (All right, I admit that the trivia questions about the Seven Dwarfs and various movies were exaggerations - but only slight ones!) The object was to show that there is a very large difference in the public perception of service contracts and warranties (extended or initial) and the more complicated situation that exists in the real world.

To evaluate these contracts, you must balance their costs against the cost of downtime. (I know this is obvious, but it has to be stated.) This should include such esoteric numbers such as the loss of customers due to lack of service and delivery, the loss of vendors due to missed payments, in addition to the direct costs of having employees either fully or partially idle.

Another factor that is not often considered is the cost of having in-house spares. Even if you have no technical staff, having the right parts around can mean a large time saving in getting your system back up and running. For servers, for example, if you absolutely have to keep your network up, you can use dual servers. In this rather expensive option, a secondary server with a live duplicate of all of your data just sits there monitoring the primary server. If the primary stops working, the secondary server kicks in. With Novell's persistent client feature activated on the workstations, your employees might notice a short lag in network response, but will lose no data.

Because of the expense, only a few of my clients use true duplicate servers, but almost all of them use what amounts to a bargain version of the same concept. Most companies can make good use of what I call a "primary" workstation - that is a single more powerful computer used for database reorganizations, month-end closings, and other demanding tasks. I simply add extra RAM chips and appropriate disk controller cards to the machine, upgrade to a high-end network card, and the client has a "primary workstation/secondary file server". The only other requirement is to use external disk and tape drives on the primary file server. Even a "civilian" can be trained to swap the disk and tape drives and run the simple batch file needed to turn the secondary server into the primary.

While this doesn't eliminate the need for a service contract on a server, it may remove the necessity of purchasing an expensive, response-time specific contract. The computer maker's inexpensive extended warranty, combined with initial or extended warranties on the tape and disk subsystem - normally a well covered RAID (Redundant Array of Independent Disks) system, anyway - satisfies my clients' needs.

Spares are also handy for complicated devices such as RAID systems and tape changers. For tape autoloaders (also called changers), response time-based warrantees are often very expensive, so I have the client buy a single, non-autoloader tape drive to use while waiting for the repair of the autoloader. For RAID systems, I have the client purchase spare drives. Various controller cards and even motherboards can also be purchased to avoid having to pay for a high-end service contract.

For workstations, I include the inexpensive manufacturer's extended warranty unless the client requests otherwise. It may not be very cost effective to get a high-end contract on a large number of similar workstations. It is often simpler to just buy one or more extra workstations (and maybe spare

parts, too) and let them sit in storage until they're needed. Just don't succumb to the temptation to use them for new employees without replacing them, or you'll lose your safety margin.

Another factor that may help is to avoid proprietary servers where possible. If some of your workstations have the same motherboards or other parts as your servers, you could use them as a sort of spare parts farm. It is quite handy to be able to cannibalize a workstation to get a file server in a pinch.

The same holds true for the drives in your RAID system. If they are standard (and hopefully high performance) drives, you can get replacements almost anywhere.

If you need to purchase a service contract or extended warranty, make sure you know both with whom you are dealing and what you are actually getting for your money. For instance, I make it a standard practice to reveal to the client when I am working in a joint effort with another firm, whether I am bringing in the other firm to my client or they are bringing my company into theirs.

I constantly find that new clients were, at least initially, unaware that they were dealing with multiple companies when purchasing service contracts. (I've also seen this type of hidden multiple-company relationship in sales of cabling, programming, and specialized hardware and software.) Even computer manufacturers aren't always forthcoming with information on who will actually visit the client site when a call for warranty service is made.

Occasionally, I've even been asked to lie about myself and my staff being employees of other firms. I've declined, even when it meant losing business. (I am in no way saintly - and have witnesses who can prove it - I just have a conscience. Or, as the Church Lady might say, "Ethics aren't meant to be followed only when it's convieeeeeeeeeeeeenient.") It seems to be common practice to lie to the client in these situations. I saw that happen regularly in my old mainframe days. You never knew if a contract programming firm was sending you their own employee, an employee of another company, or an independent contractor.

One of the biggest trends in computing is the increase in partnering of firms to provide for their clients' needs. There's nothing unethical about it. It's been done since the beginning of commercial computing. That practice is even a part of the original definition of the term "systems integrator".

I do, however, have a lot of problems with the idea of hiding the fact that a multiple-company relationship exists. If I send someone to a client, that client has a right to know, from a liability, a tax, and an evaluation point of view, what my company's relationship is to that person. I don't want my client to find out that I am using a "third party" firm only after he or she sees an unfamiliar name on the truck. That takes away my client's ability to evaluate the purchase of services or goods prior to delivery of the service or goods. You must be able to independently evaluate the individual companies within this type of joint effort or partnership to make a valid judgment about whether you are going to purchase a service contract.

Even computer manufacturers aren't always forthcoming on who is providing warranty service for their products. Exactly what is an "Authorized Service Center", anyway? What are the minimum qualifications to be one? Or, if you buy through mail order, who provides the service? Certainly it's not employees of the mail order firm.

At the very least, you must find out the qualifications of both the people you will speak to on the phone and those who will come to your site to fix your computers. What is their skill level? How much experience do they have? Are the lighter-weight members of the team backed up by more technicallly qualified staff members? What are the criteria for moving up to the next higher level of technical expertise when a problem becomes intractable?

Finally, remember that vendors are not offering service contracts and extended warranties out of the goodness of their hearts - they need to make money. If their offer seems too good to be true - it probably is! Somewhere, somehow, corners are being cut to create an unrealistic offer. You can't get effective service from a company that is either unethical or out of business!

SOFTWARE & INTEGRATION SERVICE CONTRACTS

by Wayne M. Krakau - *Chicago Computer Guide*, February 1998

Wayne is a partner in Krakau Business Computer Systems, a systems integration firm that is a Novell Gold Authorized and Novell IBM Host Master Authorized Reseller. He has been working with computers for 21 years. He holds an M.B.A. in Marketing and a B.S. in Information Science and was originally educated in Computer and Information Systems Analysis Engineering. He holds CNI, MCNE and CNTE ratings and all five of the original LANDA Certifications. Wayne can be reached for questions, comments, or topic suggestions at (847)298-7695 or via E-mail to wayne@mcs.com.

After having been reprimanded for leaving out the specifics of software and what I call integration service contracts out of my series of columns on service contracts and extended warranties, I decided I had better cover the details of those contracts while the subject matter was still fresh in my mind. Obviously, the quality-of-service warnings in the previous articles also apply here.

On the pure software side there are different classes of contracts, with subtly different criteria for evaluation. The first is for contracts covering custom software. The equation is rather simple. If the cost and/or risk of using a pay-as-you-go method of support is too high, pay the price asked or switch to another product. For true custom software, you have very little recourse. If the software is very open and closely follows industry standards for hardware, software, and user interfaces, a third party might be able to generate a work-around to a serious problem without help from the developer.

The potential waiting time and cost of emergency service without a contract can be way too high, especially for smaller programming firms or individual freelancers. Small software companies must allocate most of their resources to programming related personnel. Their tech support resources are sized to handle the people who have support contracts plus a little extra to assist the occasional non-contract problem. If they don't estimate their tech support load correctly, purposely shortchange tech support to save money, or simply have a random burst of high-tech support usage, non-contract customers might not get a call back for many days.

If you decide that you must have a contract, you are totally at the software company's mercy. If they offer a ten dollar per month contract for the first year and switch to a $1000 per month contract price for the second year, you'll just have to pay it if you need that software for business survival.

For vertical market software, the equation is similar, but with a bit more flexibility. You must determine how common the software is, and how many dealers, VARs (Value Added Resellers), and consultants are available to provide support. If you have local support available, and the developer of the software has a large enough staff that they can handle ad hoc queries promptly, then you can either skip the contract or at least go with a simpler, less expensive version (assuming, of course, one is offered).

The key here is the openness of the software. If VARs or, better yet, customers, can get access to the source code (normally for an additional fee) or at least extremely thorough technical documentation, you can more easily risk going without a contract.

At this level (and higher), another cost factor comes into play. Many of these software contracts are bundled with a free supply of updates, upgrades, patches, and fixes. This alone might make the purchase of a service contract economically feasible.

Another factor is the dealer's part in the picture. Some dealers will bundle their services with the contract from the software company. If the dealer has access to the program code or technical reference material, and has extensive training and experience in the product, then it's worth evaluating this bundle. Even then, the real expertise resides at the software company, so, while the dealer may be an important part of fixing a problem, it takes a case-by-case judgement call to decide if a bundled contract is practical.

There is a closely-related type of software that is usually classified as subset of vertical market software, but really consists of specialty products with limited customer bases that apply to most types of

businesses. A document and image management system (DIMS) is a good example of that genre. Almost any business can use one, but, except for some inexpensive, simplistic products, they are not widely used. The market is also split amongst many companies, so no individual company has many customers.

When I sell a DIMS, I always recommend that my client purchases the highest level of support contract and renews it every year. The DIMS is often purchased in conjunction with a project to either destroy or at least store in an inconvenient location, the original documents that go into the DIMS. This makes the DIMS vital to accessing those documents. I don't bundle my services, even though it is the software company's policy to use VARs as their pipeline to their customers. I just can't justify to my clients paying my company enough to cover us in case a serious problem arises in a product whose internal functions are totally inaccessible to us. I have seen time-consuming disasters happen with other products, so I know the potential is always there. This is not to say however, that another VAR might see the equation differently. That's why there are entrepreneurs. Each sees things his or her own way.

For common commercial software, most small to medium size companies don't purchase service contracts. Larger companies use them to reduce the size of their internal technical staffs, or buy them as part of enhanced update and licensing offers. Third party phone support contracts are also available from specialty support companies.

Just as in the case of pure hardware contracts, there are things you can do to reduce or eliminate the need for service contracts. One is to educate your staff, especially any technical support personnel. Another is to companies' own technicians use when someone calls them for help. They contain all of the latest patches and fixes along with extensive troubleshooting databases, all updated every month. If you have anyone on staff capable of using these CDs, you can avoid most service calls.

Yet another is to establish a close relationship with your system integrator. Someone in that line of business should be inclined to avoid nitpicking you to death with billings for quick simple questions and should also be inclined to treat you fairly in billing for larger, more complicated problems. They also have the advantage of being intimately familiar with your particular computer system. (Warning: The preceding paragraph was an incredibly self-serving paid political announcement!)

The final category I want to cover is what I call an integration contract. This would cover at least all of the software, including system software (the desktop operating system, the network operating system, and various utility programs), and sometimes hardware as well.

This is another category in which I have turned down offers for service contracts from clients because of the potentially company-bankrupting amount of time that my staff would have to spend if a client had major problems due to, for example, an operating system bug. There are no bug free operating systems available - not DOS, Windows 3.x, Windows 95 (Yikes!), Windows NT, and even my old favorite NetWare.

Just as there are insurance companies that only do health insurance, others that do auto insurance, and still others who do both, some companies do offer integration service contracts. Just as with insurance companies, you have to decide whether they would support or abandon you if things get tough.

On the other hand, if your whole computer system is so badly designed that you feel you need to buy this high level of integration support for it, then maybe you should be dealing with a different integrator.

Over the last few years I have had to change health insurance companies due to massive price increases, poor service, after-the-fact changes in exclusions, and inadequate treatment (by an HMO as determined by second opinions from other doctors). I find it hard to believe that anyone should be less picky than I am about my health insurance when choosing what amounts to health insurance for their computer systems.

HOW TO PIRATE SOFTWARE - Part One

by Wayne M. Krakau - *Chicago Computer Guide*, March 1998

Wayne is a partner in Krakau Business Computer Systems, a systems integration firm that is a Novell Gold Authorized and Novell IBM Host Master Authorized Reseller. He has been working with computers for 21 years. He holds an M.B.A. in Marketing and a B.S. in Information Science and was originally educated in Computer and Information Systems Analysis Engineering. He holds CNI, MCNE and CNTE ratings and all five of the original LANDA Certifications. Wayne can be reached for questions, comments, or topic suggestions at (847)298-7695 or via E-mail to wayne@mcs.com.

I've named this column in honor of an elective class offered when I was in engineering school, "How to Lie With Statistics." That class was really a practical guide to detecting others lies and avoiding inadvertent lies of your own. This article is really a practical guide to avoiding piracy.

First, a major warning: I am not a lawyer and this is not meant to be a comprehensive legal document. Take it for what it is - a report from the field from an interested party. If you want serious advice, go to an attorney specializing in the software aspects of intellectual property. It's such a narrow sub-specialty within the legal profession that my legal clients (who are not specialists in software law) constantly ask me for my opinion. (Talk about desperation!)

Reseller-induced piracy is the initial introduction to piracy for some businesses. (A reseller in this discussion is one who purchases goods for subsequent resale.) To undercut competitors' bids, some companies will include software with computer systems and either not bill for them or bill at artificially low prices that nobody can beat. While I have encountered a couple of individuals who might have purposely asked fewer questions than they should have (Wow, a $200 Rolex watch - what a bargain!), the overwhelming majority of people were simply taken in due to lack of knowledge.

The reseller copies a single copy of software onto multiple clients' computers. The original might have been licensed to the reseller, or it might have been acquired surreptitiously from a previous customer. In one variation on this, the reseller sells the software to the customer, but simply walks away with it on the way out the door. In another, the reseller sells one copy of software to a customer and then loads it either on a file server or on multiple workstations.

Sadly, I have caught resellers of all types, from individual freelancers to big-name chains and franchises, playing these tricks on their customers with products ranging from PC-DOS to NetWare. Regardless, the customer is left with pirated software and all of the legal complications that come with it.

The most obvious evidence of reseller-induced piracy is missing media, documentation, licensing, or registration materials. Sometimes a polite phone call is all that is necessary to get your materials back. The hope is that the reseller is so surprised at being caught that the materials will be returned just to avoid further hassles. (As a reminder to my fellow native-born Chicagoans, it is now considered gauche to send an enforcer equipped with a Louisville Slugger to get your stuff back.)

If more than that is needed, you might be better off contacting legal counsel. You don't want to get into a situation where you get accused of extortion or assault because you phrased your request inappropriately in the heat of anger. (People exposed to the double whammy of being ripped off AND being left dangling as a possible target for both civil and criminal legal actions are usually VERY angry.)

If you don't get an immediate positive response from the offending reseller, your only way to cover yourself is to purchase the appropriate software and then go after the reseller for reimbursement later. Remember that, here, the appearance of impropriety is as bad as actual impropriety. If anyone could misinterpret your actions as having taken advantage of an offer that was obviously too good to be true, you could be accused of piracy. Even if you convinced everyone that you were an innocent party, you might be coerced into purchasing replacement software directly from the publisher at full list price as a show of good faith. Buy now or pay lots more later.

Another type of piracy is rationalized piracy. That occurs when someone thinks that they have a good excuse to pirate software and rationalizes the act. The price is too high. The license agreement is too restrictive. The software company doesn't provide adequate support. The dealer mistreated me. The head of the company already has too many billions of dollars. (That's a tempting one!) My boss told me to do it. (My personal favorite, also known as the Nuremberg Defense. It didn't work after World War II and it won't work now.)

I'm sure that these excuses will go over great in law offices and courtrooms. I've also heard that one of the qualifications to be a federal judge is to have a great sense of humor.

Next month I will continue with this guide to piracy. Meanwhile, my parrot made off with my bottle of rum, but I've got my cutlass at the ready and I'll hunt him down.

HOW TO PIRATE SOFTWARE - Part Two

by Wayne M. Krakau - *Chicago Computer Guide*, April 1998

Wayne is a partner in Krakau Business Computer Systems, Inc., a systems integration firm that is a Novell Gold Authorized and Novell IBM Host Master Authorized Reseller. He has been working with computers for 21 years. He holds an M.B.A. in Marketing and a B.S. in Information Science and was originally educated in Computer and Information Systems Analysis Engineering. He holds CNI, MCNE and CNTE ratings and all five of the original LANDA Certifications. Wayne can be reached for questions, comments, or topic suggestions at (847)298-7695 or via E-mail to wayne@mcs.com.

This is a continuation of my column on piracy. It's not really about how to do it. It's about how to avoid it, as viewed by a non-lawyer working out in the trenches. This article is meant as a practical guide, not a legal treatise. If you try to treat my advice with the same authority that you would attribute to true legal advice from a lawyer specializing in this field, you have my sympathy.

The first type of piracy on this month's menu involves custom or semi-custom software. ("Semi-custom" being my term for a base, standard set of software modules that have customized enhancements or add-ons.)

The critical factor for this type of software is the basic concept of ownership of software. If, as a normal part of an employee's job, as described in that employee's official job description, while on company premises, using company supplied computers and software tools, a person writes new software for the company or enhances or adds onto existing company-owned software, then the employer owns the software. If the employee writes software under any other circumstances, the outcome could be debatable, though, from the examples that I've seen, the employer has an advantage in these debates, depending, of course, on the circumstances.

If the person writing the software is not an honest-to-God, W-2 issued, employee, the software belongs to the writer, not the company, and that person can decide exactly what rights the company has to the software. The only way around this is through a written agreement giving either complete ownership or specific rights of use of the software. A verbal agreement is explicitly excluded from qualifying. Even if you videotape the verbal agreement, only a written agreement counts.

What does this esoteric theoretical nonsense mean in the real world? It means that the freelancer or even large programming firm that you hire to write software for your company owns that software and has complete control over its use unless you get your rights to that software stated in writing. Just paying for the labor of writing is not enough. Licensing issues must be spelled out in the written contract. If actual ownership, as opposed to licensee's rights is at issue, a separate clause should be added to the agreement showing that ownership is transferred in return for something specific, such as a particular dollar amount allocated to the purchase. (The amount doesn't have to be large.)

All of this means that, without such an agreement, the software that you thought you "bought" is really only has a limited implied license. Even if you have some type of agreement, it would typically be written by the software company and be so totally one-sided in their favor that it would be worse than useless.

If your firm adds more standalone PCs or network workstations at your existing site, then you might be covered under that implied license, subject, of course, to a judge's opinion. If you install the software at another branch of your company, install it on a second LAN at your existing site, or decide to market the software, then you are in trouble. Even if you have a copy of the source code, you might not have the right to modify it. Even showing it to another programmer in an attempt to bail yourself out of a lack-of-support problem might be prohibited. Reverse-engineering the data structure in order to extract data to allow you to either switch to a different program or simply create custom reports can be prohibited. Those one-sided license agreements can be killers. From my observations, there are very few

written agreements in the real world, and most of those that do exist are severely skewed in the writer's favor.

Frankly, if you do anything to upset the writer, you are in trouble. Suppose you are extremely dissatisfied with the response to bug reports and/or enhancement requests from the writer and you decide to withhold payment - even long after the original software was "purchased" (at least from your point of view). The writer can go after you in court, and might even be able to stop you from using the software at all. Even without going to court, the writer can simply stop helping you. Without a current, up-to-date copy of the source code and the legal right to use it, you are stuck. Most people caught in this situation don't have any copy of the source code, much less a current one.

This series of columns was inspired by an episode of this type. I received a call from a potential client whom I had met previously while demonstrating specialty software. They were having reliability problems with the firm that both maintained their network and wrote their custom software. Coincidentally, they had plans to improve the software and sell it to similar organizations in the future.

I asked the obvious questions about ownership, contracts, and licensing rights, and got all the wrong answers. As I write this, the potential client is seeking legal help. They don't have the full source code for the program that drives their organization and the developer won't give it to them. Their contract explicitly forbids them from showing the portions of source code that they possess to non-employees. They are also forbidden from reverse-engineering the file format to extract data. They have no distribution rights, but the software company can sell the software (which was designed based totally on information provided by the client) to anybody, including competitors of this client.

The kicker is that even if the client's legal council (not software specialists) reviewed the contract prior to signing (I haven't had a clear answer as to whether they reviewed it or not), they probably wouldn't have fully understood the implications of the contract. I would like to assume that they would have realized that a specialist was needed. Based on the "Pay Me Now, or Pay Me Later" concept, some legal specialist will eventually make money out of this situation.

Next month, this piratical adventure will continue. I haven't caught my parrot yet, or retrieved my rum, but I have had a "pointed" lesson on not running with sharp objects - like my cutlass. Ouch!

HOW TO PIRATE SOFTWARE - Part Three

by Wayne M. Krakau - *Chicago Computer Guide*, May 1998

Wayne is a partner in Krakau Business Computer Systems, Inc., a systems integration firm that is a Novell Gold Authorized and Novell IBM Host Master Authorized Reseller. He has been working with computers for 21 years. He holds an M.B.A. in Marketing and a B.S. in Information Science and was originally educated in Computer and Information Systems Analysis Engineering. He holds CNI, MCNE and CNTE ratings and all five of the original LANDA Certifications. Wayne can be reached for questions, comments, or topic suggestions at (847)298-7695 or via E-mail to wayne@mcs.com.

Yo ho! Yo ho! The pirates' life for me! Oops! That's the *Disney Channel* take on of piracy. (So, maybe you were expecting the *Discovery Channel* or at least *A&E*?) This is actually another installment in my series on the ins and outs of software piracy. As in the previous articles in this series, any attempt to use these semi-coherent ramblings as actual legal advice would be foolhardy at best, borderline suicidal at worst.

License misinterpretation is, in my experience, the most common piracy. Often, people don't read the license that comes with their software, and I have a hard time blaming them. Assuming that you can find the license agreement among the tons of advertising and other extraneous literature commonly bundled with software, and, that your eyesight is strong enough to read the fine print, you will often see something like the following (With apologies to Jerry Pournell who published a pair of similarly themed paragraphs in his *Chaos Manner* column in *Byte* magazine, sometime in the mid 1980s.):

Paragraph 1 - This software is such absolutely worthless junk that it cannot be reasonably assumed by anyone with an IQ higher than a grapefruit, to be useful for any actual business purpose, or, for that matter even physically safe to install on a PC. Therefore, even if this software cause flames to shoot out the back of your PC, trashes all of your data, and permanently puts you out of business, you can't blame us.

Paragraph 2 - This software is so incredibly valuable to human civilization that if you so much as think of violating this license agreement, you agree to pay the authors (Nobel Prize, Pulitzer Prize, and Sainthood pending) enough money to design a time machine and go back in time to prevent your great-great-grandparents from meeting, thereby eliminating both your existence and your ability to pirate this software. Assuming that you run out of funds prior to the completion of said research, in lieu of the above penalties, you will surrender your firstborn child, who will then be properly trained to respect sacred agreements such as this.

The opening paragraphs are often so outrageous that people assume that the whole license agreement is completely unenforceable nonsense (definitely not a safe assumption in these litigious times). And, after you've made it through these paragraphs, you have to try to figure out under what conditions you can actually use the software. As I've said in the first installment of this series, even the attorneys that I have as clients can't interpret many licenses. That's why so many users give up and just skip reading the agreement.

Instead they rely on secondary sources like ads, summaries printed on packages, software company salespeople, their reseller, or the implications of the title of the product (Network Version, 5-User Pack, etc.) Any one of these sources can provide misleading information, and none of them is guaranteed to be a valid legal defense in case there is a dispute (though they might make good "mitigating circumstances" in the right circumstances).

From my point of view, as a reseller, I've frequently received misleading information from my distributor sales rep, my distributor dedicated product manager (for the software in question), and even

from my sales rep at the software company. To solve this, I use the same technique that I use when talking to tech support representatives. I simply won't get off the phone until I get a straight answer.

For licensing issues, I request the exact product numbers and quantity that I need for the specific client in question. If that answer doesn't correlate to the verbal explanation of the license restrictions, I ask for another explanation. To double-check this information, I request a fax of the written form of the license agreement along with any other related policy statements. Sadly, I sometimes have to make multiple phone calls to various people within the software company to get an accurate answer. However, using this method, the worst case would be that I still didn't understand the exact specs of the license restrictions, but at least the client would have the right combination of products to keep them out of trouble.

Next month I will continue with some of the variations on licensing that you might encounter. Meanwhile, I negotiated a truce with my parrot. He can't get the bottle of rum open by himself, and he can't drink very much anyway (besides, there's nothing more pitiful than a drunken parrot who has fallen off his perch), so we agreed to share the rum.

HOW TO PIRATE SOFTWARE - Part Four

by Wayne M. Krakau - *Chicago Computer Guide*, June 1998

Wayne is a partner in Krakau Business Computer Systems, Inc., a systems integration firm that is a Novell Gold Authorized and Novell IBM Host Master Authorized Reseller. He has been working with computers for 21 years. He holds an M.B.A. in Marketing and a B.S. in Information Science and was originally educated in Computer and Information Systems Analysis Engineering. He holds CNI, MCNE and CNTE ratings and all five of the original LANDA Certifications. Wayne can be reached for questions, comments, or topic suggestions at (847)298-7695 or via E-mail to wayne@mcs.com.

This is yet another episode in my continuing saga of piracy on the high Cs (as in the programming language C) - and you thought that *War and Peace* was long. As in previous installments in this series, this is a practical guide, not qualified legal advice.

This article covers the various types of licensing agreements that I have encountered. I will assure you ahead of time that I am not making up any of these agreements. The descriptions are based on actual products.

The most basic license is for free software, or its closely related companion Freeware. They're not the same thing. Freeware is a more specific term with its own special rules, while just plain, free software can have any rules. Except in Unix environments, free software usually comes with some restrictions on commercial use. It is often aimed at individuals. If companies want to use it, they might have to pay.

At the very least, the license forbids use within a commercial program. It prohibits distribution in return for money, though you can give it away. Even if you give it away, you usually have to provide the complete package, including advertising, documentation, and any related programs that were bundled with it. It is important to remember that you can be sued or even prosecuted for violating the license on free software or true Freeware. (And it wouldn't hurt to put it high on your candidates' list for files to be scanned for viruses prior to turning them loose on your computers or network!)

The next step up the software food chain is Shareware. This is also a very specific term and its rules and even its use are governed by the Association of Shareware Professionals. (ASP - Isn't that what killed Cleopatra?) Shareware is a try-it-before-you-buy-it system. You get a limited amount of time to try the software and then must either delete the program from your disk or fill out the included form and pay the appropriate price for your usage (for example, single-user noncommercial, single-user commercial, multi-user, site license, corporate-wide site license, and so forth).

As you might expect, Shareware has restrictions on use and distribution similar to those of Freeware - no sales, except by authorized agents and no reuse within other products without a pre-signed agreement. Shareware holds my personal observation record for most abused software license category. For instance, it seems that everybody I encounter, seems to have a copy of PKZip, but I seldom meet people who have actually paid for it.

Shareware vendors go after license violators much more frequently than free software or Freeware vendors. Their association allows them access to much more legal resources, and provides additional validity to them when they talk to prosecutors. As ASP grows and becomes more organized, I expect them to increase their license enforcement efforts.

Next up are the fully commercial categories, starting with simple single user licensing. If you want their software, you buy enough copies for everyone who uses it. This category and the others that follow can be tricky in their wording regarding what constitutes a user. For some, a user is an actual person in the act of using the software. For others it is a person with the potential to use the software because it is loaded on a particular computer. The former allows you to load the software on multiple workstations as long as it is not simultaneously in use while the latter requires you to physically remove

the software from one workstation before using it on another. The two methods are often referred to as per-user versus per seat licensing, though I have seen these terms used differently.

Some software automatically allocates itself for the per-user licenses, but many others require some type of license management utility to control access. In fact, the lack of license management software for companies using honor-system (as opposed to self governing) per-user licensing can be enough to start a piracy investigation. Many companies, including Microsoft, are very suspicious of customers who have more workstations than licensed users but don't control them with utility software. I've even seen cases where the presentation of evidence of regular and thorough use of license management software was successfully used to ward off impending piracy investigations.

While there is usually a quantity discount of some type offered for these licenses, there is no special pricing for networks. You have to get network-specific versions of software to get network pricing. Both per-seat and per-user variations also show up in this category.

A common pattern is to buy one copy of the software that covers the server and the first user (or workstation), and then buy enough network node licenses to cover the rest of the users (or workstations). That initial copy might be the same as the single-user software or it might be a special network master software. Falling behind on the user (or workstation) license count in a growing company is both very common and very illegal.

Quantity-based "packs" (as in 5-pack, 10-pack, etc.) are also quite common for network-based schemes, as are site licenses. Sometimes the "packs" are more restrictive in that they make you pay separately for the server and the client side of client-server software. A 25-user license for the server might include the server software and an allowance for 25 workstations, but if you need only 25 concurrent users, but have 50 workstations, you would have to purchase 25 extra workstation licenses. (One of my favorite specialty programs uses this method, and it makes it a very hard sell to cost-conscious potential buyers.)

A more generous type of network license is the per-server variety. This usually is very cost effective and has the added benefit of being easy to budget for future planning. You just count your servers and ante up.

In addition to these common variations on licensing, there are both precalculated and custom site licenses as well as corporate versions of site licenses that cover multiple locations. There is a lot of activity here, since, by using the latest license management software, you can allocate per-user licenses over a Wide Area Network (WAN) to save money. Software companies, particularly Microsoft, have been changing their licensing structure specifically to restrict this trick, much to the displeasure of corporate America (like Bill Gates really cares about offending customers - or for that matter, even the government).

Next month, we'll play the software piracy version of Truth or Consequences, with an emphasis on the consequences. Meanwhile, I'm trying to find an AA chapter that will accept my parrot. That rum seemed to bring out the worst in him.

HOW TO PIRATE SOFTWARE - Part Five

by Wayne M. Krakau - *Chicago Computer Guide*, July 1998

Wayne is a partner in Krakau Business Computer Systems, Inc., a systems integration firm that is a Novell Gold Authorized and Novell IBM Host Master Authorized Reseller. He has been working with computers for 21 years. He holds an M.B.A. in Marketing and a B.S. in Information Science and was originally educated in Computer and Information Systems Analysis Engineering. He holds CNI, MCNE and CNTE ratings and all five of the original LANDA Certifications. Wayne can be reached for questions, comments, or topic suggestions at (847)298-7695 or via E-mail to wayne@mcs.com.

Finally, the seemingly never-ending saga is coming to a close. This is the last of my series on the ins and outs of software piracy. Again, I must warn you that this is a report from the trenches, not a legal guide. This, as they say on all of those legal discussion shows, is the penalty phase of the trial.

Listing many different ways to pirate software begs the question, "What happens if I get caught?" The short answer is, to paraphrase the common expression, "Don't let the cell door hit you in the butt on the way in." Yes, piracy at a high enough level (two grand at retail) is a federal felony, carrying with it all of the usual benefits of a felony conviction. For this offense that includes up to two hundred thousand dollars in fines and a couple of years' free room and board in a federal "hotel" - and that's per occurrence! I'm sure it also does wonders for your employability and credit rating.

And, just to cheer you up, on top of the criminal implications, there are also some unpleasant civil law possibilities - try company and individual lawsuits usually starting out at a minimum (based on my observations) of fifty thousand dollars!

The term "individual" here means you. There is a big push on to pursue the people in the chain of command who actively pirated software, gave the order to someone else to pirate software, or even passively allowed piracy to happen when it was within their sphere of responsibility to prevent it. This means that even if you hate your employer and have a good enough resume to get another job if your employer is bankrupted by assorted fines and settlements, you still need to be aware of piracy issues just to cover your own butt. Even something as seemingly harmless as neglecting to prevent employees from copying software onto their home computers can make you eligible for lawsuits and possibly even prosecution.

The first victim of the legal attack is the system administrator. Please note that possessing the actual title "system administrator" is not required to put you in jeopardy. Whatever your real title, if you're the one who has been shanghaied into monitoring the network, you're elected as the system administrator and, therefore, have a target painted on your back.

The important thing to remember is that for the subordinate, the Nuremburg Defense ("I was just following orders.") simply doesn't work. (If you don't get the historical reference, then get thee to a library! Try looking up a little event known as WWII.) (In fact if you don't get the literary reference within the previous parenthesis, then look up "nunnery" in Shakespeare's *Hamlet* - and no, that's not Babe's little brother.) (Ha! And you thought that you'd never need all that "useless" information that you pretended to learn in your high school English and History classes.)

For the superior, in turn, the Sergeant Schultz Defense ("I know nothing.") also doesn't cut it. (If you don't get the television reference, then get cable and look for *Hogan's Heroes* reruns. If you do get this one, then you are probably from my generation and wasted your formative years watching too much television - just like I did.) If management of the computer system is anywhere below you within the chain of command, you might also have a target on your back. Barring a successful implementation of "plausible deniability" - a term invented to describe the Eisenhower administration's policy toward its lying about the U-2 spyplane flights over the Soviet Union. Used, that is, until they got caught in the lie

due to the downing of one of the planes. Oops! (All right, I'll admit it. This reference didn't come from high school. It came from watching *PBS* and *The Discovery Channel*. I guess I'm still a child of television.) Essentially, it was a precursor of the Sergeant Schultz Defense.

If you think that you don't need to keep exacting records, create and enforce company anti-piracy policies, or purchase software tracking products, then you could have a problem. Following these procedures may be the only way to preemptively stop a minor inquiry from becoming a full-blown investigation. They can be used to counter an unfounded accusation or to unravel a misunderstanding without a lengthy and expensive court battle. They are cheap insurance against spending a lot of time and money defending yourself.

With the apparent risks on the table, what are the odds that you will get caught? Well, for example, think about that non-English-speaking cleaning lady that you yelled at the other day. She tells her daughter about her terrible employer. The daughter just happens to be a computer whiz in her high school. She carefully instructs her mother on what to look for and what "garbage" to save. Then they call one of the 1-800-FINK phone lines. (No, that's not a real number.)

The next thing you know, the daughter has a college fund - we're talking ivy league here, not community college. Meanwhile, you are hip deep in Software Gestapo backed by federal agents, are missing your file servers and half of your workstations, your employees' home computers are being scanned, your company is out big bucks, and you are facing some serious personal legal fees. Yes, you guessed it. There are reward programs for people who turn in software pirates. As you may have also surmised, disgruntled employees are frequently sources for leads. I'm sure that dissatisfied customers, competitors and even estranged spouses are also popular sources.

The one last bit of advice I can give to anyone who has pirated software, either accidently or on purpose is to KEEP YOUR @#$%^&* MOUTH SHUT! I've always idly wondered why people brag about some petty theft like stealing towels and knickknacks from hotels, which to me seems like wearing a sign that says "I don't have personal ethics, so don't deal with me, since I might not have business ethics, either." I am totally at a loss as to why anyone would cop to or even brag about software piracy, potentially a felony!

For example, I recently attended at a non-computer-related social event at which I overheard a specialty software vendor brag (at high volume in a public place) about being equipped with all of the tools necessary to crack any CD-ROM-based copy protection scheme. He described how he had used that technology to illegally copy dozens of different software CDs onto new writeable CDs for both his personal and his employer's use. Finally, he offered to sell this technology, at cost, to any of the new friends that he had made at this function. He even offered to direct them to the appropriate piracy-oriented Web sites so they could adapt the system to any new protection schemes that might thwart them in the future.

After finding out that I am in the computer business, people often go out of their way to tell me about the various ways in which they have ignored or at least stretched the limits of licensing rules. They even tell me details of sometimes blatant piracy. Please - I don't want to hear about it, and you shouldn't want anybody to know about it. I would suggest that you don't even talk about weaknesses in record-keeping at a company where there is absolutely no piracy, unless it is part of a discussion on how to improve things. Loose lips, even in this law-abiding situation, could result in an incredibly inconvenient and possibly expensive investigation.

Meanwhile, my parrot (now clean and sober), inspired by the recent release of a movie starring a parrot, has winged his way (in his case, literally) west to Hollywood in search of stardom. I didn't have the heart to tell him that the parrot movie seems to be tanking, reducing the demand for intelligent parrots in future flicks.

OS UPDATE - Part 1

by Wayne M. Krakau - *Chicago Computer Guide*, August 1998

Wayne is a partner in Krakau Business Computer Systems, Inc., a systems integration firm that is a Novell Gold Authorized and Novell IBM Host Master Authorized Reseller. He has been working with computers for 21 years. He holds an M.B.A. in Marketing and a B.S. in Information Science and was originally educated in Computer and Information Systems Analysis Engineering. He holds CNI, MCNE and CNTE ratings and all five of the original LANDA Certifications. Wayne can be reached for questions, comments, or topic suggestions at (847)298-7695 or via E-mail to wayne@mcs.com.

This article contains an update on the latest operating system issues. The most obvious operating system happening is the recent release of Windows 98 by Microsoft. It will be loaded on most new computers and, of course, is available as an upgrade to either Windows 95 or Windows 3.x.

I had great hopes for Win 98 as the salvation for my clients who are drowning in Win 95 bugs. Microsoft even reported to the national press that Win 98 would fix more than 5,000 outstanding bugs in Win 95. (I'm sure the marketing braniac who thought of this sales tactic - "Buy our new product because our old one stinks" - is probably now assigned to an assembly line putting Win 98 CDs into boxes.) There was no mention of either the ethics of having released an operating system with 5,000 bugs in the first place or what the original bug count was when Win 95 was released. Since Win 95 is already in its second major release (and God and Bill Gates only know how many unmarked slipstream releases), and, there are bug-fixing Service Packs available for it, the original bug count could logically be assumed to have been way more than the currently admitted 5,000.

My hopes, alas, are not to be realized. Front page articles in the national press are already reporting difficulties with Win 98. Many computer manufacturers are issuing warning of disastrous consequences for those using Win 98 on various models of their computers. Some corporations are forbidding Win 98 upgrades and are even erasing the hard disks of new computers equipped with pre-installed Win 98 and installing Win 95 in its place.

As you might guess with the release of a new desktop operating system, the more proprietary the computer, the more likely it is to run into hardware interfacing bugs. Since notebook computers are almost 100% proprietary, they are the ones getting hit with the most problems. In some cases the problems are so severe that you simply can't run Win 98. In most cases, manufacturers are scrambling to create and distribute patches.

For desktop computers, hardware interface related bugs are appearing in direct proportion to the degree of uniqueness in the hardware. This, by the way, is one of the reasons that I recommend that any purchase of new computer equipment include some type of "percentage of proprietariness" rating as part of the evaluation, with major deductions in the overall rating given for using potentially incompatible, proprietary subsystems when they aren't really necessary.

This is especially true given that the performance figures that I've seen on proprietary systems such as RAID arrays, disk controllers, and network cards indicate that commonly available aftermarket products (that is not purchased from the computer manufacturer) usually outperform the manufacturers' products, often at a lower cost. I will hedge my bets on this piece of advice (translation: weasel out) by stating that, obviously, these types of performance comparisons aren't the only things to think about when planning a purchase, and, even then, should be evaluated on a case-by-case basis.

In addition to the hardware interface bugs in Win 98, there are many other non-hardware-related bugs being reported, including bugs in new machines and as well as upgraded ones. Disabling other vendors' DLL files is one outrageous reported bug that is really a reprehensible design decision, not actually a bug. It harks back to the old days when the alleged motto for DOS developers at Microsoft was "It's not done until it blows up Lotus 123." The spontaneous lockups caused by Internet Explorer, however, are the bugs that shocked me the most.

My pre-release impression of Win 98 was that it was mostly a giant bug-fix for Win 95 that, ethically speaking, should have been offered as a free download, not as a paid upgrade. The only feature that seemed to be much of an advance over the latest release of Win95 was the finalization of the integration of Internet Explorer with the operating system, a process which started in later versions of Win 95. This is the main issue currently under dispute in legal action by the Justice Department against Microsoft. It is ironic that this advancement, the one feature most publicized by Microsoft in general and Bill Gates, in particular, is one of the major causes of bugs for Win 98. You'd think that they would at least make sure the most bragged about feature really worked

Now for the really bad news - Microsoft has officially declared that no Service Packs are currently being planned for Win 98! I know Microsoft has earned a bad reputation for being unresponsive in the past, but this is getting really blatant. I guess its attitude is now "You will not only pay for this bug fix, but you'll have to put up with new bugs, too!"

That sound you hear in the background sounds suspiciously like the plaintive bleating of sheep being led to slaughter. Tune in next month to see if Babe (the pig not the baseball player) can be recruited by the Justice Department as a prosecutor to help protect us sheep. Baaaaaa!

OS UPDATE - Part 2

by Wayne M. Krakau - *Chicago Computer Guide*, September 1998

Wayne is a partner in Krakau Business Computer Systems, Inc., a systems integration firm that is a Novell Gold Authorized and Novell IBM Host Master Authorized Reseller. He has been working with computers for 21 years. He holds an M.B.A. in Marketing and a B.S. in Information Science and was originally educated in Computer and Information Systems Analysis Engineering. He holds CNI, MCNE and CNTE ratings and all five of the original LANDA Certifications. Wayne can be reached for questions, comments, or topic suggestions at (847)298-7695 or via E-mail to wayne@mcs.com.

Here's the second in my series on the latest operating systems issues. This article covers the upcoming release of Windows NT Version 5.

As with Windows 98, I have great hopes of NT 5 fixing a lot of problems and omissions in its predecessor, though the problems I reported last month in Win 98 have me running scared.

In a manner similar to what was done with Win 98, some braniac put a notice in one of the technical areas of Microsoft's Web site that NT 5 would fix more than 10,000 outstanding bugs in NT 4! (Yikes! Is there something in the water supply at the Microsoft corporate campus that drives people to commit ritual employment suicide?) Microsoft management immediately had this notice modified to have a more marketing-friendly phrasing - well, almost immediately - not until the national press found and published it. Oops! Again, no mention was made of the ethics of having released a product with that many bugs, or of the original bug count, prior to the release of the various patches and fixes that were subsequently released.

In this case, it was a technical person, not a marketing or sales person who released the embarrassing data, so he or she might not be demoted to box-stuffing duties, as a technical person would not reasonably be expected to be educated on the benefits of the Microsoft Creative Truth (in English: lying) policy. However, I suspect that the techie who did it might never again be allowed to even browse the Web, much less alter a portion of Microsoft's own site.

Another thing that worries me is that the commercial-release version of Windows NT Version 5 is a ghost, an ethereal spirit, a figment of Bill Gates imagination - in other words, vaporware. This is not a piece of software that is within a few weeks of release, with all of the features pretty much locked in, and only the last minute removal of test code and the addition of the last few bug fixes pending. This is software with an expected (as opposed to officially announced) release over a year from now (as I write this in August of 1998).

Microsoft has incorrectly estimated the release date so many times that these inaccurate announcements now seem less like overconfidence or even incompetence and more like a specific marketing policy of what was termed in IBM's glory days as FUD - Fear, Uncertainty, and Doubt. A major portion of this policy is to keep people from buying your competition's product by dangling the imminent release of your allegedly killer product in front of them, even though it's nowhere near ready for release. The idea is to either get them to buy your current not-so-great product in the hopes of an easy or inexpensive transition to the new one, or, to make them so worried about making a decision that they hold off buying altogether. If imitation is the sincerest form of flattery, then IBM must be very proud. (Of course, you need to remember that this aspect of IBM's FUD policy was a factor in getting antitrust investigators to go after IBM years ago!)

Any product as far away from its release date as NT 5 is, could easily change by the time it is released. A particular feature that sounds appealing now might end up not working as advertised, be buggy to the point of being unusable, or simply be omitted from the final product at the whim of the developer. The last-second addition of other, as yet unnamed features, (especially in a company driven by marketing needs as opposed to technology or customer service) might also destabilize or just slow down the base product. You'll have to consult your favorite psychic for the true facts.

From what I've seen, Microsoft is not going to change one important characteristic of NT that has always been one of my pet peeves. That is the use of a desktop operating system as a server. I still lean strongly in favor of true honest-to-God network operating systems that were designed and optimized as such as opposed to a jack of all trades (and potentially master of none) operating system that must be both a desktop and network operating system. (Or as a fake commercial on *Saturday Night Live* once put it, "It's a dessert topping AND a floor wax, too!") I would be much happier to see Microsoft split NT into two separate products with two separate development teams, each with its own unique goals. A while back, Microsoft claimed that they were different, but some serious techies (with more expertise in such matters than I have) exposed that lie by examining both products with specialized software tools and finding that their kernels (the base underlying programs) were identical. Oops!

Microsoft is finally addressing the performance deficiencies of NT, but I am worried that there are limits to what can be done while still keeping this dual mode network/desktop policy. With NT 4, all of the information I have been able to gather from software companies, very large resellers (who have dealt with both NT and NetWare for some time), and system administrators indicates that the performance ratio between NetWare and NT 4 for file servers is 4 to 1, and for application servers (database, mail, etc.) is 6 to1. That is, if you have a maxed-out NetWare 4.11 file server and replace NetWare with NT 4, you will need to buy three more equivalent servers to handle the load (hence the 4 to 1 ratio). If you have a maxed-out NetWare 4.11 application server and replace NetWare with NT 4, you will need five additional servers (for a 6 to 1 ratio). I have seen similar performance ratios in the national press though with some variation in the actual numbers (3 to 1 for file servers in one article, for example). Obviously, anything Microsoft can do to reduce these ratios will help.

Interestingly, these ratios are part of the motivation for many NetWare resellers (as evidenced both by my own personal interviews with them as well as many nationally published reports) to jump on the NT Server bandwagon in spite of believing that NetWare is a far superior solution. They like the profit in selling more servers to handle a given need. They also like the additional profit on the labor side of the equation, too. They tell me that the additional time needed for installation and integration leads to some additional billable hours, but that the additional ongoing maintenance and management costs of NT versus NetWare are where the real profits are.

Next month, I will continue with my NT 5 coverage. In the meantime, please note that while I am a rabid NetWare fan (feel free to take my opinion with many large grains of salt), I know full well that market demands will require me to actively sell NT Server Version 5 (and maybe even NT 4) even though I only currently sell NT Workstation Version 4. I have multiple employees undergoing certification training in NT 4 now, and have the materials (though, at the moment, not the time) to get certified myself. While I have strong technical and ethical opinions, my MBA specialty was in Marketing with a sub-specialty in Marketing Research. Just as I am authorized to sell multiple brands of computers in order to satisfy my clients' demands, even though I have an overwhelming favorite brand, I will need to satisfy my clients' needs for alternative network operating systems. Besides, every NT 4 server that I've encountered so far has been configured just as poorly as almost all of the NetWare servers that I see. Opportunity knocks!

OS UPDATE - Part 3

by Wayne M. Krakau - *Chicago Computer Guide*, October 1998

Wayne is a partner in Krakau Business Computer Systems, Inc., a systems integration firm that is a Novell Gold Authorized and Novell IBM Host Master Authorized Reseller. He has been working with computers for 21 years. He holds an M.B.A. in Marketing and a B.S. in Information Science and was originally educated in Computer and Information Systems Analysis Engineering. He holds CNI, MCNE and CNTE ratings and all five of the original LANDA Certifications. Wayne can be reached for questions, comments, or topic suggestions at (847)298-7695 or via E-mail to wayne@mcs.com.

Here's the third in my series on the latest operating systems issues. This article completes my coverage of the upcoming release of Windows NT Version 5.

To continue with my ramblings on NT 5, another hurdle that it must overcome is the potential conflict with its customer's Year-2000 compatibility efforts. With all of the expected delays, NT 5 is expected (as opposed to announced) to be released in mid-to-late 1999, just as most companies are going on full Y2K alert. Programming and support personnel are scheduled to be working overtime, trying to keep up with any last-minute Y2K-related errors.

Remember, Y2K errors can pop up at any time a calculation looks ahead beyond January 1, 2000, not just after that actual date. As we approach that magic date, applications that calculate receivables, payables, loans, leases, individual or project schedules, court dates, or just about any other date-related fact are eligible for failure, or worse yet, subtly inaccurate data. With all of this to worry about, it's no wonder that companies are less than anxious to make a major upgrade to their network and desktop operating systems right in the middle of this mess.

Now, word is out that Microsoft's own closest corporate partners are rebelling at the idea of doing simultaneous Y2K and NT 5 implementations. These are the major corporations that would normally have agreed to unconditionally implement NT 5 immediately upon release after having been top-level beta testers, all in return for major price discounts and the promise of greatly enhanced support services from Microsoft.

Microsoft can't guarantee a specific release date, but expects these companies to agree to implement NT 5 en masse on whatever day Microsoft chooses. This issue has yet to be sorted out, though I expect that many of these corporate partners may jump ship, resulting in major embarrassment for Microsoft's marketing department, and a major headache (in form of less feedback from the trenches) for Microsoft's techies.

My biggest fear here, is that Microsoft will be tempted to do with Windows NT 5 what it did with Windows 98. That is, it arbitrarily picked a date and decided that for marketing purposes (or as they say in the Far East, to save face) they would simply release their product, ready or not. With Windows 98, one week there were published reports of a release date six months in the future, with lots of bug-fixing work left, and the next week the software was released! Windows 98's reliability has suffered greatly from its rush to shipment, and I'm sure that NT, being an inherently more sophisticated product, would have even more problems as the result of such a premature release.

Another issue is, as NBC so eloquently puts it, is the "It's new to you" syndrome. A particularly difficult to implement technology such as a network directory services system, takes much experience to get right, and if you haven't done one before, as in Microsoft's case, then it's new to you. That means you're a beginner, likely to make all of the mistakes that anyone would reasonably expect from a beginner.

When Novell was a beginner in directory services, they came out with NetWare 4.0 (along with its descendants, 4.01 and 4.02) which was elegant in conception but a real disaster in terms of actual implementation. They aggravated things by releasing it before it was really ready, just as I am afraid that Microsoft will do with NT 5. NetWare 4.0 was so bad, in fact, that I declined to offer it to my clients, and

actually gained a few new clients when I was called upon to uninstall it and put back a previous version of NetWare for some corporations.

Novell didn't get things right until they released version 4.1. That product followed up on the promises inherent in the design of 4.0. As is usual for Novell marketing campaigns, no serious effort was made to explain to the public that 4.0 was a mess and that 4.1 was the "really working" version and was radically improved as compared to 4.0. That's why, to this day, I still run into applications developers and sometimes even civilians (non-computer folks) who are afraid to move from NetWare 3.x to 4.x for compatibility reasons. They fear that their application programs won't work with anything higher then NetWare 3.x.

I have no reason to believe that Microsoft, which already has a shaky reputation as far the reliability of major software releases is concerned, should magically be able to create a stable directory services system on their first try. Of course, there is always the hope that they have learned from Novell's mistakes, but so far, they haven't even learned anything from their own, so don't hold your breath.

Whatever happens with NT 5, I am resigned to the fact that, due to marketing-driven customer demand, I will probably have to sell and support it. As long as I don't get blamed for its problems, my company will survive and probably thrive on the increased number of billable hours.

Next month I'll cover NetWare 5. Notice that I didn't say "the upcoming release of" NetWare 5? That's because I just received E-mail that notified me that NetWare 5 has been released more than one week early! Wow! How's that for innovation? What will those people at Novell think of next? Maybe they'll even start marketing their products properly - or is that too much to hope for?

OS UPDATE - Part 4

by Wayne M. Krakau - *Chicago Computer Guide*, November 1998

Wayne is a partner in Krakau Business Computer Systems, Inc., a systems integration firm that is a Novell Gold Authorized and Novell IBM Host Master Authorized Reseller. He has been working with computers for 21 years. He holds an M.B.A. in Marketing and a B.S. in Information Science and was originally educated in Computer and Information Systems Analysis Engineering. He holds CNI, MCNE and CNTE ratings and all five of the original LANDA Certifications. Wayne can be reached for questions, comments, or topic suggestions at (847)298-7695 or via E-mail to wayne@mcs.com.

This is the fourth installment of my series of articles on operating system updates, with NetWare 5 being the subject at hand. It's a bit of a good news/bad news story, with the emphasis on the good news. (The bad news isn't all that bad, I'm just a nitpicker as far as NetWare is involved. Remember, "You always hurt the one you love.")

The newly-released version of Novell's network operating system product is called NetWare 5. It is the immediate descendent of IntraNetWare 4.11. The first part of the story of this product is in the name. In what's been called in advertising circles the worst marketing disaster since New Coke, Novell, during the introduction of Version 4 of their main product, renamed it "IntraNetWare", a step that single-handedly blew off a rather large percentage of its customer base, losing untold numbers of potential repeat customers. I've seen independent studies indicating that many people believed that "IntraNetWare" was meant only for companies having a leased line to the Internet and possibly even an in-house Web server, too. Oops! ("Marketing ability? Marketing ability! We don't have to show you no stinking marketing ability!" - Novell's original company motto.) Finally, Novell has seen the light. They went back to good old "NetWare" for Version 5. (New company motto - "Okay, we'll show you SOME marketing skills, but don't push it!")

The next bit of news about NetWare 5 is the best. It actually works! (Quite a good trick since it was released ahead of schedule!) While there are obviously some bugs in any new product, first reports indicate that NetWare 5 is stable and reliable. Whew! That's a relief. I will admit that I was worried about having another mess like the initial release NetWare 4.0. In this case, however, Novell had an advantage in that NetWare 5, in spite of all the new features, is still conceptually a derivative of IntraNetWare 4.11, whereas when comparing NetWare 4.0 to its immediate predecessor, NetWare 3.11, it was more like, as the *Monty Python* gang put it, "And now for something completely different!"

Also, Novell, unlike a certain multibillion dollar operating system publisher, which shall remain nameless, is actively responding to bug reports and suggestions by publishing patches, fixes, and various hint files. At the Novell seminars that I have attended, technical questions were usually met with a response that included a verbal summary along with a reference to an already-published document or patch file that provided a solution or work-around. The answers to questions, posed during the seminars, that didn't have previously published answers have been showing up on the Web within a week or two after the seminars.

The new feature that has gained Novell the most publicity is protocol independence. NetWare is no longer tied to the old IPX/SPX system of communications, though it still can use it to support older workstations and other network devices. The default communications protocol is true, standards-following TCP/IP (as opposed to another operating system vendor, based in Washington State, whose server and workstation operating systems violate multiple standards when using TCP/IP).

Old NetWare could use TCP/IP only by encapsulating IPX/SPX packets inside of TCP/IP packets - effective, but grossly inefficient and somewhat limiting in terms of compatibility. It was like being forced to take your mail, already stuffed into a normal business envelope, addressed in English, and being forced to, to stuff that envelope into another, larger one, addressed in another language, in order to get the

mail delivered. At the other end, the recipient would have to open both envelopes to get to your message. What a pain! Now you just fire up NetWare with real TCP/IP.

The irony is that NetWare 5's implementation of TCP/IP even includes nonstandard work-around routines that are required to deal with communicating with other operating systems' standards violations (see above)! While the Novell personnel at these seminars did their best to follow Novell's new corporate policy of not being at war with any competitors, but instead trying to be cooperative and avoid criticism (what they used to call "coopetition"), the attendees, despite being scolded by the Novell folks, were not so restricted. Those who had worked extensively with the competition's offerings were quick to point these and other deficiencies, often quite sarcastically. I expected this reaction from the long time NetWare resellers in attendance, but I was somewhat surprised to hear this from the corporate administrators in the crowd. I had thought that they had all fallen under the competition's propaganda avalanche. Oh well, I guess free thought still survives!

In the spirit of free thought, the audience, as a whole, was equally vocal (actually MORE vocal) in expressing their displeasure whenever some installation and/or configuration routine was either not documented well enough, or, was inadequately automated. (I'd like to think that NetWare users and even resellers are not either some sort of passive zombies, dedicated to idolizing some demigod from Utah, or, alternately, Luddites, sticking with what they know out of blind fear of trying something new.)

For example, the handouts and overhead at one seminar showed a page with six steps required to do an across-the-wire upgrade from NetWare 4.1x to NetWare 5. Note that across-the-wire means that you have acquired a new server to replace the old one at the same time as you are upgrading NetWare. The information from the old server must go literally across the network wire to its replacement. After counting something like twenty-eight distinct (admittedly easy) steps needed for the speaker to do the upgrade, many attendees noticed that he was reading his own detailed list of steps! After a mass protest erupted over the inadequacy of the handout, he offered to put all of his "insider" information on floppy disks that would be sent to all of us. After further - and much louder - protest, he agreed that the proper place for such information was Novell's Web site, where anybody, not just those at this seminar, could get at it. He also agreed to pass on our suggestions on how to further automate the procedure. While most of these steps are in the manual, it's a lot easier to use a summarized list, especially if that list includes the latest insider tips.

There's more NetWare 5 coverage coming next month. Meanwhile, I am scheduled to burn a goat on my red altar and I have to find my prayer rug. Now, which way is facing Orem (Utah)?

OS UPDATE - Part 5

by Wayne M. Krakau - *Chicago Computer Guide*, December 1998

Wayne is a partner in Krakau Business Computer Systems, Inc., a systems integration firm that is a Novell Gold Authorized and Novell IBM Host Master Authorized Reseller. He has been working with computers for 22 years. He holds an M.B.A. in Marketing and a B.S. in Information Science and was originally educated in Computer and Information Systems Analysis Engineering. He holds CNI, MCNE and CNTE ratings and all five of the original LANDA Certifications. Wayne can be reached for questions, comments, or topic suggestions at (847)298-7695 or via E-mail to wayne@mcs.com.

This is the fifth column in my series on operating system updates. It continues my coverage of the newly released NetWare 5, the latest version of Novell's flagship product. As I mentioned last month, it's a good news/bad news story, with an emphasis on the good.

Much of the publicity surrounding NetWare 5 has been about its new Java Virtual Machine (JVM). Java is intended to be a near-universal language that can run on any operating system. The folks at Microsoft have attempted to dilute Java's popularity by coming out with their own incompatible version of Java, but, so far, have not succeeded. The only thing they've succeeded in doing is getting themselves sued by Sun, the originator of Java, for violating the terms of the agreement between them regarding Microsoft's use of Java. (Gee, Microsoft in court - what a novel concept!)

Microsoft has lost the first round in this battle, but the final results are not yet in. If the final decision goes against them, I fully expect the Microsofties to start a program modeled after the U.S. Government's farm program in which the government pays farmers to NOT plant certain crops. In this case Microsoft could afford to pay huge fees to Java programmers to NOT write Java programs. They would get these fees in addition to any money that they could earn by programming in "permitted" languages, with bonuses going to those who choose Microsoft's languages. Over time, Java would wither and die. That'll teach Sun (and any others) to avoid going up against Microsoft! How dare they! (This speculation assumes, of course, that the Justice Department is totally - as opposed to the current partially - asleep at the switch.)

Novell's Java Virtual Machine is a program that allows Java-based applications to run on a NetWare server. Novell also provides a version that can run on a properly equipped workstation, but, for now, that is really a side issue. The JVM provides an easier way for developers to create server applications, and Novell enhances their capabilities by throwing in a bunch of programming tools. (As for JavaBeans for NetWare, who comes up with these wacky names? Maybe it's the radiation from monitors.) JVM is Novell's bid for the application server market as well as its way to counter criticisms about the lack of graphical management tools on its server.

Right now, the only serious applications that run on Novell's JVM are the new installation programs and a management program called ConsoleOne. The installation program's benefits are obvious, given that many of the new features added to NetWare need to be configured before use. The old character-based INSTALL routine was already complicated enough in previous NetWare versions.

ConsoleOne contains a subset of the features that, up until now, have only been found in the workstation-based program NWADMIN. Some of these features are also present in the DOS-based NETADMIN, but nobody is clamoring for more network management via text-based, non-GUI (Graphical User Interface) programs. NETADMIN has been relegated to being a tool for only old die-hard NetWare geeks (like me!) when doing quick, experimental changes to things like login scripts.

ConsoleOne's big advantage is that it alleviates the need to have a workstation handy when some NWADMIN-type task must be done while you're hidden away in the server room. A JVM disadvantage, from corporate management's point of view, however, is that while you remain unobserved in the file server room, supposedly working on the network, you might really be playing Java-based games on the

server! (I suspect that the average time span between the invention of a new computer programming language and the design of games for that language is measured in seconds.)

ConsoleOne takes up tons of memory, and as with any GUI, can be processor intensive, so most people will turn it off when not in use, whereas a non-GUI tool like MONITOR can be left running continuously. ConsoleOne is also currently kind of slow, and it doesn't yet completely duplicate the functionality of NWADMIN, or even take the place of the old faithful MONITOR program, but it is a major step in the right direction as far as system administration is concerned. Luckily, the underlying JVM itself is very fast, and I'm sure that ConsoleOne can be improved in subsequent versions, so the initial shortcomings are probably only temporary. (Hey, if Microsoft can wax eloquent about the features of an as-yet nonexistent product, then I can speak about expected incremental improvements in an honest-to-God real, delivered product!)

Obviously, except for Novell's own management and installation tools, the usefulness of Novell's JVM depends on the impact of Java on the industry as a whole. If Java is a success (in spite of Bill Gates best efforts), then having the fastest Java engine in the industry would be major advantage.

Another addition to NetWare 5 is a 5-user version of Oracle8, the popular relational database. From what I've seen, Oracle8 is tied with one other, much less common, product for the "most" relational database engine on the market. That is, it follows more of Dr. Codd's rules than the vast majority of other relational and allegedly relational databases. Note that there are many database products out there that claim to be relational, but don't follow ANY of these rules! (Sorry to rave out, but this has been a major pet peeve of mine since a totally non-relational product won the award as the best relational database in a major national periodical a few years ago.) The addition of Oracle8 to NetWare is pretty much an "either/or" proposition for Novell customers. If you use Oracle-based programs, it's great. You just add additional user licenses as necessary. If you don't use Oracle-based programs, then it's a waste of a couple of tracks on an installation CD, and the source of a few more annoying questions on your next NetWare CNE or CNA test.

This saga will continue next month. Meanwhile, I will be trying to figure out why I didn't have enough marketing sense to plug my own business when given the opportunity to at the end a recent radio interview. My MBA concentration was in MARKETING! My marketing professor must be rolling over in his grave! That thumping sound you hear is my head banging into the wall.

OS UPDATE - Part 6

by Wayne M. Krakau - *Chicago Computer Guide*, January 1999

Wayne is a partner in Krakau Business Computer Systems, Inc., a systems integration firm that is a Novell Gold Authorized and Novell IBM Host Master Authorized Reseller. He has been working with computers for 22 years. He holds an M.B.A. in Marketing and a B.S. in Information Science and was originally educated in Computer and Information Systems Analysis Engineering. He holds CNI, MCNE and CNTE ratings and all five of the original LANDA Certifications. Wayne can be reached for questions, comments, or topic suggestions at (847)298-7695 or via E-mail to wkrakau@krakau-inc.com.

Okay! This is it. Here is the sixth and FINAL column in my series on operating system updates. (Hold the applause and try to keep the jeering to a minimum.) This article contains the remainder of my coverage of NetWare 5 and a totally unbiased summary of the various operating systems.

NetWare 5 introduces a new file format called Novell Storage Services (NSS). As with some of the other features that I've mentioned, NSS is a bit of a good news/bad news story. The good news is that NSS handles large volumes with incredible efficiency. Memory use for disk management and acceleration is down. The minimum amount of RAM required to mount a given volume had dropped to way below V4.11 levels. Obviously, having more total memory in a server allows for more caching and therefore more speed, but you might be able to get away with less total RAM, especially when dealing with volumes that are only occasionally mounted, such as optical disks that only contain archival information, or CDs containing seldom-used reference databases. Also, if you have a RAM chip failure, you may be able to remove the offending chip and limp along at a slower speed while waiting for the replacement chip.

The other advantage in NSS improved mounting speeds. I have heard reports form the field that 50GB NSS volumes mount faster than 1GB NW 4.x style volumes in the same server! For those running 50GB or larger volumes, this is a fantastic improvement. For most of those running smaller volumes, this improvement is handy, but not that important. The exception is for those running dual, redundant servers whose recovery cycle is dependent upon mounting speed. One dual-server NetWare 4.11 system that I've designed mounts a 22GB, 60% full volume in less than 90 seconds. While this is not exactly snail-like performance, users may get impatient during that delay and reboot their workstations, potentially losing their work. A 15-second NSS mount time (my guess based on the examples that I've seen) would keep the impatient types from either giving the three-fingered salute (CTRL-ALT-DELETE) or pushing the RESET button.

The bad news about NSS is that it doesn't support software mirroring (which isn't that common anymore), file compression or the Transaction Tracking System (TTS). If you depend on the extra space gained by file compression, you will be disappointed. Because of some problems with early versions of NetWare's file compression routines, it has gained a bad reputation, even though it has been very reliable in its recent incarnations. I suspect that this is the reason that I haven't been able to find any confirmation that Novell is working to add file compression to NSS. Either they aren't doing it, or their keeping fairly quiet about it. (My apologies if I simply missed some announcement about this).

The lack of transaction tracking, however, will be a temporary inconvenience, as Novell is working on adding that to NSS. Meanwhile, the SYS volume, which must have TTS, has to use the NW 4.x style format. Also, any volume that contains data explicitly protected by TTS should avoid NSS. With NSS, you will also lose the seldom mentioned, but quite handy implicit TTS protection that attempts to protect non-TTS files whenever possible. (Work with NetWare servers for a long time and you'll see that implicit protection in action.)

A handy improvement to memory management NetWare 5 is the more precise control of the dump memory function after an ABEND (abnormal end - Geekspeak for certain types of crashes). In NetWare 4.x and below, you needed enough room on your DOS partition to hold a copy of all memory

after an ABEND, with room for multiple copies being desirable. Since I am now frequently selling servers with 1GB of RAM, I had to configure huge DOS partitions. NetWare 5 dumps just the critical portion of RAM needed to debug the system, so I can now fit many dumped copies of memory onto a small DOS partition.

Oh, if you'd like to buy just about every disk drive ever made, NSS can handle up to eight "ZetaBytes" where one "ZetaByte" is defined as one million terabytes. (Note that this definition of ZetaBytes only appears one website other than Novell's. According to all other references, the international standard for prefixes is as follows: Kilo, Mega, Giga, Tera, Peta, Exa, Zetta, and Yotta, respectively, with no mention of Zeta with one "t". This leads me to believe that the correct prefix would be "Exa". Aha, one last Novell standards violation needing eradication, as well as proof that the Web truly is the great black hole of wasted time!)

And now on to the unbiased summary of operating systems that I promised you. Oops. That's actually a lie. You won't be getting an unbiased opinion. You'll have to settle for mine. If you've read my columns for very long, you have probably realized that I don't believe there is such a thing as an unbiased opinion on ANYTHING. There are only opinions from people who admit their prejudices and technological limitations, and those from people who don't. In my case, since I have worked on mainframes, minis, individual PCs, LANs, and WANs, including more terribly obscure desktop and network operating systems than even I want to remember, I hope you will feel that my years working on NetWare with DOS, Windows 3.x and Windows 9x workstations do not completely invalidate my opinions. Perhaps they could be considered at least interesting, if not necessarily accurate or universally espoused.

First, regarding workstations, I wholeheartedly agree with the Gartner Group and many other parties who believe that NT Workstation has too many hardware and software compatibility problems to be effective for most organizations. Almost all of my clients have software that won't run under NT or hardware that is either not supported by NT, or simply not powerful enough for it. Even the engineering clients who I expected to adopt NT Workstation in droves can't run it due to incompatibilities with existing software. If you are starting with a clean slate, and can afford its demands and limitations, then you should know that NT workstation will eventually be THE desktop operating system, so you might want to start with it now, but you don't have to.

As to Windows 98, if you buy a new machine already equipped with it, then use it. If you have computers equipped with OS Release 2 of Windows 95, then keep it, since you already have most of Windows 98's speed and interface advantages. Even if you have an earlier release of Windows 95, I'd suggest that you wait until at least the first Service Pack for Windows 98 is released before you upgrade. (Microsoft finally relented and is creating a Service Pack.) Windows 98 is currently far too buggy to try without a good reason.

How about the network operating system? You can probably infer from the previous articles in this series that I still prefer NetWare. Besides the vaporware issue (NetWare 5 is here, NT 5 isn't), from what I've seen of NT 5 Server, most of its key "advances" are merely implementations of features that NetWare has had for some time and has and perfected over the years. For some of these features, with directory services being the most important, the underlying design isn't as good as Novell's current product, as evidenced by suggestions in the national press that even if you standardize completely on NT Server, you should use Novell's NDS to manage it. Also, Novell is certainly not going to sit still and let Microsoft catch up.

As to market issues, after NT 5 Server is relatively stable (anywhere from six months to two years after its release) I will probably offer it as an alternative to NetWare for those clients who demand it, as long as they understand Microsoft's sad history with handling bugs. (See Part 2 of this series for more details.)

What about the name change from Windows NT 5 to Windows 2000? Am I the only one who sees this as an implication that Microsoft is embarrassed by NT 4 and wants to dissociate NT 5 from it? Also, what were they thinking about when they came up with that name? These days, the only

computer-related thing that civilians (normal folks as opposed to computer geeks) know about the phrase “2000” is the Millennium Bug! Why associate your product with a disaster? What’s next, a Titanic II cruise of the North Atlantic?

WHY 2K? - Part 1

by Wayne M. Krakau - *Chicago Computer Guide*, February 1999

Wayne is a partner in Krakau Business Computer Systems, Inc., a systems integration firm that is a Novell Gold Authorized and Novell IBM Host Master Authorized Reseller. He has been working with computers for 22 years. He holds an M.B.A. in Marketing and a B.S. in Information Science and was originally educated in Computer and Information Systems Analysis Engineering. He holds CNI, MCNE and CNTE ratings and all five of the original LANDA Certifications. Wayne can be reached for questions, comments, or topic suggestions at (847)298-7695 or via E-mail to wkrakau@krakau-inc.com. Past columns and company information are available at http://www.krakau-inc.com.

Why? That's the universal question that normal people (the non-computer-geek population) ask me about the infamous Year 2000 Bug (misnamed the Millennium Bug). How could supposedly intelligent people - viewed with awe by those mystified by computers - be so incredibly short-sighted that they could neglect to allow for something as obvious as the upcoming turn of the century? That's the main question I'm addressing, with some Y2K hints thrown in for good measure.

Let's start out with a little computer horror story. In the late 70's, while employed as a mainframe programmer, one of my first assignments was to act as part of an ad hoc team assembled to fix the 1980 Decade Bug! Yes, many of the programs in that corporation could not even handle a decade change, much less a century change. While articles today are whining about the difficulty of dealing two-digit years, we were facing one-digit years!

Now for the really scary part - our specific assignment was to apply a decade-specific emergency patch that would allow the program to be run ONLY during the1980s! Let's see now. When faced with the shortsightedness and possible downright incompetence of their predecessors, information systems management chose to - you guessed it - add yet another layer of incompetence! And no, this wasn't just any plain vanilla company. This company was a combination computer service bureau/computer consulting firm/computer facilities management firm. Note that the operative word here is "computer". All of the employees, from the president on down, were computer literate, and the vast majority of them were computer professionals. In other words, there were no excuses that could have been made regarding non-computer managers interfering with computer-related tasks.

Well, orders are orders, but ethics are also ethics, so, not being too fond of the Nuremberg Defense ("I was just following orders."), I did some research on date handling (the fiscal kind, not the relationship kind). After examining some existing programs, asking more experienced colleagues, and spending some time (my own) at the 70's equivalent to the World Wide Web - the local library - I came up with an algorithm (geekspeak for a formula) to handle all dates up to the year 10,000 correctly, at which point, it would run out of digits. It was an easy task, once I had the appropriate information.

With the algorithm in hand, I proceeded to examine the group of programs assigned to me. What I found was worse than expected. The programs written in the 60s already had a decade-specific patch to handle 1970! Even if I'd wanted to, I couldn't have applied the ordered patch - there already was one there. Just to make things more interesting, I also found a few post-1976 programs that couldn't handle a Leap Year! I fixed all of them, even the ones that would have worked through 1999.

What did I get for my extra effort? I got yelled at. No, to be more accurate I got YELLED at. The Director of Information Systems took me to his office and absolutely blasted me for not following orders, for not being a team player, and, most of all for wasting the company's valuable time and money in taking the extra time to implement and test the more complicated, thorough fix even though the programs involved would "obviously" not be used for more than another couple of years before being replaced

Hmm. I'll bet that's what they said ten years prior when they installed the "temporary" 1970 patch that I found in some of those programs. Oh, and in case you are wondering, I ran into a manager

from this former employer in 1996 and he confirmed that they were still using lots of programs that were "obviously" going to be replaced in the early 80s.

Now, it's time for a little quiz. What was the common factor, the common word, the theme used in describing changes to programs in the passage above? The magic word is "patch" with some points given to those who selected the related word "temporary". These words describe the underlying theme of the computing industry - the idea that much of the day-to-day work is just a temporary, "quick and dirty" patch.

Complete programs and even suites of programs are written based on the idea that they are temporary stopgap measures that will, when the immediate crisis has passed, be quickly replaced with more carefully planned and tested programs. Ah, but when one crisis passes, the next materializes, as if predestined, to take its place, so the replacement or at least reworking of the original quick and dirty programs is put off again and again, until the project is finally forgotten.

Much has been made of the original reason for shortening the date, space considerations, but by the 70s, mainframe languages could compress the two extra year digits into one byte of storage. Mainframes were also well on their way in the transition from cards to tape and disks as primary storage, so that extra byte (or even two) wasn't nearly as important as it originally was. There's more to the problem than just space for extra digits.

Next month I'll continue my analysis of the Y2K issue. For now, I'll get back to digging my bunker and stocking up on supplies.

WHY 2K? - Part 2

by Wayne M. Krakau - *Chicago Computer Guide*, March 1999

Wayne is a partner in Krakau Business Computer Systems, Inc., a systems integration firm that is a Novell Gold Authorized and Novell IBM Host Master Authorized Reseller. He has been working with computers for 22 years. He holds an M.B.A. in Marketing and a B.S. in Information Science and was originally educated in Computer and Information Systems Analysis Engineering. He holds CNI, MCNE and CNTE ratings and all five of the original LANDA Certifications. Wayne can be reached for questions, comments, or topic suggestions at (847)298-7695 or via E-mail to wkrakau@krakau-inc.com. Past columns and company information are available at http://www.krakau-inc.com.

Well, we're back to that critical question of why in the world did the computer industry get itself into this Year 2000 mess. Last month's example showed that only the young and naive are foolish enough to try to buck the "quick and dirty" doctrine within the computer industry if management is really determined to embrace that shortsighted attitude.

In the PC world, this doctrine is exemplified by that old favorite program (sarcasm intended), MS-DOS. Never in the history of computing has there been a program that has invited more criticism. Complaints over the years have included such subjects as memory management, physical memory limitations, and the lack of management utilities, but most are directly related to weaknesses in the original design. Even today, a lot of the problems with Windows 95 and Windows 98 can be traced back to the fact that they are still running with pieces of MS-DOS within them.

If you buy into Microsoft's and especially Bill Gate's claim that Microsoft is a bastion of innovation, (yeah, right) you might think that MS-DOS deficiencies were some project management error in the early days of the company by Mr. Gate's himself. Well, while that innovation story plays well in the press as a potential defense against government intervention in Microsoft's affairs due to alleged illegal monopolistic practices, it doesn't match the facts.

MS-DOS is a prime example. Microsoft purchased it to fill IBM's request for an operating system for Intel 8086/88 processors. It was written by one person gradually taking form over a couple of weeks to test the new Intel 8086 processor. It was never intended to be a full-featured operating system. It was only a bit of test code hammered out to accomplish one specific, limited task.

Does this sound familiar? I'll give you a hint. MS-DOS's original name was Quick and Dirty DOS (Disk Operating System), a truly appropriate moniker. Sure, Microsoft polished it up a bit before releasing it, and added various improvements over the life of the product, but mostly, it was still the quickie test program that a lone programmer originally created. (For you Oliver Stone fans, think of Bill Gates as playing the part of the second programmer on the grassy knoll.) Thus, the entire PC industry was held back by the limitations built into the design of a program that was never intended to be used beyond those first couple of weeks of testing.

There are two other closely related factors involved in the spread of the Y2K bug. The first is what I have previously referred in my columns as the Programmer as an Artist Syndrome. This is a pervasive attitude that is spread from programmer to programmer and is significantly aggravated by the fact that the non-computer public at-large holds programmers in awe due to their knowledge of computers. Programmers have not been properly indoctrinated with the idea that they are craftspeople, charged with the intrinsically creative task of making useful programs, not merely artists making programs strictly either for their esthetic value or simply as an outlet for their creative urges with little or no thought of usefulness, maintainability, or reliability. The most obvious sign of this is when a programmer is personally offended when either a supervisor or a peer review team constructively criticizes his or her precious code. It's like telling Picasso that his representation of women in his paintings is not realistic enough.

The second factor is what I call the Macho Programmer Syndrome, though it is definitely not exclusive to males. This attitude problem at its most severe instills programmers with the belief that the more convoluted, obtuse, and downright unreadable their code is, the better. Their perceived job is to come up with incredibly unique, imaginative and original algorithms (formulas), and if nobody else can understand them, so much the better.

An example of this is an interview with Apple's lead programmer for the original Mac that I read in a national magazine years ago. He openly bragged that when his team found a bug in a particular subroutine (a logical subsection of a program) more than about two weeks after the subroutine was written, even the programmer who wrote it, much less any other team member, couldn't understand or debug it. His solution was to simply throw out the entire subprogram and rewrite it from scratch. He was very proud of his team's ability to write this stuff. Luckily, this attitude is not as common as it used to be.

Going against these attitudes, is the trend, sadly appearing mostly in the corporate world, but not in smaller companies, toward programming methodologies. These methodologies include such practices as establishing corporate standards, teaching structured programming, design, and especially testing techniques, including peer review teams on all projects, encouraging general business education among computer specialists, and various other quality-related policies. These policies increase direct development costs but greatly reduce future cost of debugging, maintenance, and enhancements.

In a perfect world, these methodologies would win out. In our real world of downsizing, mergers, acquisitions, and especially the popularity of the next-quarter-results-only style of management, there's lots of backsliding. In addition, computer departments are considered an expense, almost all programmers are salaried, and computing experience is not valued by management. Just to make things even more interesting, the Peter Principle butts in so that expert programmers are sometimes promoted, only to become lousy (And often frustrated) managers.

Because of these management factors computer training is often the first thing to go when budget cuts are needed. Staffing budgets come next. Why pay for extra staff or for outside temps when you can simply force the existing salaried staff to work 12 hours a day, seven days a week, rotating between day and night shifts to avoid overloading your mainframe or your LAN.

This is not an exaggeration. It is from personal experience. It was, in fact, the final straw in my decision to leave the corporate world and turn my part-time business into a full-time proposition. I figured that if I am going to get an ulcer or even a heart attack from overwork, I might as well do it in a comparatively pleasant environment of my own choosing. Dilbert is really a lot closer to nonfiction than most people realize.

If the older, more experienced programmers leave for better jobs elsewhere, then just hire lower-priced beginners (or import foreign employees with special visas who you can oppress even more than your old employees), with the added benefit of having avoided paying some pensions. Or, you could just make the remaining staff work even longer hours. If the older group doesn't leave fast enough, then find ways to either make them even more miserable or even to outright fire them, as long as you document the firings well enough to avoid age discrimination suits, of course.

While some of this might sound like programmer-bashing, I intend it to be a criticism more of the education, ongoing training, and management of programmers. I am definitely in favor of increased self-responsibility and higher ethical standards for individual programmers, but I think that I have made a case for providing a better programming environment to avoid messes like the Y2k Bug.

Dealing with programmers who have never worked in a supervised team environment is basically a gamble. Some of them have realized the necessity of quality control techniques on their own. Others have learned it in classrooms. But many haven't been exposed to these techniques at all, and still others left the corporate world specifically because they felt oppressed by ham-handed attempts to implement quality control methods. It's a throw of the dice.

Next month this series will continue. For now, does anybody know of some good land with a defensible perimeter for sale somewhere in the wilds of Montana?

WHY 2K? - Part 3

by Wayne M. Krakau - *Chicago Computer Guide*, April 1999

Wayne is president of Krakau Business Computer Systems, Inc., a systems integration firm that is a Novell Gold Authorized and Novell IBM Host Master Authorized Reseller. He has been working with computers for 22 years. He holds an M.B.A. in Marketing and a B.S. in Information Science and was originally educated in Computer and Information Systems Analysis Engineering. He holds CNI, MCNE and CNTE ratings and all five of the original LANDA Certifications. Wayne can be reached for questions, comments, or topic suggestions at (847)298-7695 or via E-mail to wkrakau@krakau-inc.com. For past columns and company information go to http://www.krakau-inc.com.

To summarize the answer to the title question, as explained in the previous columns of this series, the world of Dilbert (except perhaps, for the lack of facial hair on the male techies) is real. Management has run amok, employees are treated horribly, and the entire system is based on giving rewards for pumping out poorly designed, buggy programs. Those who espouse quality are whacked on the head with a mallet, just like the gophers (or whatever they are) in a popular arcade game, and forced back down into their holes (also called "cubicles").

All right, I'm exaggerating a little, but not much. Here are a couple of examples that are not exaggerated. First, how do you create quality programs when there is a complete reorganization of Management Information Systems (the computer department) every two weeks? This included changes of managers, supervisors, team members, and areas of responsibility approximately every two weeks for a period of months.

And, remembering my remarks about the Programmer as an Artist Syndrome, how do you keep productive between dodging rubber band wars when you walk in on a December morning and find two four-foot tall electrically animated 1890's-costumed boy and girl figures (the kind you see in department store windows) flailing away at each other under the Christmas tree while locked in a sexual position best known for its numeric designation? (I hurt myself laughing at that one!) Each day thereafter, they were in a different position. Later the Christmas tree was decorated with red panties.

Doesn't that sound like something straight out of Dilbert? It's amazing what people (myself included) will do to try to avoid burnout when faced with bizarre working conditions.

So much for explanations of Y2K-related (and many other) bugs. Now, what are you going to do about it? First, you have to decide on which of two major strategies to take, using either free or commercial testing software. If you have a small, very structured environment with a limited number of tightly controlled programs, you can use a mostly manual approach. In this case you start out by using one or more of the dozens of free programs available to test individual PCs for Y2K compatibility. Almost every company that makes a full, commercial Y2K program also makes a free, simple one to test the underlying PC hardware for Y2K compatibility. In addition, many national magazines publish their own free Y2K utilities. Needless to say, Web access is required.

This free software will automatically put your PC into one of three categories, compatible, incompatible but patchable (basically a system clock problem), and simply incompatible. If the tests result in a compatible rating, then you're home free. If your system is patchable, then many of these free testing utilities will automatically install a free software patch. From that point on, every time your computer boots up, the patch will override incorrect information and provide valid dates. If the system is rated incompatible, you might still be able to salvage the system by replacing the BIOS (Basic Input/Output System) chip on the motherboard, if an appropriate chip is available.

You will have to use your own judgement as to whether it is worth the trouble to keep an older machine that needs a new chip, or just toss it. Note that even older machines may be useable as hardware assembly and repair training aids for some charitable organizations and that PCs are considered

hazardous waste. Give a call to the National Christina Foundation at 1-800-CHRISTINA or check their Web site, www.cristina.org for further information.

After testing the PCs themselves, you must individually research each and every piece of software on every PC. You then must upgrade, patch or reconfigure each as specified by the publisher to make that software Y2K compatible. Now you know why I specified that this fully manual method was only practical for environments with limited complexity.

For more complicated and less structured environments, commercial Y2K software is required. (Of course you could always take the traditional MIS attitude which would hold that if your System Administrator is on salary, working that Administrator 24 hours a day, seven days a week to avoid paying for commercial Y2K software is permissible!) Not only will this type of software do the hardware test, and, apply a patch if necessary, but the better products will even coordinate version checking and will scan data for possible problems.

In version checking, the product is supplied with a database of Y2K facts about various versions of common software, updateable via the Web. It checks your systems for programs and uses its database to tell you what patches, fixes, or upgrades are required. While you still have to apply these patches, fixes, and updates yourself, this still saves a lot of time and potentially some trial and error experimentation.

Data scanning finds documents, spreadsheets, and databases which contain date-oriented data that is stored in ways that might make it vulnerable to Y2K problems. A human will still have to examine the results of the scan and decide what to do about it.

We have come full circle and are back to the WHY in the title. If you are informed by a software publisher that a specific version of their product is Y2K compatible, are you safe? Not necessarily! The people who allegedly cured the Y2K bugs are the same people who put it in there in the first place. Wait - it gets worse! Those people who created the bugs in the first place now have enough seniority to get assigned to new development projects, so the ones doing maintenance tasks like bug-fixing are the beginners right out of college or trade school, so they are probably even more likely to mess up the bug fix and maybe even generate their own new bugs.

Remember that Windows 98 was supposed to be Y2K compatible right out of the box, but now has a Y2k fix available. Just because a company says its software doesn't have bugs doesn't mean it's true.

Finally, as in the example I gave in the first article in this series in which I explained how I was ordered to put in a limited decade-only date fix and found previously entered limited fixes, there is a right way and there are many wrong ways to fix the Y2K bug. The right way is to put in a fix that covers September 9, 1999 (9999 in computer format), January 1, 2000, February 29, 2000, all future leap years, and any artificially generated date limits created by programs which arbitrarily pick a start date for the creation of the universe and count up from there, potentially running out of digits at some arbitrary date in the future.

Sadly, lots of companies, government agencies, and especially consulting firms are allowing limited fixes that are marginally quicker to install and test, but will break again in the future. In many of theses cases, the internal MIS folks or the external consultants are not telling management that they are taking these shortcuts. Their excuses are the usual ones. "It looks good in my quarterly budget to fix this problem cheaply." "This program won't be around long enough to break." And the ever popular, "I'll have another job somewhere else by the time this program breaks, so I won't be around to get blamed." Of course, the outside consultants have the added motivation of getting paid to fix the problem again in the future.

Don't you just love this industry? Oh, and in case you think I am being too cynical about this issue, I've got many stories to back up my attitude, from both personal experience and first-person accounts, covering mainframes, minicomputers, and PCs, and involving everyone from in-house staff, to management, to small individual freelance programmers, and to some of the largest consulting firms in the country. I'm sure that Dogbert, the ultimate consultant (and a companion of sorts to Dilbert), would be proud.

HOBBY 95/98 - Part 1

by Wayne M. Krakau - *Chicago Computer Guide*, May 1999

Wayne is president of Krakau Business Computer Systems, Inc., a systems integration firm that is a Novell Gold Authorized and Novell IBM Host Master Authorized Reseller. He has been working with computers for 22 years. He holds an M.B.A. in Marketing and a B.S. in Information Science and was originally educated in Computer and Information Systems Analysis Engineering. He holds CNI, MCNE and CNTE ratings and all five of the original LANDA Certifications. Wayne can be reached for questions, comments, or topic suggestions at (847)298-7695 or via E-mail to wkrakau@krakau-inc.com. For past columns and company information go to http://www.krakau-inc.com.

If you use Microsoft Windows 95 or Windows 98, you probably have a hobby that will take up endless hours of your time, and will keep your mind off of those annoying day-to-day concerns that can really take their toll - like getting your work done. That hobby is the care and feeding of these desktop operating system siblings. Keeping Windows 95 or 98 relatively stable (totally stable being realistically unattainable), can, and in some cases has, eliminated any efficiency gained by transitioning from a manual system to a computerized one in the first place.

For most people, it is a hobby by necessity, not by choice, excluding, of course, those who actually get so much of a kick out of constantly tinkering with their systems that they would rather do that then get their work done. (Hmmm. That last save didn't go as fast as usual. Maybe I'd better tweak some of the system performance parameters.)

In a network environment, you have even more potential destabilizing influences to worry about. This includes network cards, client software (to communicate over the network), and the added complication of using the network versions of applications software which may have their own unique idiosyncracies

The first line of attack that I use against Windows instabilities is a set of utilities sold as part of Network Associate's (www.nai.com) McAfee Office. Network Associates bought the companies that produced several of my favorite utility programs and now sells them packaged together at a far better price than I paid when I acquired them separately. The availability of network licensing based on the number of workstations make McAfee Office even more cost effective. Realistically, almost anything that works is cheaper than spending hours (or paying for some else's time) trying to beat Windows 95/98 into submission.

The main disadvantage to using any repair-oriented utility such as this one is the added potential for the utility itself to destabilize Windows, especially when that utility is installed by nontechnical users. Contrary to the propaganda put out by all of the utility program publishers, these programs are definitely not as carefree as they are advertised to be. This holds for McAfee Office and all of its competitors. In addition, when multiple utilities are loaded on one computer, they can fight with each other for control of the safety mechanisms that they have in common, causing additional crashes.

While McAfee Office contains many different products, including the popular VirusScan anti-virus program, the key utilities for keeping Win 95/98 healthy are Uninstaller, First Aid, and Nuts & Bolts, itself a collection of multiple utility programs.

Uninstaller does exactly what its name implies. It removes programs and files in a reasonably safe and complete manner. The automated uninstall features included in many programs often blindly remove only a base set of files, leaving behind potentially harmful files and directories. They frequently neglect to check if the files that they remove are potentially used by other programs on the computer. Also, almost none of them have an Undo feature to use in case the uninstallation causes problems, or even if you simply change your mind. These problems can cause mysterious crashes that are very difficult to debug.

Uninstaller eliminates these problems by adding cross reference checks and an Undo feature. While its analysis is not always perfect, in that it occasionally errs on the side of safety and doesn't delete some truly removable files and directories, it's the best uninstall package that I've found.

Uninstaller also some handy features for general system cleanup that allow you to find useless files and registry entries. The useless items are then graded for safety of removal and backed by an Undo feature. If you only delete the green-labeled items and leave the yellow and red ones alone, your system should remain healthy. Of course it takes human judgement to decide if a particular file that is marked as not absolutely necessary for Window's survival is actually required because it contains valuable data, or is related to some function that you may need in the future.

One function of Uninstaller, and all of its competitors, that causes more trouble than its worth is its installation monitor. All of the installation monitors that I have tested, including this one, have a tendency to fight with and confuse installation programs causing either incomplete, and therefore dangerous, installations, or outright crashes. I now turn this feature off. The only real disadvantage to turning it off is that every time you launch Uninstaller (and again, any of its competitors), you have to wait for the program to scan your system to find out the current state of all programs. Since this program is not used that often, I would rather put up with this minor inconvenience than risk the mayhem that this feature can cause.

The second important feature of McAfee Office is First Aid. It is an all-encompassing diagnostic program which includes an added crash prevention and recovery safety feature. Its main function is to scan all of the hardware and software in your system, looking for potential problems. This could be a hardware conflict, a missing program file, or an invalid registry entry. Let me warn you that this thorough scan takes a long time to complete. After scanning, First Aid shows the problems, rated by seriousness, and offers both automatic and manual fix buttons, as well as an ignore option.

In most cases, the automatic fix works just fine. The manual fix option, with hints provided, works in most other cases. Repair programs are not perfect, however, and each has its own strengths and weaknesses, so occasionally you hit a dead end. First Aid then offers a guided question and answer session leading to more advice. If that fails, it leads you to connect automatically with its massive troubleshooting database on the Web. This also gives you access to support newsgroups.

Finally, possibly the most valuable feature of First Aid is Windows Guardian, a crash recovery and prevention tool. It sits it the background and monitors Windows. When a program starts to crash, it intercepts the crash and pops up a window offering to either attempt a recovery or to close the offending program. Much of the time, the recovery option works. When it doesn't, at least the close option almost always (again, none of these programs are perfect) allows you to shut down the program and keep Windows alive.

After closing the program successfully, sometimes you can just restart it and continue working. At other times, it won't restart, or, you might notice that Windows starts misbehaving. In this case, the damage to Window's environment has been too great, so you have to save your work from other programs, close all of them and restart Windows.

While Windows Guardian can't fix every crash, it is, by far, the most effective crash recovery program that I've seen. Even in cases where it can't fully recover, it allows you to click on a Details button to see exactly what happened to the crashing program. If you have the appropriate technical background, this can provide valuable information toward finding a long term solution to that particular program's problem.

Next month I will complete coverage of my favorite utilities and then move on to various other things you can do to try to keep Windows 95/98 behaving itself. (FYI: Whacking the monitor on the side while cursing vehemently doesn't help. I've already tried that.)

HOBBY 95/98 - Part 2

by Wayne M. Krakau - *Chicago Computer Guide*, June 1999

Wayne is president of Krakau Business Computer Systems, Inc., a systems integration firm that is a Novell Gold Authorized and Novell IBM Host Master Authorized Reseller. He has been working with computers for 22 years. He holds an M.B.A. in Marketing and a B.S. in Information Science and was originally educated in Computer and Information Systems Analysis Engineering. He holds CNI, MCNE and CNTE ratings and all five of the original LANDA Certifications. Wayne can be reached for questions, comments, or topic suggestions at (847)298-7695 or via E-mail to wkrakau@krakau-inc.com. For past columns and company information go to http://www.krakau-inc.com.

This is the second installment in my series of columns about the hobby that many people have (usually unwillingly) these days. That is, keeping Windows 95 and 98 systems running. In particular, it continues with coverage of additional programs within one of my favorite stability-enhancing tools, McAfee Office, from Network Associates.

Nuts & Bolts is a major component of McAfee Office. It is a conglomeration of individual utility programs collected within a menu system equipped with an incredibly annoying sound scheme (which I would dearly love to disable). The menu divides the programs into four categories, Repair/Recover, Clean/Optimize, Prevent/Protect, and Secure/Manage. It's obvious that this menu is simply a front end to a bunch of unrelated tools, but the underlying programs are good enough that you are willing to put up with this minor disadvantage. While Nuts & Bolts has 23 different programs within it, I'll just be covering the most valuable of the stability-related ones.

The first interesting program within Nuts & Bolts is PC Checkup. As its name suggests, it performs a diagnostic check of your PC's hardware and software. In theory, this action duplicates the main function of First Aid, one of the other components of McAfee Office (covered in last month's column), however, PC Checkup is just different enough that its strengths and weaknesses in both diagnosis and repair complement rather than compete with First Aid's capabilities. The idea here is that if one program can't either find or fix the problem, sometimes the other can.

PC Checkup will run either in normal or advanced mode, the choice of which determines just how thorough (and time-consuming) the tests will be. It also presents the various categories of tests and allows you to enable or disable each. This is a great way to narrow down the tests when you already have some idea of what's wrong. As with any of this type of program, it isn't perfect at either diagnosing or repair, but it does catch and fix a lot of problems, including some that First Aid misses. Note that the same could be said for First Aid. It sometimes succeeds in cases where PC Checkup fails.

Discover Pro is a Nuts & Bolts program that provides extensive, detailed information about your PC that could be valuable to technical types when analyzing problems. It also includes a limited testing feature. Sadly, this program is somewhat of a troublemaker in that it has a nasty habit of either destabilizing Windows 95/98 or completely locking up the whole system. Because of this, I seldom use it, and when I do, I make sure that nothing is running that I wouldn't mind crashing.

There is a handy little Undelete program. It's simple and it works. If your problem is related to a missing file, it's there to help.

The Cleanup function isn't as good as the one in Uninstaller, another Nuts & Bolts component, so I usually don't bother with it.

The Registry Wizard is an important troubleshooting and safety tool within Nuts & Bolts. It allows you to backup and restore the Registry, clean up unneeded Registry data, repair the Registry, and optimize the Registry. It often finds and fixes Registry problems that other programs miss. Its only weakness is that it frequently reports false errors when it encounters Registry items which contain perfectly legitimate sub-parameters. It is somewhat unnerving to see, for instance, more than 300 errors reported and only a dozen fixed, until you recognize this weakness.

Shortcut Wizard finds and attempts to fix any Windows shortcuts that no longer work. If it doesn't succeed in fixing a shortcut, it assists you in manually fixing it.

Bomb Shelter is made to hide in the background and intercept upcoming crashes. First Aid's Guardian feature is superior and far less likely to cause problems, so I don't use Bomb Shelter.

WinGuage dynamically monitors and displays several critical internal Windows statistics. It occasionally crashes, and its information usually isn't vital, so I seldom use it.

EZ Setup is a detailed control panel for many Windows 95/98 settings. The most valuable information it presents is in its Run/Add/Remove screen. It displays all programs that run automatically at Windows startup, including those programs stashed in various obscure corners of the Registry. For each program, it allows you to add, remove, or edit the command that executes that program. It gives the same information and options for programs in Window's Add/Remove Program Menu. EZ Setup also contains a lot of other options for various aspects of the system, though none as valuable as its Run/Add/Remove screen.

As you can see, Network Associates McAfee Office (Yikes! What a mouthful.) is a valuable troubleshooting and stability-enhancing tool. I sell it with every computer (usually as part of a site license) and I offer it as an add-on to existing systems.

Next month, I'll cover some of the free things that you can do to help keep Windows 95/98 alive and ticking. (Threatening the PC with a hammer doesn't seem to help at all.)

HOBBY 95/98 - Part 3

by Wayne M. Krakau - *Chicago Computer Guide*, July 1999

Wayne is president of Krakau Business Computer Systems, Inc., a systems integration firm that is a Novell Gold Authorized and Novell IBM Host Master Authorized Reseller. He has been working with computers for 22 years. He holds an M.B.A. in Marketing and a B.S. in Information Science and was originally educated in Computer and Information Systems Analysis Engineering. He holds CNI, MCNE and CNTE ratings and all five of the original LANDA Certifications. Wayne can be reached for questions, comments, or topic suggestions at (847)298-7695 or via E-mail to wkrakau@krakau-inc.com. For past columns and company information go to http://www.krakau-inc.com.

This is the third installment in my series of columns about arguably the most popular hobby in the world, as measured by the number of participants as opposed to the number of willing participants. That hobby is trying to keep Windows 95 and 98 systems stable. Having already covered my favorite commercial stability-enhancing product, I'll now cover some free methods of partaking in this hobby. Free, that is, only if you consider your own time worthless.

The first tool in your free arsenal is Microsoft's REGEDIT, the Registry Editor. It is also the most dangerous tool.

On the good side, it can take a snapshot of your System Registry and save it as a file. If your computer starts acting wacky, (a technical term, of course) REGEDIT allows you to restore the Registry from one of those backup files, hopefully restoring sanity to your system.

Any time you are about to make any serious change to your system, you should use REGEDIT to back up your Registry. The most obvious change would be the installation of either a new program or an updated version of an existing program. However, so many things happen automatically and, in some cases, accidently that you might want to do one backup at the beginning of each day for a machine that regularly exhibits instability. If you see an approaching thunderstorm, you might want to back up your Registry. Considering some of the strange sensitivities that I've seen on Win 95/98 computers, you might want to back up your Registry even if you just want to look at your computer from across the room!

Now we get to the dangerous aspects of REGEDIT. First, restoring one of the backups that I've just encouraged you to make is definitely not without hazard. If you restore your Registry at an inopportune moment, you can screw up your system worse than it was in the first place. You should make yet another backup just before restoring in case the restore initiates a disaster. That way you have at least a fighting chance of getting back to where you were in the first place.

The other dangerous aspect of REGEDIT lies in its name, the Registry Editor. Anytime you can edit a configuration file as little understood and poorly documented as the Win 95/98 System Registry, there will be plenty of opportunity to mortally wound the file. One tiny misunderstanding or typo and your Registry could be rendered irretrievably corrupted. Editing the Registry is a high-risk activity, and the result cannot always be undone by restoring the aforementioned backup copies. Edit at your own risk, not without due consideration of both your own skills and those of the person or document advising you to perform the edit.

One of the most valuable free "tools" is access to patches and updates via Microsoft's web site (entry point www.microsoft.com/support). For owners of a single Win 98 system, using these files is almost a complete no-brainer, since automatic update is a built-in feature of Win 98.

For owners of multiple Win 98 systems, accessing Microsoft's site and downloading each separate update and fix file for each individual computer is ridiculously labor and time intensive. For whatever reason, Microsoft has been very secretive about the cure for this problem. (Just to be ornery, perhaps? I am reminded of Mel Brooks' line in his movie, *The History of the World, Part I*, "It's good to be King!") The solution is to go to www.microsoft.com/windows98/downloads/corporate.asp and download the individual files. After that you can apply them to your computers, one by one.

There is one caution to this method. While Win 98 updates are programmed to automatically know when they are appropriate (they are supposed to refuse to install when they are not needed), it is still possible to waste time and possibly cause problems by just randomly applying these updates.

When you allow Microsoft's Web site to directly update your system, this all sorts itself out automatically. When you download the files and apply them yourself, you are responsible for organizing the updates. Microsoft makes this harder by not including separate documentation or README files with the updates. If you execute the downloaded files, they immediately launch their installation programs without a thorough explanation. (I can't remember any other company that neglects the trivial task of self-documenting its downloadable files.) Just prior to downloading each file you must either print out the Web page containing the documentation, or make copious notes. Then you must keep the printouts or notes with the files so you can coordinate between them. ("It's good to be king!") One odd thing that I've noticed is that I've had better luck accessing Microsoft's site (and not just the updates section) with Netscape's Navigator than with Microsoft's own Explorer, regardless of whose computer I'm using! Welcome to *The Twilight Zone*.

Next month I'll continue covering update files which will lead to tips on handling the most common Win 95/98 problems that I've encountered. (I'll report later on the effectiveness of melting down all of your jewelry and building a golden idol in the image of Bill Gates.)

HOBBY 95/98 - Part 4

by Wayne M. Krakau - *Chicago Computer Guide*, August 1999

Wayne is president of Krakau Business Computer Systems, Inc., a systems integration firm that is a Novell Gold Authorized and Novell IBM Host Master Authorized Reseller. He has been working with computers for 22 years. He holds an M.B.A. in Marketing and a B.S. in Information Science and was originally educated in Computer and Information Systems Analysis Engineering. He holds CNI, MCNE and CNTE ratings and all five of the original LANDA Certifications. Wayne can be reached for questions, comments, or topic suggestions at (847)298-7695 or via E-mail to wkrakau@krakau-inc.com. For past columns and company information go to http://www.krakau-inc.com.

Here's the fourth and last in my series on my least favorite hobby, the upkeep of Windows 95 and Windows 98. Since word on the street is that, contrary to previous announcements, Microsoft is preparing two more versions of Windows 98, this hobby could still be quite common for several years to come.

One of the most common problems I see is the gradual corruption of Windows 95 due to the improper application of patches. (I must plead guilty to having done this more than once before I caught on.) Unlike patches for Windows 98, Windows 95 patches are inherently dumb, and will happily overwrite perfectly good files with inappropriate ones. You must read the on-screen instructions very carefully, and print them prior to download, since most of these files don't have internal README or other explanatory files. Even then, many don't have complete descriptions of their limitations.

The nastiest of these poorly documented limitations has to do with Windows 95 Operating System Release 2 (OSR2). OSR2 is the version of Windows 95 that came bundled with the vast majority of computers that have Windows 95.

According to three different sources within Microsoft's Tech Support, OSR2 already has ALL of the communications-related patches within it and most of the other patches as well. There are very few patches which should be applied to this release. The Windows 95 patches you find on Microsoft's Web Site are almost exclusively meant only for the original version of Windows 95, OSR1. That version came bundled with computers prior to the release of OSR2 and also all in separately purchased copies. OSR2 has never been available at retail. It was only available bundled with new computers.

Note that all three of the Tech Support personal that I talked to, as well as those who answer questions in Microsoft's Tech Support Forums denied having access to a comprehensive list of what patches were appropriate for which version of Windows 95/98, or any detailed information about what combinations individual communications-related files were valid. They could only offer to walk me through fixing individual systems at $35 a pop.

Before fixing this problem, locate the WINSOCK.DLL file in the \WINDOWS directory and search the \WINDOWS\SYSTEM directory for any files with "SOCK" as part of their names and you'll find the communications-related files in question. When these get out of synch, the most common symptoms are randomly losing both the printing and the "SAVE AS" functions within your browser. Once lost, these functions stay lost until at least until you shut down your browser and reestablish your communications with the Internet. Sometimes you are forced to reboot, too.

When this condition is at its worst, all browsers will simultaneously lose functions. In a less severe version, Internet Explorer loses functions, but you can start up a Netscape session and keep on saving and printing.

Another common symptom that occurs much more frequently than the browser problems strikes only America Online (AOL) users. They lose the ability to print their e-mail messages. Sometimes even rebooting doesn't help that one.

A solution that has worked for me (found by combining the advice obtained during the three separate Microsoft Tech Support calls previously alluded to with info gathered from my own experimentation), is to try to get all of the SOCK files except for WSOCK32N.DLL to the same date as

they originally were. Extract files from the cabinet files (*.CAB) on the original Windows 95 CD, if necessary. For WSOCK32N.DLL, keep the latest version that you find on your computer. You may have to experiment a bit to find a combination of communications files that works, since some of the files you find may have originally come from patches or other programs.

Your experiments may come to naught, however, unless you know the magic process that completes the fix. Without it, you might alter the frequency of the problems, but you won't actually stop them. This is particularly insidious if you correctly determine the combination of SOCK files that should work. In that case the problems will occur less frequently, but will keep coming back to haunt you.

The solution is to toast the young farm animal of your choice on an altar adorned with the likeness of Bill Gates. What? You say that your religious or ethical convictions forbid such ceremonies? All right. I'll let you in on the alternate procedure. Delete and reinstall your default printer. That's the key process that is needed to make your file manipulations truly take effect. Don't just switch to another printer as the default and then switch back. You must delete the printer. Hey, I never said this stuff was either easy or obvious.

An important preventative measure for Windows 95/98 problems is to use a disk imaging product to take a snapshot of the hard disk. This option is available for either those computers attached to a LAN or those with enough local auxiliary storage to hold a compressed duplicate of their disk drives. With the licensing options available for quantity purchases for a LAN, imaging software can be quite inexpensive.

This type of software allows you to take a snapshot of your disk prior to installing some software or making some other change to Windows 95/98. It's also a good idea to make a baseline copy occasionally to cover you in case some program makes unannounced changes to the Windows 95/98, a process that is, sadly, becoming more common. If Windows 95/98 starts misbehaving, just back up your changeable data files (documents, spreadsheets, databases, etc.) and then restore the image file. Once you put your data back, you got a stable system.

Yet another stability assuring tool is Novell's Z.E.N.works Starter Pack. It comes bundled with NetWare 5 and is downloadable for those with NetWare 4.x. Z.E.N.works stands for Zero Effort Networks, one of the more grandiose product name in computing. The Starter Pack version is free. You have to pay if you want the complete Z.E.N.works version, but the Starter Pack version has enough features to make it worthwhile for both Windows 95/98 users and NT users.

Z.E.N.works Starter Pack allows you to automatically repair applications, so that a missing, corrupted, or version-switched program file won't cause a crash. The product monitors the original installation of a program to detect all of the files that the program needs to run. It then stores a copy of those files on the server.

When an individual workstation is about to experience a crash due to the aforementioned problems, Z.E.N.works steps in and replaces the appropriate files from its storehouse on the server, thereby allowing the application to continue running normally after only a slight delay. That means one less unsatisfied user and one less call to a harried system administrator.

Another valuable feature of Z.E.N.works Starter Pack is its ability to ease management of the desktops across the network. You can force standardized configurations to the desktop based on both computer and user ID. You can also restrict access to many potentially dangerous aspects of Windows 95/98 and NT. (Think of it as a greatly enhanced version of Windows 95/98 Policy Editor.) Obviously, this trick only works in companies where these restrictions are politically feasible. (Didn't the French Revolution start this way?)

My final suggestion is to look at upgrading to Windows 98 Second Edition (SE). While it's not perfectly stable, at $19.95 for Windows 98 users and $109 for Windows 95 and Windows 3.x users, it might be a lot cheaper than spending a lot of time either debugging Windows 95/98 or reinstalling from scratch. The mere act of applying the upgrade over a previous version of Windows 95 or 98 might be enough to eliminate errors due to the gradual corruption of the overall configuration of the operating system. If you think of the upgrade as merely a stability-enhancing tool as opposed to the creator of a perfectly stable environment, the decision to upgrade might be a bit more palatable.

Now I'm going to count the gray hairs that I've gotten trying to keep Windows 95/98 systems alive and kicking. I am anxiously awaiting future versions of Windows 98 and of the Workstation version of Windows 2000 in the hope that I won't have to worry so much about workstation stability in the future.

DESIGNING FOR SAFETY - Part 1

by Wayne M. Krakau - *Chicago Computer Guide*, September 1999

Wayne is president of Krakau Business Computer Systems, Inc., a systems integration firm that is a Novell Gold Authorized and Novell IBM Host Master Authorized Reseller. He has been working with computers for 22 years. He holds an M.B.A. in Marketing and a B.S. in Information Science and was originally educated in Computer and Information Systems Analysis Engineering. He holds CNI, MCNE and CNTE ratings and all five of the original LANDA Certifications. Wayne can be reached for questions, comments, or topic suggestions at (847)298-7695 or via E-mail to wkrakau@krakau-inc.com. For past columns and company information go to http://www.krakau-inc.com.

Afer observing a new client company unnecessarily brought to its knees by the crash of its main file server, I decided I'd better publish some information on designing a LAN so that a single point of failure will be only an inconvenience, as opposed to a catastrophe. The irony is that most of these safety-related design techniques are comparatively inexpensive, so implementing them is really a form of cheap insurance that won't break most budgets.

Despite my company's current NetWare centric attitude, most of the techniques and products covered here are applicable to any Network operating System (NOS). (Note that since Microsoft's Windows 2000 Project Manager has publicly claimed, as published in the August 18 issue of *Computer World Online News*, www.computerworld.com, that "Windows 2000 users would experience thirty times fewer crashes than they do with NT," one simple safety technique would be to run NetWare at least until a fully patched version of Windows 2000 is available, presumably six months to one year after Windows 2000 is released. This is, of course, assuming that you have an overwhelming need to switch your servers to Windows 2000.)

The first design technique I'll suggest is to use true dual servers. I know that this suggestion seemingly violates my "inexpensive" claim, but for those companies that lose large amounts of money any time that their servers are down, anything that improves uptime is cheap.

My favorite software for implementing true dual servers is Vinca's StandbyServer (www.vinca.com) software. It comes in both NetWare and NT versions. Novell likes it so much that they now sell it. With StandbyServer, a secondary server is set up with special drivers that make it look like a mirrored disk drive on the first server, even though it is really an independent server running a separate copy of the NOS. Its disk drives are kept in synch with the primary server's disk drives via a direct network link between the two machines. For added safety, the two servers can be separated so that major disasters (fires, floods, etc.) won't affect both.

When the secondary server notices that the primary server is no longer responding via either the dedicated link or any of the general-purpose network segments that the two servers share (presumably because the primary has stopped functioning), it immediately reboots and wakes up using the identity of the primary server. Since it has the original primary server's identity, all workstations and other network-connected devices think they are still talking to the original primary server.

On a system that I set up for an engineering firm using NetWare 4.11, the 22GB disk array on the former secondary (now primary) server mounted in less than 90 seconds. That's easily fast enough for the NetWare client software on the workstations to automatically reconnect each workstation and reestablish all drive mappings. This means that users merely have to wait a little while and then click on "RETRY" to continue working. Those who don't happen to read from or write to the server won't even have to do that. Note that NetWare 5 should, theoretically, mount volumes even faster than NetWare 4.11.

If you can't justify the expense of true dual servers, you can use the concept of twin servers (my name for it) that most of my clients use. With this method, you purchase one workstation at the same time as you purchase your file server. It should be a twin of the file server, with the same motherboard, cards, and disk non-NOS (to be explained later) disk drives.

If that machine is used on a day-to-day basis as a workstation, you only have to worry about justifying the incremental cost increase over a regular workstation needed to make this machine a twin. That might include the cost of a better overall model of computer. It could also include upgrading to a faster CPU. It would certainly include the added cost of extra RAM, and probably a pair of SCSI controllers (one for disks and one for a tape drive). It might also include either a high-performance or even a multi-port network card. It takes more time to set up and maintain twins, too, so be prepared for some additional labor expenses.

The key to twin servers is using external tape and NOS disk drives in concert with one or more (for redundancy) internal non-NOS drive. If the primary file server crashes and can't easily be revived, just connect the external tape and disk drives and the UPS (You do have one attached to your file server, don't you?) to the secondary file server and reboot it, selecting the predefined server mode. In a matter of minutes, you've gone from having a completely idle network to just being short one workstation, a survivable event for most businesses. I've taught this drive-switching trick to complete computer amateurs, and have observed them using it successfully. It's not difficult to learn.

External tape drives don't cost that much more than internal tape drives. While I'll cover the details of disk drive options later, I will state now that external disk enclosures that can handle plain, mirrored, or duplexed drives are inexpensive. Also, most aftermarket RAID (Redundant Array of Independent Disks) systems are only available as external units. If you use RAID system from a computer manufacture that uses internal RAID systems, you will have to manually transfer each individual drive from one machine to the other and probably have to manually reconfigure the RAID controller, too, eliminating much of the inherent ease-of-use of the concept of twin servers.

To make twin servers work in a NetWare environment, you must have identical MS-DOS partitions on the non-NOS internal drives of the two servers. Since a NetWare server's identity is contained on its NOS drives, moving the external drives from one machine to the other will also move the server's ID, so that a quick switch might be fast enough for the workstations to automatically reconnect. Realistically, it takes a few minutes of testing (unless there is smoke pouring out) to determine that you should give up on the main server, so rebooting workstations is often required. Keep in mind, however, that the alternative is to sit and stare at a dead network while waiting for the primary server to be fixed. Even if you have in-house techies and a stock of spare parts, this could take some time.

Since MS-DOS is no longer a viable workstation operating system for most people, I use PowerQuest's (www.powerquest.com) PartitionMagic to split the internal drives into one MS-DOS partition at the beginning and multiple 16-bit (for compatibility) Win 95/98 partitions. (Since I also like PowerQuest's DriveImage, I usually sell DriveImage Pro, which contains both DriveImage and PartitionMagic, since it is available at an incredibly low per seat cost when purchased for the whole network.)

PartitionMagic also contains PowerQuest's BootMagic, which allows you to select which partition (and operating system) to boot on. I set the primary server to default to booting MS-DOS, which will then automatically start NetWare. I set the secondary server to booting Win 98. In each case, I use BootMagic to display a menu with a 30-second countdown that allows you to override the default operating system selection.

After initial configuration and testing, my only concern is that the non-NOS disks in each of the servers stay in synch. If I update NetWare, I may have to copy information between the MS-DOS partitions. If I update Win 98, or add programs to it, I must copy any changed information between Win 98 partitions. The easiest way to do the Win 98 copying is with DriveImage. I can create an image of one machine, copy it to the network, and "plant" it on the other computer, though I would want to manually synchronize the two MS-DOS partitions first.

I finally got tired of hearing clients protest the purchase of an old, theoretically dead, operating system like MS-DOS, so I stockpiled packages of MS-DOS and I now give them away with new servers. That solves the licensing problem without any misunderstandings. There are other MS-DOS-like systems out there, but for maximum compatibility, I prefer using the real thing. Note that many server

manufacturers, especially (though not limited to) mail order houses, explicitly instruct you to pirate MS-DOS onto their NetWare servers, both in their documentation and via phone calls to tech support. Giving away MS-DOS also eliminates the necessity of having my company assist a client in doing that piracy, thereby keeping both organizations out of court.

Next month I'll continue covering safety-related design issues for networks, including detail on the disk options mentioned above. Meanwhile, thinking about smoking computers made me remember an emergency phone call I made to my boss in the middle of the night while working as a mainframe computer operator years ago. When he answered the phone, still half asleep, I said, "The good news is that the fire is out!" I don't think he ever forgave me for phrasing it that way.

DESIGNING FOR SAFETY - Part 2

by Wayne M. Krakau - *Chicago Computer Guide*, October 1999

Wayne is president of Krakau Business Computer Systems, Inc., a systems integration firm that is a Novell Gold Authorized and Novell IBM Host Master Authorized Reseller. He has been working with computers for 22 years. He holds an M.B.A. in Marketing and a B.S. in Information Science and was originally educated in Computer and Information Systems Analysis Engineering. He holds CNI, MCNE and CNTE ratings and all five of the original LANDA Certifications. Wayne can be reached for questions, comments, or topic suggestions at (847)298-7695 or via E-mail to wkrakau@krakau-inc.com. For past columns and company information go to http://www.krakau-inc.com.

This is the second in my series covering creating safer networks. I have gotten tired of seeing networks go down and stay down for extended periods because no techie bothered (or knew enough) to tell the business folks how relatively easy it is to design a network to be at least partially fault tolerant. That is, it resists or at least is readily brought back up after some type of system failure.

The key word here is “design” - as opposed to “throw together” or to pick a “canned” system from an overly simplistic menu. The intense primordial urge of the techie (presumably since prehistoric times) to skip all of that planning “nonsense” and just roll up his (or her!) shirt sleeves (or in the old days, pelt sleeves) and start implementing something (or in non-Techie English, doing something) must be resisted. And no, resistance is definitely NOT futile!

Right now, the topic is backing up your system. Realistically, we’re talking about tape as the media of choice for the vast majority of systems. The cost per megabyte stored is absolutely unbeatable, and its reliability is proven. Even for single-user systems, using removable-disks, such as Iomega’s Zip drives, isn’t as cost effective. I’ve found that business people are incredibly cost sensitive about tape cartridges. At anywhere from $20 to $50 per cartridge, having enough cartridges to provide a thorough backup sequence can cost more than the actual tape drive! Using any media with a higher cost per megabyte would encourage a less than adequate backup pattern. Note that there are special cases where high-capacity magneto-optical disk cartridges or other technologies may be practical, but those are rare exceptions.

I recommend that, as a minimum, my clients use at least a three-week rotating backup pattern that uses 15 tapes for those companies that don’t do weekend backups. An additional 3 (or more) tapes can also be used as a separate “fiscal month” (consisting of 3 weeks) rotating backup. If one of the weekly tape groups (5 tapes) is kept offsite, the business can conceivably survive a building-destroying disaster.

Each of these tapes contains a complete backup of the system. This implies that you should have tape drive with the capacity to backup the whole system at once. If necessary, you can use either multiple drives or an autochanger (a jukebox for tapes) so that one nightly backup actually spans across more than one tape. I do NOT recommend manual spanning for any but the smallest single-user system, as the manual effort required to change tapes usually results in the backup process being skipped when it’s not convenient.

An absolute requirement is to manually check the backup results report (usually automatically printed if the backup software is set up appropriately) every day to make sure nothing unusual happened. I tell my clients to immediately fax me that report if there is an unidentified anomaly within it. If the backup is suspect, I want to know about it right away. Blindly running backup jobs and changing tapes as scheduled is a great way to set yourself up for trouble the next time you need to retrieve data from your tapes. I see this done a lot, but, usually don’t find out until after a disaster that requires retrieving data from tapes.

As to the tape drive, itself, I strongly encourage my clients to use an external drive. If the computer has any sort of internal electrical or heat related problems, an internal tape drive could be toasted along with the computer. Then, not only would you have to worry about repairing or replacing the

computer, you would also have to scramble to repair or replace the tape drive. Since tape drive repairs usually take some time, you might end up buying a second drive. Also, if you've had a tape drive for a while, you might not be able to easily find a compatible replacement. This is after all, the computer industry, where a "generation" of products can be defined in months if not weeks.

Now you have to decide on what type of drive to use. Obviously, your current and predicted future capacity requirements may limit your choices. Usually, I sell DDS (Digital Data Storage) 4mm DAT (Digital Audio Tape) drives. In the past, I sold DDS-2 drives which hold put about 8GB on a 4GB tape by using hardware-based compression. Now I sell DDS-3 drives, which put about 24GB on 12GB tapes, also using hardware compression. DDS-3 drives can also read and write DDS-2 tapes in case you want to exchange data with someone who owns an older drive. My favorite brand of drive is Gigatrend (www.gigatrend.com) because of both their support and their use of a diagnostic LCD panel on their drives. Their drives are actually HP, SONY, Quantum and other big-name drives using various technologies with Gigatrend electronics and enclosures.

DDS autochangers are available with 6, 12, and even 24 drives if you either need the extra capacity, or, just want to automate your backups over an extended period. Autochangers are also handy if you want to give various departments or individuals the ability to make backups of, for instance, an accounting system just prior to month-end closing.

For larger capacities, I sell DLT (Digital Linear Tape) drives in 30, 40 and 70GB capacities. DLT autochangers are available, too. These drives are noticeably more expensive than DDS drives, so I don't sell many, even though the cost per megabyte of the tapes is less.

An up and coming tape technology is AIT (Advanced Intelligent Tape) which uses a new AME (Advanced Metal Evaporative) tape material with MIC (Memory in Cassette) technology that stores the tape directory on a chip withing the tape cartridge. I haven't used it yet, but I have heard very nice things about these drives. The come in 50 and 100GB capacities (including compression) and autochangers from 4 to 40 cartridges are available. That's up to 2TB (terabytes) before compression for a 40-cartridge library (the name used for a really big autochanger) with one or more 50GB drives if you're keeping score! Soon, that library will be available with 100GB drives in case you have some really serious offline storage needs.

One technology that hasn't worked well on servers despite a good track record on individual computers is Travan. There's something in the interface that keeps Travan tape drives from keeping up with the new faster servers and more efficient Network Operating Systems. They tend to get out of sync in communicating with the server, essentially freezing the drive. At minimum, it takes a server reboot to revive the drive. Often, it requires multiple, obscure, commands to regain communications with the tape drive, a task well beyond most users.

Another thing to avoid is bizarre proprietary technology. I've lost track of how many new, innovative tape technologies, supported by only a single firm, have come and gone. Does anyone remember the Pereos drive from Datasonix? It was a drive smaller than my hand that used cartridges based on the micro cassettes used in handheld recorders. It stored an amazing 1.25GB compressed and cost only $695! Today, a Pereos drive would make a unique paperweight. With a pair you've got bookends!

I'll continue with my design suggestions next month. Now, inspired by my memories of the Pereos Paperweight, I'm going to look through my hardware and software archives to find some other archeological relics.

DESIGNING FOR SAFETY - Part 3

by Wayne M. Krakau - *Chicago Computer Guide*, November 1999

Wayne is president of Krakau Business Computer Systems, Inc., a systems integration firm that is a Novell Gold Authorized and Novell IBM Host Master Authorized Reseller. He has been working with computers for 23 years. He holds an M.B.A. in Marketing and a B.S. in Information Science and was originally educated in Computer and Information Systems Analysis Engineering. He holds CNI, MCNE and CNTE ratings and all five of the original LANDA Certifications. Wayne can be reached for questions, comments, or topic suggestions at (847)298-7695 or via E-mail to wkrakau@krakau-inc.com. For past columns and company information go to http://www.krakau-inc.com.

Here's the third article in my series on designing safer networks. While most people wouldn't purposely design unsafe ones, lots of folks, including those who should know better, do so through the neglect of basic design issues.

Now, on to a server's disk subsystem (geekspeak for various combinations of disk drives). The first decision is where to put the disk. If you put them inside the computer enclosure, they are vulnerable to electrical and heat problems that can affect the motherboard, power supply, and other components within that enclosure. This means that if the computer itself fails catastrophically, it could take your disk drives with it. Internal drives also make it difficult, if not impossible, to create an easy-to-use twin server as I described earlier in this series.

Finally, being good at making computers doesn't mean that a company has expertise in making fast and efficient disk systems. For most high-performance internal disk systems, you are stuck with the proprietary disk and controller combinations offered by the computer manufacturer. These tend to be nowhere near as fast, efficient, or cost effective as aftermarket systems available from manufacturers who specialize in making high-performance disk systems.

Because of these factors, I recommend using an external disk subsystem for the NOS (Network Operating System) drives of a server. An external system is not subject to the internal electrical and temperature variations of the server enclosure. It can also be moved to a twin server (as defined in the first article in this series) by any civilian (that is non-computer-geek) with only a trivial amount of training.

There are several choices available in safe disk technologies. At the low end, you can start out with mirrored disks. You use two disk drives attached to a single controller. The two drives hold duplicate data, so that if the first one fails, the other can automatically take over. There is a disadvantage here in that the NOS must do all of the detail work of writing to both drives, thereby incurring a performance penalty of anywhere from 5 to 15% as compared to a single drive.

The next level up from simple drive mirroring is an enhancement of mirroring known as disk duplexing. This method uses two controllers as well as two drives, with each drive connected to its own controller. The most obvious advantage is redundancy of controllers, though that advantage has been reduced over time due to the increase in reliability of controllers. The big advantage is that while plain mirroring is saddled with a performance penalty, duplexing brings a performance increase of 5 to 15% over a single disk.

There are two reasons for this performance enhancement. The first is the fact that the two controllers can take over much of the overhead of writing to both drives, so the NOS doesn't have to do it. The other is the NOS can track the relative idleness of each drive and also the relative position of the drive heads in relation to the desired data and can split the read commands between the drives to optimize performance. (At least that's the technique that NetWare uses.)

The ultimate in safe disk subsystems is an external (see above) RAID (Redundant Array of Independent Disks) system, typically using Level 5 RAID. In this system, a minimum of three disks are tied together to hold data. The data is spanned across these disks in such a way as to provide redundant information so that the array can keep running even if one disk fails. In that case, the information that

would have been on the failed disk is recreated on the fly using the redundant data on the other two disks. Note that there is a performance penalty when running with a failed drive due to the extra effort involved in recreating the missing data.

The basic formula for calculating the total available storage in a RAID system is N-1, where N is the total number of disks. Therefore, a three-disk array has the effective capacity of two disks. Similarly, a four-disk array has a useable capacity of three disks, and so on. The redundant information is spread across all of the disks, not placed on a single disk within the array, so the array can tolerate any single drive failure. For an extra measure of safety, you can add additional drives, called hot spares, that automatically activate if one of the live drives fails.

Typically, RAID systems include a hot-swap feature. This means that you can physically disconnect and remove an individual (presumably failed) drive from the array and insert a replacement drive while the array is running. You do, however, have to be careful not to shake the system while swapping drives to because of the potential to destroy more drives while replacing a bad one.

Most traditional RAID systems are in a tower configuration, where multiple drives fit in an enclosure with one or more power supplies. They are available in fixed sizes, with seven drives being the most common. You can also choose between using a hardware RAID controller in the enclosure or one that fits in a slot in the server. Software RAID has become unpopular due to reductions in the cost of hardware-based RAID controllers. I strongly prefer the embedded (inside the RAID system as opposed to inside the file server) RAID controllers, as these work no matter what NOS brand or version runs on the file server. Using an embedded controller eliminates possible compatibility problems and tends to perform faster, too.

The RAID system that I sell most often is the Radion system from Peripheral Technology Group ((www.ptgs.com & www.radionsystems.com). It is a modular system that eliminates the limitations of fixed-size RAID enclosures. Each module contains a disk drive, a power supply, a fan, a handle (for pulling it out), and a hot-swap connector. You start out with a stack of three hot-swappable drives with a four-channel hardware RAID controller in the base. You can add up to 24 more drives, distributed across four stacks, with each stack attached to its own separate channel.

Next month I'll continue covering safety issues. For now, I'll play with my newest toy. No, it isn't the latest electronic gadget, it's a Y2K bug (as in insect) that I got at the local Hallmark card shop. When you drop it (or spank it if you're into that sort of thing) it makes a loud crashing sound. I'm thinking of sending one to each of my clients who is still procrastinating about Y2K upgrades.

DESIGNING FOR SAFETY - Part 4

by Wayne M. Krakau - *Chicago Computer Guide*, December 1999

Wayne is president of Krakau Business Computer Systems, Inc., a systems integration firm that is a Novell Gold Authorized and Novell IBM Host Master Authorized Reseller. He has been working with computers for 23 years. He holds an M.B.A. in Marketing and a B.S. in Information Science and was originally educated in Computer and Information Systems Analysis Engineering. He holds CNI, MCNE and CNTE ratings and all five of the original LANDA Certifications. Wayne can be reached for questions, comments, or topic suggestions at (847)298-7695 or via E-mail to wkrakau@krakau-inc.com. For past columns and company information go to http://www.krakau-inc.com.

It's time to close out the coverage of designing safe systems. This is the fourth of the series. The goal is to inspire the creation of safer networks with a corresponding lessening of critical errors and subsequent down time.

A safety factor that is becoming more prominent over the years is power infrastructure. That is, we should consider not just the quality of the juice coming from the power company (or in this area frequently NOT coming from Com Ed), but also the quality of the power bouncing around within our buildings.

As faster and more sophisticated CPUs and support chips are created, they dump more and more garbage (electrical irregularities) onto the electrical wiring and in turn into other devices including computers and other network devices. Non-computer electric devices such as flourescent lighting, copiers, refrigerators, etc. also pollute the electrical system within a building.

Standards violations within the wiring also contribute to electrical problems. A lot of the infractions are so obvious that even I can find them, but when you talk to power experts, their horror stories are positively mind-bending. Considering the number of (electrical) safety-related violations that they encounter on a regular basis, I'm quite surprised that we haven't started recalling whole buildings the same way we recall cars with faulty wiring. Perhaps that will happen after enough people get zapped or barbecued.

I have found that, increasingly, the underlying electrical wiring is so bad the even a high-quality UPS (Uninterruptible Power Supply) isn't enough to keep a file server or other network device running. A plague of mysterious glitches and gremlins is usually the first sign. The lockups and ABENDS (Abnormal End, an old term passed on from mainframes) are truly random from a computer processing point of view, though if you expand your thinking to include the possibility of power problems, you might notice some causal relationship with a power system-related event. These events could include motors cycling, air-conditioning systems starting, or my favorite, everybody switching on their PCs almost all at once the first thing in the morning.

Initially you usually test and/or replace various components within the server, trying to pin down the source of these non-repeatable errors. You might even swap out the UPS itself. Eventually you move on to manipulating electrical devices yourself. If you're lucky, you will find that you can affect a computer's reliability be turning on or adjusting things like fans, lights, photocopiers. You might even have neighboring offices or even buildings cycle their major machinery, looking for a clue.

Rewiring is the most obvious answer, but that is often not financially feasible, or not even available as an option because the client has leased space and doesn't have the authority to alter the electrical system.

I finally had to start adding a power conditioner to the UPS, or, if there wasn't a UPS there in the first place, adding a combined conditioner/UPS. This gives near-perfect power to the server in question, but it does add additional noise to the common electrical system, so it might aggravate problems that other network devices are experiencing. In this case you might have to put power conditioners on all network devices.

If the power infrastructure is in really bad shape, or the network spreads out across multiple electrical areas within a building or even to other buildings, grounding may be an issue. Computers can tell only tell the ones from zeros on the network cable if they have approximately the same idea as to what constitutes a one or a zero. The baseline is the electrical ground. If the ground is "bad" or simply different, communications can randomly or even completely fail.

Again, assuming that rewiring is not an option, you can use a device called a Ground Guard (PowerVar, www.powervar.com). It is available as an add-on device for an existing power conditioner or as an integrated unit built into a power conditioner. It will provide a rock-solid ground for the devices plugged into it. Its single disadvantage is that this new ground will almost certainly not match any other ground, thereby eliminating any chance of communications between protected (by a GroundGuard) and unprotected equipment. The only sure cure is to put a GroundGuard and power conditioner (combined or separate) on every device connected to the network. It isn't cheap, but it is often a lot less expensive that either gutting the electrical wiring system and starting over, or putting up with an unstable network.

Naturally, if you are protecting the computer from electrical problems coming over the power lines, you should also protect it from power coming over the network cable. Remember that protection from both dangerous overvoltages and potentially disrupting noise operates on the "weakest link" theory. Miss one piece of the chain and you get nailed.

The final terminus of LAN safety (returning to the use of "safety" in relation to network reliability) is the workstation itself. Having recently covered some of the details of making a Windows 9x workstation more reliable, I will summarize my past suggestions.

First, keep the system patched and the software up to date.

Second, use your favorite utility software to solve and even prevent problems.

Third, think seriously about locking either some are even the whole user interface with either Policy Editor (included with Windows 9x) or Novell's ZenWorks Starter Kit (included with NetWare). Specifically, I have found that Windows 98 and especially Windows 98SE (Second Edition) are particularly sensitive to cris-crossing changes made to the Display and to the Desktop Themes. It seems like these two programs within the Control Panel were written by different programming teams at different times, with conflicting goals. Changing your Windows environment using both programs makes Windows very unhappy. It demonstrates this by adding a new level of instability that will drive the user (and the system integrator) nuts.

I have been requesting that clients try to encourage their users to avoid using Desktop Themes until such time as Microsoft issues a fix. Display is an older program that is less likely to screw things up if it is used alone.

Again, the best way is to simply standardize on an interface and lock it in, but I have been all but threatened with severe bodily harm when I suggest this option to users, as they seem to feel that their civil rights are being violated if they can't fiddle with the interface. Until somebody figures out a way to charge my fees for repairing a workstation (after too much "fiddling") back to a user's salary, I don't think the option of protecting the interface will catch on. After all, even Windows 3.x had a limited version of this problem (troubles caused by inappropriate customizing of the interface), and hardly anybody other than schools ever locked that interface.

Now, I'm waiting to see the effect that the upcoming Y2K-related made-for-TV movie will have on the public's perception of computing. I have heard that it is incredibly inane with regards to ignoring facts in favor of overdramatizing. Already at least two major motion pictures on the same subject have been cancelled, presumably due to the fact that movie studios are owned by major international conglomerates who might be adversely financially affected by the panic potentially generated by these movies.

Oh well, I still haven't been able to overcome the myths about computing that were initially spread by the 1983 movie *War Games*! I'm particularly offended by the stereotype of the overweight, bespeckled, face-fur equipped guy spending all of his time on a computer, to the exclusion of any social

life. Where do they get such silly ideas? Now, let me clean up these candy wrappers and set aside my glasses so I can trim my beard.

AFTERMATH

by Wayne M. Krakau - *Chicago Computer Guide*, February, 2000

Wayne is president of Krakau Business Computer Systems, Inc., a systems integration firm that is a Novell Gold Authorized and Novell IBM Host Master Authorized Reseller. He has been working with computers for 23 years. He holds an M.B.A. in Marketing and a B.S. in Information Science and was originally educated in Computer and Information Systems Analysis Engineering. He holds CNI, MCNE and CNTE ratings and all five of the original LANDA Certifications. Wayne can be reached for questions, comments, or topic suggestions at (847)298-7695 or via E-mail to wkrakau@krakau-inc.com. For past columns and company information go to http://www.krakau-inc.com.

Okay. It's over. Year 2000 is here. The survivors can come out of their bunkers in northern Idaho and western Montana now, presumably to enjoy the post apocalyptical glow. However, considering the huge supply of dried beans that were in their larders (as reported during several news broadcasts), I suspect the insides of their bunkers had much more of an Armageddon-like fire and brimstone odor than the outsides.

Of course, if you want to be picky, the Millennium won't occur until 2001. (Arthur C. Clarke didn't put "2001" in his book and movie title just for the hell of it!) And, if you want to be really picky, these numbers are off by three years due to a calculation error by a sixteenth century monk who must have gotten into the sacramental wine. That would mean that the REAL Millennium was two years ago.

After an incredibly intense few weeks of catching up on major Y2K-related system overhauls for clients who waited until almost the last minute, I encountered only two problems.

The first was that one of the Y2K fix utilities that was supposed to protect some older machines simply didn't work. This was only a minor inconvenience, as the solution was to correct the date manually only once in the new year. After that, the computers ran just fine.

The second was a three-hour blackout several days into the new year. Since having random three-hour or longer blackouts is a normal occurrence for Com Ed customers, I would judge this incident to not be related to Y2K. Oh, and in case you think that some future competition among electrical power companies might solve this problem, just remember that the weakness is in Com Ed's power lines and transmission infrastructure, the same lines and infrastructure that any competitor would also have to use.

On a more global note, we've had our spy satellites start speaking in tongues for the first three days of the new near, just in time to miss covering the expected period of increased terrorist activity. In a classic case of bureaucratic obfuscation, officials claimed that the images that were lost during the malfunction contained nothing critical. Let's see now. Pictures that they haven't seen are judged as unimportant? I guess those ESP experiments at the Pentagon must have gone better than I thought. Of course, the ultimate irony is that it was a Y2K patch program that destabilized the satellites' communications. Oops!

On another front, a bunch of slot machines went nuts. Sadly, for the patrons of the casino, they didn't spew money. They simply didn't work.

There was at least one report of a temporary boost in wealth when a mutual fund listed a shareholder's share value inaccurately by two decimal places. A $26,000 investment suddenly became a $2,600,000 windfall! Alas, it only lasted for a few minutes, and the newly minted millionaire didn't have time to transfer the funds to the Caymans.

Thousands of people all over the country were hit with multiple duplicate charges against their credit cards. The credit card companies were successful in updating and testing their own computers, but they couldn't force the individual merchants that are their clients to uniformly upgrade their point-of-sale credit terminals. Merchants who didn't download and install the appropriate Y2K patch were (and in some cases, still are) inadvertently entering multiple transactions for their customers. Luckily, the extra

transactions were detected by credit card companies' the standard duplicate-detecting programs. So only a few customers actually got overcharged.

In an unrelated incident, more than ten thousand credit card merchant terminals in and around London would not accept either credit cards or ATM cards.

A Y2K voice-mail glitch occurred at Staples, even though the company had paid thousands of dollars in early 1999 to allegedly eradicate such errors. Their DOS-based voice-mail system was deleting any messages with a date prior to 12/31/1999 because it "thought" they were too old, presumably by 100 years. It was also producing annoying warning messages about these "old" messages.

These are just a few examples of the reported troubles. Rather than feeling fooled in some way by the dire predictions of various industry prognosticators, we should probably be thankful that they had what it took to scare people into testing and fixing their systems in time. Just as I have noticed that it takes a couple of gun shots in the area to get people to attend Neighborhood Watch meetings, it seems to have taken some VERY worst-case forecasting to get people to prepare for Y2K.

The really scary part of this story is that many of the program fixes were of the "quick and dirty" variety that I outlined in a previous Y2K article. In a technique called windowing, you can add a temporary, time-limited patch to a program to get it past a critical date. I first saw it when fixing programs to handle the 1980 decade change. Some programs written in the sixties already had a windowing type patch top handle 1970!

Now, we have to worry about programs whose windowing-style Y2K patches will expire on some random, probably undocumented date in the future. It was faster and cheaper to put in this type of temporary patch than to analyze and fix the programs properly. Obviously, many computer management and programming folks haven't learned their lesson from the Y2K crisis. One would hope that corporate management would have caught on by now and demanded proof that permanent fixes were being installed, but, realistically, the "this quarter's bottom line is all that matters" attitude seems to mask all future-oriented problems.

Now for a quick quiz. What's the one thing that is worse than being called in late October and being asked to either rewrite a critical program from scratch or completely replace an entire LAN (or both) before January 1? The answer is being told on a Tuesday that you have to do one (or both) of the aforementioned tasks by the next Monday in order to keep a client from going out of business. (Is that your FINAL answer?)

Well that's just what is going to happen. Many small and medium sized companies (and possibly a few large ones) decided that just crossing their fingers and hoping was the appropriate method to prepare for Y2K. Maybe their systems will fail and maybe they won't. (What I call "The Peter-Paul Principle" - Sometimes you feel like a nut, sometimes you don't!) If they do, it could be at any time. If it's a major failure, they could be out of business quite rapidly. Only time will tell.

Meanwhile, I'll be selling burnt offering kits, complete with live goats, to assist those business owners and managers in their pleadings to the gods of computing.

HIGH HOPES

by Wayne M. Krakau - *Chicago Computer Guide*, March, 2000

Wayne is president of Krakau Business Computer Systems, Inc., a systems integration firm that is a Novell Gold Authorized and Novell IBM Host Master Authorized Reseller. He has been working with computers for 23 years. He holds an M.B.A. in Marketing and a B.S. in Information Science and was originally educated in Computer and Information Systems Analysis Engineering. He holds CNI, MCNE and CNTE ratings and all five of the original LANDA Certifications. Wayne can be reached for questions, comments, or topic suggestions at (847)298-7695 or via E-mail to wkrakau@krakau-inc.com. For past columns and company information go to http://www.krakau-inc.com.

The phrase "high hopes" pretty much describes my attitude toward Microsoft's recent introduction of the production version of Windows. There is a lot of pent up need for both new and upgraded systems out there and, on a very selfish note, I want some of that business.

On the other hand, there's Microsoft's reliability history, or lack thereof, to consider. Already, a leaked Microsoft memo complains to development teams about the absurdity of having customers pay hundreds of dollars for an operating system with 63,000 known bugs!

In support of the validity of the memo, all of the scheduled interviews between Mary Jo Foley (Smart Reseller magazine, www.zdnet.com/sr), the reporter who broke the story, and various Microsoft personnel were abruptly cancelled due to scheduling difficulties. Didn't anyone learn anything from the WaterGate and MonicaGate scandals? Covering up and stonewalling only make things look worse than they are.

On a more reasonable note, Microsoft claimed that the memo refers to the results of a comparatively unsophisticated, automated scanning program that inherently issues many false warnings. I must admit that this is a believable argument against there being a full 63,000 bugs.

The comparison I would make is to running spell-check and grammar-check routines against the articles that I write. Almost every technical term gets flagged as invalid. All of the maddeningly cutesy made-up company and product names that are so common to the computer industry, also get flagged. Finally, because the grammar-checker "thinks" that my style is too informal, it flags some of my more user-friendly phrasing.

These are, of course, in addition to the real spelling and grammar errors. Since these routines are not infallible, they also miss some errors. I suspect the same is true of the results of Microsoft's scanning program. The memo even refers to a lesser number of "real" known bugs of somewhere in the 21,000 to 28,000 range (depending upon how you interpret the memo), while encouraging the development teams to get that total down to zero bugs by the "next" release.

Huh? Do you want to run that by me, again? Bugs that are, by a rather obvious implication, fixable, are being left in until the "next" release? Is this just to meet some arbitrary release date (which Microsoft claimed it wouldn't do), or is it to provide further motivation to pay the price for a future upgrade (which Microsoft also claimed it doesn't do)? In my past experience in both corporate and freelance programming and programming project management, allowing known bugs to get published was grounds for firing.

I've already taken a major financial hit by being blamed for the bugs in Microsoft's 9x series of products. The typical reaction is "Microsoft couldn't have been so successful with programs that are that buggy." I also get "It can't be that buggy because someone would have already sued them out of existence for that bad a program." I really don't want to go through the same thing with Windows 2000.

Luckily, reports from the field indicate that Windows 2000 is less likely to crash when compared to either NT 4 or Windows 98SE. It seems that most of those outstanding bugs are not being hit on a day-to-day production basis. Still, I would suggest that you seriously consider waiting until at least the

first full service pack is issued before you implement Windows 2000, especially one of the server versions.

Of course, there is a strong temptation to churn my clients. I am using “churn” in the same sense that it is used in the stock trading business. There, it is used to denote stock sales made simply to generate greater transaction fees for the broker. I’m using it here to denote upgrading or replacing a client’s operating system simply to generate additional software, hardware, and consulting income.

That has been a popular method for increasing income used by computer resellers of various types over the last few years. I have lost count of the number of potential clients who have approached me with malfunctioning and/or inefficient systems, who are being pushed by their current reseller to switch operating systems. Simply fixing their problems doesn’t seem to be profitable enough.

I am also approached by those potential clients who have been brainwashed by Microsoft (or other software companies) into switching operating systems just to “keep up” with industry trends. In some cases, upgrading was justified on a business-decision basis, though not necessarily to the product that they originally intended. More often, the “gee-whiz factor” just wasn’t enough justification for an expensive upgrade. Note that this applies to upgrades within the NetWare family of products, too.

While Windows 2000's server versions are still not as good as Novell’s NetWare (even after the bugs are fixed), they hold the promise of being the first Microsoft-based servers that I can ethically recommend as a reasonable alternative. I still won’t recommend switching just for the hell of it. Active Directory, for instance, in its eventual final form, with all of the bells and whistles added some time in the future, won’t be as good as Novell’s NDS is today. (Read the reviews if you don’t believe me.)

Even if you have a special application which must run on Windows 2000, Microsoft (and virtually all application software developers) recommend that you put it on a separate Windows 2000 server all by itself. In that case, you still don’t have to switch any of your current NetWare servers to Windows 2000 (including doubling up on memory and quadrupling processors in order to maintain the same relative performance).

For those already running NT 4, the upgrade decision depends on whether you absolutely have to have one of the new features in Windows 2000. For instance, if you have large, complicated system and managing Domains is driving you nuts, switching to Windows 2000 to get Active Directory might be worthwhile. If you are having extreme reliability problems in a specific situation in which you have successfully tested Windows 2000, then it also might be worth upgrading. Either way, if at all possible, wait for the release of more bug fixes. Remember that this is a Version 1.0 product.

Oh, and here’s a quick wake-up call to those Microsoft apologists who keep harping on the alleged “fact” that you can’t write bug-free programs. This is from someone who spent a large chunk of his career being called in by other departments (when I was an employee) and by programming firms (after I had my own company) to review other people’s designs or to actually manage their projects.

Yes you can write bug-free programs. It just takes the proper attitude by programmers as inspired by their management. It’s a little trick discovered many years ago and known by various names, but most commonly as “structured testing methodology.” It’s a royal pain in the posterior, but it’s worth the effort.

You can even protect your program from being blown by an unknown future program with which it must interface. You just have to validate any data passed to your program in the same way that an accounting program validates an account number. If you enter an account number that doesn’t exist, the accounting program doesn’t blow up, it merely displays an appropriate message, hopefully in readable English as opposed to some wacky code number. Similarly, you must validate data passed to your program. Then, instead of just blowing up, your program can display an appropriate message, too.

Also, if you allow for replaceable, shared subprograms and program libraries (like DLLs and VXDs), then take the time to write a management program to prevent errors. Windows 2000 is supposed to be way better than its predecessors at this task.

Meanwhile, I’m crossing my fingers and hoping that Windows 2000 will be a stable and reliable alternative for my clients because, frankly, I could use the bucks!

HOT GADGET

by Wayne M. Krakau - *Chicago Computer Guide*, April, 2000

Wayne is president of Krakau Business Computer Systems, Inc., a systems integration firm that is a Novell Gold Authorized and Novell IBM Host Master Authorized Reseller. He has been working with computers for 23 years. He holds an M.B.A. in Marketing and a B.S. in Information Science and was originally educated in Computer and Information Systems Analysis Engineering. He holds CNI, MCNE and CNTE ratings and all five of the original LANDA Certifications. Wayne can be reached for questions, comments, or topic suggestions at (847)298-7695 or via E-mail to wkrakau@krakau-inc.com. For past columns and company information go to http://www.krakau-inc.com.

I'm a certified gadget freak. I love it when I find a new, useful gadget. Now, the people at Microtest (www.microtest.com) have delivered a gadget fix that will keep me flying for quite some time. It is their DiscZerver VT, a Virtual CD Tower.

This device is virtual in that it simulates a CD tower in the same way that a virtual reality program simulates the real world (minus the aliens and monsters, of course). From the point of view of an individual workstation, it looks just like a separate file server with an attached CD tower. It is, however, easier to manage and way faster.

Not only can the DiscZerver VT look like a server, but it can look like whatever kind of server your client workstation wants, and, will, if desired, adapt to the appropriate native security system. For workstations running a NetWare client, it can look like a NetWare 3.x or 4.x server tapping into either the bindery or NDS, and talking via either IPX or IP. It can also support Microsoft clients (mimicking an NT server and using Domains), OS/2 clients (as an IBM LAN Server), UNIX or Linux clients (using UNIX NFS), and even Macintosh computers (using AFP over EtherTalk). If you want to really get generic, you can also access the DiscZerver VT's virtual CDs directly from your Web browser. Most of the time you will probably use Map or Net Use commands to access the CDs.

The DiscZerver VT is a little box with a blue faceplate equipped with various status lights. It is obviously based on the size of an external drive enclosure for a single 5.25" hard disk drive. Rack mountable and drive-bay mountable versions are also available.

It contains an IDE hard disk drive for storing 7, 14, 28, or 49 full CDs, depending upon which model you purchase. Since it supports compression, and most CDs don't actually use the full 650MB available, you will probably be able to fit more CDs than advertised. It has RAM for caching and the appropriate CPU and logic chips to run a stripped-down, embedded version of Linux. This means that your current file server won't have to control or cache CDs. In addition, it has a SCSI-2 port in the back so you can attach and control a real CD Tower, a single CD-ROM drive, or a CD burner (to create your own CDs).

The DiscZerver VT comes with four types of workstation software. The first is called ZerverView . (Points lost for overuse of cutesy, contrived names!) It allows monitoring and limited management of the device.

EazyImage allows uploading and optionally compressing images of CDs loaded in the workstation's CD-ROM drive. This is the method that my clients are using to put CD images on the DiscZerver VT. I tend to place servers, especially tiny, easily moved ones, in protected locations that make it very inconvenient to accessing an attached CD-ROM drive.

EazyImage Builder (a separate program with a confusingly similar name) allows you to build a CD image from individual files and directories. This can be used with an attached (to the DiscZerver VT) CD burner to create CDs or to upload a 2GB max pseudo CD image (my term) directly to the DiscZerver.

The main program is actually an internal Web site accessed via your Web browser. I usually teach my clients to get to it via a menu item ("Manage Via Browser") within ZerverView, though you can

create a shortcut within the Favorites section of your browser if you prefer. Using ZerverView first allows you to confirm the status of the DiscZerver VT prior to managing it.

It has the same disadvantage of many real Web sites in that it was obviously created by technical gurus, and its overall ease of use suffers a bit from it. The people with the greatest technical skills aren't always the ones who can design the best interface. Luckily, the interface is good enough that I can, with some effort, train civilians (non-geeks) to do the most common day-to-day management tasks, such as adding and deleting CD images.

This Web-based software is the main configuration and administration tool for the DiscZerver VT. It is here that you can set the IP address manually, set it to get an address from a DHCP server, or set it to be the DHCP server to allocate addresses to other devices on the network.

You can also set it to ignore the built-in security system of the file servers on the network and use its own groups and users that you define for security. If your CD's are essentially self-restricting in that you can't use their databases without the appropriate application software, or, you have a site license for them, you can leave the DiscZerver VT's user security wide open. This is what most of my clients are doing. Naturally, you will want to be a lot pickier about administrator security.

Note that licensing issues are potentially the biggest problem with either the DiscZerver VT or real CD Towers. You have to contact all of your software vendors and find out if shared licenses are available and exactly how they are handled. Vendors are constantly coming up with new, and sometimes confusing, ways to arrange and enforce licenses, so don't assume that you can just load up a CD and share it.

The DiscZerver VT is actually one part of a whole family of existing and yet to be released products based on this little blue-faced box. (That's why I keep using its complete name!) The WebZerver, for instance, is already out. Naturally enough, it's a Web server (duh). Their FileZerver will be available soon. I'll let you guess what that does. (Is that your final answer?) Actually it's a bit more complicated than you expect. As well as its obvious use, its also being marketed as an alternative SAN (Storage Area Network) solution. They're even coming out with a LinuxZerver which, naturally enough will run Linux applications.

Now all I need is a BugZerver that would automatically detect bugs in servers and workstations, download patches and fixes, and install them. Any news on that product, Microtest?

SHARE AND SHARE ALIKE

by Wayne M. Krakau - *Chicago Computer Guide*, May, 2000

Wayne is president of Krakau Business Computer Systems, Inc., a systems integration firm that is a Novell Gold Authorized and Novell IBM Host Master Authorized Reseller. He has been working with computers for 23 years. He holds an M.B.A. in Marketing and a B.S. in Information Science and was originally educated in Computer and Information Systems Analysis Engineering. He holds CNI, MCNE and CNTE ratings and all five of the original LANDA Certifications. Wayne can be reached for questions, comments, or topic suggestions at (847)298-7695 or via E-mail to wkrakau@krakau-inc.com. For past columns and company information go to http://www.krakau-inc.com.

Sad, but true, whether we like it or not, clunky old modems attached to POTS (Plain Old Telephone System) lines are often still appropriate, even on a modern LAN. (How about clunky old clichés?) In those cases, following your parents' advice and sharing your toys (modems) may be the best policy.

There are still a lot of companies out there without dedicated internet connections. They are still using dial-up connections. For those with dedicated connections, having one or more dial-ups lines as a backup is quite useful, especially since the ever-more-common DSL lines are still having teething pains.

Remote control applications are frequently used, and, until Virtual Private Networks (VPNs - a connection to a LAN via the Internet) over dedicated lines become more common, require a phone and modem for each connection. Remote control can be used to connect with technical support personnel, for employees calling in from their home computers, or for employees calling in from on the road.

Remote node applications, in which the phone line becomes, in effect, a long extension cord for the network. Your computer acts like it was plugged directly into the LAN cabling system, though typically a lot slower. Again, an individual phone line (probably an ISDN line to be practical) is needed for each connection unless a VPN over dedicated lines is used.

Using individual modems for these applications is one alternative. However, while modern internal modems are relatively inexpensive, they aren't free. In addition, they tend to be a lot more temperamental than their more expensive external brethren, especially in these days of Plug and Play (or Plug and Pray, as it is often called). This is especially true if you use no-name or off-brand internal modems. External modems cost a little more but are less likely to have problems and are much easier to debug. Whichever you choose, at some point buying individual modems does add up.

Once you've installed the modems, you need a phone line for each one. That involves an up front wiring cost to get a socket next to each computer plus the monthly tariff on the phone line. That can also ad up to serious bucks.

A better solution is to use shared modems controlled by the software that comes free with all of the major Network Operating Systems. To implement these shared modems, you need a multiport serial device along with the appropriate adapter cables. You can also use multi-modem devices, but then you run into the same manageability and debugging problems that you get with internal modems, with the added potential of obsolescence (unless you are absolutely sure that modem standards will never change).

If you use a single, non-redundant server, an internal multiport serial board is appropriate. The smaller boards will have either separate sockets directly on the boards, each with its own cable, or, more often, a fanout cable to attach to modems. Boards with a lot of ports, or at least the potential to expand to handle a lot of ports, use a single cable connected to a little black box, which, in turn, has individual sockets for the modem cables. Those modem cables are frequently custom designed. If expansion is supported, the little black boxes are daisy-chained with custom cables.

Note that you can get "smart" and "dumb" boards. The smart ones have one or more CPUs and are capable of taking almost all of the considerable overhead of serial communications away from the server. The dumb boards have just the bare minimum amount intelligence to communicate with the

modems. They force the server to do almost all of the work. That can make a real dent in the overall performance of the server, so you should stick with the smart boards.

External multiport serial devices are all little (and occasionally not so little) black boxes with a network cable socket and a bunch of serial ports that frequently require custom cables. Being stand-alone devices, they are all essentially self-contained smart boards, controlled by the file server.

While they are more expensive than internal boards, they do have some advantages. The first is that they can be placed anywhere where there is a network socket, so you can put them wherever the phone lines terminate.

The second advantage only applies to those who have either true redundant servers, or who implement something similar to the concept that I call twin servers (interchangeable server and master workstation with external hard disks and tape drives). If you use internal boards with along with these designs, you have to buy two boards, one for each server. Also, when one server fails, you have to reconnect one additional large, and quite awkward, connector, with a bunch of modems dangling, to the reserve server.

However, if you use an external multiport device, that device is set to talk to one particular server, as designated by that server's internal ID. Since that ID is on the external disk drive, the reserve server comes up with the same ID as the primary server. The little black serial box doesn't know the difference. It thinks it's talking to the same server, so it just keeps working.

And now for the computer trivia question of the day. What's the other name for the fanout cable mentioned above? An octopus cable. Just don't serve it over pasta!

SHOWTIME!

by Wayne M. Krakau - *Chicago Computer Guide*, June, 2000

Wayne is president of Krakau Business Computer Systems, Inc., a systems integration firm that is a Novell Gold Authorized and Novell IBM Host Master Authorized Reseller. He has been working with computers for 23 years. He holds an M.B.A. in Marketing and a B.S. in Information Science and was originally educated in Computer and Information Systems Analysis Engineering. He holds CNI, MCNE and CNTE ratings and all five of the original LANDA Certifications. Wayne can be reached for questions, comments, or topic suggestions at (847)298-7695 or via E-mail to wkrakau@krakau-inc.com. For past columns and company information go to http://www.krakau-inc.com.

Yes, it's that time of year, again. The trade show season is upon us, and the first, and biggest, (at least for Chicago) is Spring COMDEX.

As usual, Spring COMDEX was held at McCormick Place, that Chicago wonder whose design begs the question, "Did the architects of this place own stock in shoe companies?" If Chicago is considered a city of neighborhoods, then parking, individual seminar rooms, (not to mention the Press/Media Room) and the main exhibit area for this show seemed to be in several different neighborhoods.

All limping aside, the first impression that I got from the exhibit area was "Where did everybody go?" Except for the fact that the concurrent Linux World show filled out a lot of empty space, the show floor looked very sparse, indeed. Chicago's long time (and deservedly applied) reputation within the computer industry as a bastion of computing conservatism has led to a chicken versus egg situation.

After years of being told by vendors of PCs, LANs, and related products, that their sales are fantastic on the coasts, in the south, and in every other Midwest city except Chicago, I have finally heard that Chicago sales of such products are starting to take off, especially for the more esoteric and difficult to justify products such as document and image management and network telephony. Worshipers of mainframes and minicomputers are finally losing ground. IBM might even have to cancel their plans to dedicate a shrine to their local AS/400 sales people! (That's a joke, but not far from the truth, considering Chicago's spectacular AS/400 sales record.)

Now, computer related companies have to decide on whether to wait for Chicago sales to catch up with the rest of the country's before exhibiting here, or to jump in right away to help kick-start those sales. So far, they seem to have given up on us, as evidenced by Spring COMDEX. Maybe those cows on parade didn't help our image in quite the right way. I do hope that, by next year, these companies catch on to the opportunity that they are missing in this former cow-town and start exhibiting in Chicago.

The second impression I got at Spring COMDEX was a far better one, and may bode well for Chicago's reputation. I attended a press conference, called at the last second by Mayor Daley, covering the new Chicago Education Technology Integration Plan. As a professional cynic (at least that's one of the most common repeatable things that a lot of people call me), I was genuinely shocked by what I observed.

Initially, there was the image of a politician and his associates keeping their act together in the face of a hoard of potentially hostile media in a political press conference, that unique cross between an old-fashioned Revival meeting and an Old West ambush in a box canyon. I'm more used to the far less interesting computer industry press conference. The only thing of interest in those is the anxious wait to see if some overenthusiastic entrepreneur is going to pop a blood vessel during his excited attempt to explain his new breakthrough in the C language programming for toasters or some other terribly exciting technical innovation.

The really amazing thing was the presentation of accurate, well thought out, and well-documented information in a field - computer technology planning - which is rife with opportunities for inaccuracies and outright lies. This, in a field in which I am, nominally, at least, an expert, and could,

in theory, detect most of the aforementioned deviations from fact. (Oops, I'm starting to talk like a politician!)

With some help from personnel of marchFirst (their spelling), who donated more than $750,000 worth of time to this project, the Mayor (Da Mare?) outlined a cooperative venture of the Chicago Board of Education, the Chicago Roman Catholic Archdiocese' Office of Catholic Education and the City Colleges of Chicago to coordinate their efforts in getting computer education up to speed in Chicago.

This cooperation will include the computing infrastructure of both LAN and electrical (AC power) cabling, the acquisition and maintenance of computer hardware and software, and most important, the curriculum needed to effectively use all of this technology. This last item is what impressed me the most. It is quite rare for a politician, especially a self-proclaimed beginner in computing, to fully understand the difference between using computers in education and training future computer professionals. Computers are useful as training devices for basics such as math and reading as well as a myriad of other subjects. They are useful as tools for many academic and business situations (think wordprocessors, spreadsheets, etc.). Finally, they are useful for that small percentage of people with an aptitude and an interest in to becoming computer specialists. The Mayor obviously understood the subtleties of this situation.

Finally, when both his own staff and the marchFirst folks were peppered with questions on the lack of primary emphasis on donations for used computers, Mayor Daley took over the podium. It was obvious that the civilian (non-computer) press wasn't quite getting the idea, so he succinctly summed it up by explaining the hazards of blindly accepting donations of used computers.

An educational organization accepts a computer that was discarded due to obsolescence. The donating company gets a tax deduction and avoids the expense of disposing of hazardous waste. The computer is too old to run the software used at that institution, so it is stashed in either a closet or a warehouse. Some time later, it is discovered by someone else and is finally thrown out. Meanwhile, the educational institution has wasted money (directly or indirectly) on storage, and now has to pay to dispose of the computer as hazardous waste.

The Mayor rightly emphasized that the review of proposed donations by appropriately qualified computer professionals would be required prior to acceptance. This would protect the schools from the scenario that he described. He left this cynic slack-jawed in amazement at the sight of a politician not only telling the truth, but actually knowing what he was talking about! (Oh, and by the way, the civilian press still didn't get it.)

It was also interesting from a personal point of view to recognize a lot of TV news reporters there. Up close and personal, I noticed that one of them - I'm not naming names - looked to be the original inspiration for Waldo of "Where's Waldo" fame. This is not to say that he wasn't reasonably good looking - as a TV reporter must be to survive - it was just that I swear he was a dead ringer for Waldo! This resemblance just doesn't come through via the TV screen. I'll have more on Spring COMDEX next month, after I check the Web for Waldo pictures.

SHOWTIME! - Part 2

by Wayne M. Krakau - *Chicago Computer Guide*, July, 2000

Wayne is president of Krakau Business Computer Systems, Inc., a systems integration firm that is a Novell Gold Authorized and Novell IBM Host Master Authorized Reseller. He has been working with computers for 23 years. He holds an M.B.A. in Marketing and a B.S. in Information Science and was originally educated in Computer and Information Systems Analysis Engineering. He holds CNI, MCNE and CNTE ratings and all five of the original LANDA Certifications. Wayne can be reached for questions, comments, or topic suggestions at (847)298-7695 or via E-mail to wkrakau@krakau-inc.com. For past columns and company information go to http://www.krakau-inc.com.

As the swallows return to San Juan Capistrano, and the buzzards return to Hinkley, local computer pundits (including yours truly) return to Chicago's biggest computer trade show, Spring COMDEX 2000, the main difference being the presence of mass quantities of real avian excrement at the first two locations versus the virtual bovine excrement found at the third.

My first view of the exhibit floor was from the balcony along the west wall where Mayor Daley's press conference (covered in last month's column) had just ended. My first impression was one of amazement as I realized the LinuxWorld booths covered a HUGE chunk of the floor space.

As I watched the initial inrush of the crowd at opening time, I realized that I could easily define the Linux area even without all of the signs. One side of the show floor was obviously the "suited" side and the other was the "scruffy" side, based on my observation of the attendees. After a while, there was some mixing, but the overall trend remained. The Linux half of the show had the distinct feel of the early microcomputer shows, and even, though I am loathe to admit to being old enough to have attended, some not quite so early minicomputer shows. It was definitely running about an 8.5 on the Face-Fur scale, at least for the male attendees.

As I wandered the aisles, I was impressed with the enthusiasm of the Linux crowd, some of it distinctly anti-Microsoft, but mostly pro-Linux. I was also embarrassed to be corrected in my pronunciation of the operating system. The "i" in Linux is supposed to be short, not long. Oops! I hope they won't take away my computer geek ID card.

Another example of the mood of the group was the evidence of exhibits that were throwbacks to a bygone era. A big penguin drifted about. I mean really big - something in the 5'8" to 5'10" range. While I did wonder about the health of the occupant of the penguin suit (Heat stroke anyone?), I must admit that I got caught up in the spirit and had my picture taken hugging this big fuzzy mascot of Linux. Too bad my picture wasn't ready by the time I left. I later saw the empty costume standing up (about three-quarters of the way) in a corner and wondered whether it stood up due some inherent design factor or due to the collective will of some perspiration-induced colony of bacteria.

Later, I encountered a devil mascot. Ironically, the devil suit looked a lot cooler than the penguin suit. One disappointing throwback (in the negative sense of the term) was that the devil was accompanied by a pair of scantily-clad devilettes (for lack of a better term), passing out devil horns. (I'll decline the opportunity to use the obvious pun.) I'm sure that the many female management, sales, and technical professionals both exhibiting at and attending the show did not appreciate their presence. I guess it's two steps forward and one giant step back for their business images.

Separate from all of this silliness, I was also impressed by a number of exhibits inhabited by vendor representatives who had serious answers to serious questions about how to use Linux in a business environment. They offered software and hardware solutions that specifically addressed the business concerns that I presented.

The software solutions involved the use of Linux as a server operating system, as an independent desktop operating system, and as the desktop base for running Windows emulators. The hardware solutions included bundled Linux solutions for servers, desktop systems, and various network appliances.

Network appliances are specialized “black boxes” made to do one specific task normally done by a full-featured server. The most common examples exhibited in this category were routers.

I was already a fan of Linux-kernel-based network appliances, as evidenced by my recent favorable article on a pseudo CD tower, but, prior to this show, I didn’t have a lot of interest in the other Linux solutions. The serious portion of LinuxWorld has opened my mind up to the possibilities of Linux. I’m not a true convert, but I am noticeably less of a skeptic than I was. Perhaps the “congregation’s” enthusiasm was contagious.

Do not underestimate the importance of the collective enthusiasm of this coiffure- and couture-impaired Linux crowd. This same force of nature conquered overcame the inertia of the corporate computing world to bring PCs and LANs to their current prominence. Bill Gates, be afraid!

SHOWTIME! - Part 3

by Wayne M. Krakau - *Chicago Computer Guide*, August, 2000

Wayne is president of Krakau Business Computer Systems, Inc., a systems integration firm that is a Novell Gold Authorized and Novell IBM Host Master Authorized Reseller. He has been working with computers for 23 years. He holds an M.B.A. in Marketing and a B.S. in Information Science and was originally educated in Computer and Information Systems Analysis Engineering. He holds CNI, MCNE and CNTE ratings and all five of the original LANDA Certifications. Wayne can be reached for questions, comments, or topic suggestions at (847)298-7695 or via E-mail to wkrakau@krakau-inc.com. For past columns and company information go to http://www.krakau-inc.com.

It's back to the trade shows, again. This time I'm covering the remainder of Spring COMDEX 2000 plus the recent Chicago Internet World 2000.

The major buzz in the Linux half of COMDEX was that Corel had the best Linux package. This is an observation on my part, not a technical opinion. There was also an underlying worry about the rumors, now confirmed, about financial problems at Corel. For the hobbyists and enthusiasts in the crowd, the idea of committing to a Linux vendor now, only to be forced to switch later was annoying, but not really that daunting.

Business users and their technical staff colleagues, however, have a lot more worries - like their jobs. For smaller businesses or departmental managers, the lack of in-house technical support could make a major change very expensive. For the corporate folks, whose LAN/PC support teams are typically too understaffed to handle the day-to-day running of the network, a corporate-wide conversion could be just as impractical. Hey, who ever said that business computing was easy? Besides, there seems to be a rather large personnel shortage in the food service industry. Would you like fries with that?

The most impressive bit of hardware technology at COMDEX was also the smallest. It's called the ThumbDrive, from Trek 2000 International, Ltd. (www.thumbdrive.com). The idea for this device is so clever and so handy that I'm surprised that nobody else has thought of it before. It consists of a small DIP-style (Dual Inline Package) chip embedded in a plastic housing with a male USB (Universal Serial Bus) connector sticking out one end. That connector is plugged into a USB socket on any recently manufactured desktop or laptop computer, either directly, or via a USB hub.

Once you've stuck the ThumbDrive into a USB port, what have you got? You've got an ultra high speed "disk drive" holding from 8MB to 256MB of data (with 512MB available soon), depending on the model, in a finger-sized device that is practically indestructible, based on normal office and travel handling. In principle, it should be impervious to most low and medium level electromagnetic interference. (Just, don't try to use it in place of a spark plug!) In terms of physical durability, you would probably have to stomp on one to break it.

All you need to make the ThumbDrive work is Windows 9x and the small device driver provided with the product (and also downloadable). After that, anytime you plug in the device, it immediately shows up as an additional drive within Windows Explorer. Unplug it and the drive letter disappears.

The ThumbDrive is made specifically for data transfer. Since each one cost more than an entire Zip Drive ($399 for 128MB), they are not practical for general purpose storage. However, their simplicity and durability make them ideal for transferring files between computers.

I have to admit that much of the allure of the ThumbDrive is its geek-appeal. For a gadget-freak like me, a tiny, durable, hard drive-equivalent, with a catchy name, and literally the size of a thumb, is almost irresistible. Besides, you can plug 16, 128, 265, 64 and 32MB ThumbDrives, respectively, into a five-port USB hub and pretend you are in the Outer Limits episode (the original series, not the current one) about the fugitive with the computer hand with removable fingers!

Well, how about Internet World 2000? My general opinion can be summarized in one word - yawn. How many times can you see a booth claiming to be the best and (naturally) most technologically

advanced method of [Fill in the blank: e-commerce site hosting; e-commerce software; e-commerce consulting; e-commerce payment software; e-commerce payment services; Web site hosting; Web site design; etc.].

To be honest, the fact that all of these solutions looked identical (at least within their own categories) may be more a matter of weak marketing skills than actual quality. It seems that marketing folks, unlike computer types, aren't that willing to work themselves nearly to death for a combination of minimum wages (if that) and the promise of (currently worthless) stock.

What's worse is that many booths had such generic displays that I couldn't figure out what they did! I know that cars have been marketed with commercials showing nothing more substantial than scenes of rolling hills without actually displaying a car, but a computer trade show is not the place to get so esoteric that the attendees aren't motivated to investigate your exhibit.

The only truly innovative product that I saw was a collection of wearable computing devices by Charmed Technology (www.charmed.com). Innovative - yes. Useful for any of my clients - no. Useful at a high-tech pirate-themed costume party - possibly, due to the eyepatch-looking monitors their models were wearing.

My final impression of the show was inspired by the three models at the Charmed Technology booth. They kept getting together and whispering among themselves. In my warped imagination the conversation went like this:

First Model: Wow! Get a load of that geek over there.

Second Model: Yeah. Isn't he just the worst nerd in this kingdom of nerds?

Third Model: Hey! I recognize him. He's one of those Internet millionaires.

First Model: He's mine - I saw him first!

UTILITY INFIELDER

by Wayne M. Krakau - *Chicago Computer Guide*, September, 2000

Wayne is president of Krakau Business Computer Systems, Inc., a systems integration firm that is a Novell Gold Authorized and Novell IBM Host Master Authorized Reseller. He has been working with computers for 23 years. He holds an M.B.A. in Marketing and a B.S. in Information Science and was originally educated in Computer and Information Systems Analysis Engineering. He holds CNI, MCNE and CNTE ratings and all five of the original LANDA Certifications. Wayne can be reached for questions, comments, or topic suggestions at (847)298-7695 or via E-mail to wkrakau@krakau-inc.com. For past columns and company information go to http://www.krakau-inc.com.

This month I'll be covering some interesting, and quite handy, utility programs that are of interest to both system administrators and some of the more sophisticated end-users. (I've always wondered whether the term "end-user" was purposely invented by the original mainframe geeks to allow for obvious, derogatory puns.)

The first handy utility is Screen Shot Deluxe from Parson's Technology (www.parsontech.com), now a part of The Learning Company (www.shoptlc.com), which in turn is now a part of Mattel Interactive (www.mattelinteractive.com). Whew, what a mouthful! Did you get all that?

As a side note, I have dealt with products that have been through as many as five corporate acquisitions within the last ten years. I am convinced that, at some point in the not-too-distant future, all software products will be owned by Microsoft and all hardware products, excluding actual computers, will be owned by Cisco. Then, large companies will refrain from competing with these giants due to non-compete agreements (signed for their own survival). From that time on, hit teams will be dispatched upon the emergence of any small, independent, startup companies in either category. It will be some like some weird mix of *Fahrenheit 451* and *The Running Man*. It will probably take a few more years to get to the style of living presented in *The Matrix*, but I am already somewhat suspicious of the growing, socket-like dent in the back of my head.

Despite its mixed heritage (or perhaps because of it), Screen Shot Deluxe turns out to be quite an effective utility. Its basic, underlying task matches its name - it does screen shots. That is, it takes a graphical snapshot of the screen and then prints it, just like the old Print Screen key did in the old DOS days.

In Windows, the Print Screen key is limited to dumping the resulting picture to the Windows clipboard. Then you have to manually paste the picture into another program before you can print it. That other program is also required if you decide to manipulate, annotate, or modify the picture in any way. Since the clipboard is designed to handle only one item at a time, you end up switching between programs a lot if you need multiple screen shots.

Screen Shot Deluxe, however, automates the printing and adds a lot of other options. Using the default key assignments, pressing the Print Screen key by itself produces a printout of the whole screen. Pressing the Alt-Print Screen combination prints just the active window, something that is frequently much more useful.

One major option is the destination of the captured image. You can automatically print the image, save it to the clipboard, save it to a file, save it directly into an image-editing program, or just pull it straight into Screen Shot's own editing screen.

Another is the shape of the capture. In addition to the whole screen and active window, you can grab a designated region, the object under the mouse pointer, or even a geometric shape. You can also do timed screen shots, which can be used to see messages that are going by so fast that you normally can't read them.

You can use Screen Shot's own editing for most common functions. That includes both text and image annotation, various types of size manipulation, and brightness/contrast adjustment.

I initially started experimenting with Print Screen replacement programs in order to create documentation but have since found that they can be used for some not-so-obvious tasks. For instance, the timed capture function is useful in debugging. Also, many programs have screens which have present precisely the information you need in an easily readable format, but lack any simple way of printing that information. The answer is to use Screen Shot to dump the image to your printer. In addition, when you get a Web page that simply won't print correctly no matter what you do, just nail the Print Screen key and Screen Shot will grab the screen for you, exactly as formatted.

Finally, Screen Shot Deluxe is the only screen capturing program that I've found that doesn't fight with other programs or spontaneously blow up.

The other program that I'm recommending (conditionally, in this case) is Go!Zilla, a download management program by Aureate Media (www.gozilla.com), which is now a part of Radiate (www.radiate.com). (Can't tell the players without a scorecard!) Go!Zilla is the first download manager that I have used that both works as advertised, and doesn't fight with other programs

The primary reason that I started using download managers was to continue interrupted downloads. Go!Zilla (Wow, is that a pain to type!) automatically resumes downloads with a minimum amount of overlap of new versus previously downloaded data. In addition, just prior to starting or resuming a download, it searches out the most efficient source for the file, taking into account how busy the various available Web and FTP sites are. This means you will always get the most out of your Internet connection.

Normally, Go!Zilla automatically intercepts download requests made by clicking on a file name on a Web or FTP page. If the page is itself automated in such a way as to prevent autodetection by Go!Zilla, you can right click on the file name and copy the shortcut or link, depending upon whether you are using Internet Explorer or Netscape Communicator. If Go!Zilla (Did they make it purposely hard to type?) is already active, it will grab the file name immediately. If it's not, you can activate it and paste the file name into the Add File window.

Once Go!Zilla has a file name, you can optionally redirect the file so a different, non-default directory, and then decide on whether to download now or later. If you want to download later, you can manually start the download process of one or more files, or schedule downloads for specific dates and times. You can even select download cutoff times to avoid time-based phone or link rate hikes.

Because Web sites are not perfect, and file management is always an issue when you do lots of downloads, Go!Zilla can be valuable even to those who have high speed links to the Internet. \

I have had as many as twenty Go!Zilla downloads going at once on an unattended PC using a dial-up link. When I went back to the PC, I could still use other programs without appreciably slowing or worse, blowing up the download sessions. Go!Zilla (After this article, I'm never going to type that word again!) Never even blinked when Internet Explorer crashed (Gee, what a surprise.) and subsequently took several other programs down with it.

Now comes the catch. The reason I'm making this recommendation conditional is the Go!Zilla uses Radiate's technology to pass information on what you are downloading back up to Radiate's database in a supposedly anonymous format. The anonymity of that format is currently in dispute. While nobody has yet caught anyone using this type of technology doing anything nefarious with the collected data, the potential for future misuse of sensitive data is there. In the future, I will be writing more about this trend of using what's being called SpyWare.

Just for your information, uploading a file that was previously downloaded doesn't erase it from a SpyWare database any more than running the car in reverse in *Ferris Bueller's Day Off* cranked the odometer backward. I just thought I'd warn you.

I SPY

by Wayne M. Krakau - *Chicago Computer Guide*, October, 2000

Wayne is president of Krakau Business Computer Systems, Inc., a systems integration firm that is a Novell Gold Authorized and Novell IBM Host Master Authorized Reseller. He has been working with computers for 23 years. He holds an M.B.A. in Marketing and a B.S. in Information Science and was originally educated in Computer and Information Systems Analysis Engineering. He holds CNI, MCNE and CNTE ratings and all five of the original LANDA Certifications. Wayne can be reached for questions, comments, or topic suggestions at (847)298-7695 or via E-mail to wkrakau@krakau-inc.com. For past columns and company information go to http://www.krakau-inc.com.

Just as a pair of spies masquerading as a professional tennis player and his coach could once travel the world, reporting all sorts of valuable information back to their superiors without engendering suspicion (at least on TV), various types of seemingly innocuous utility software are now infiltrating computers around the world, surreptitiously spying on the users' activities, and reporting back to their creators.

In my recent review of Go!Zilla (I thought I was finished trying to type this!), I was somewhat apprehensive about recommending the program due to its use of what's being called SpyWare. That's becoming the common name for software that gathers information about the PC it runs on and sends that information back to some central database. Apple thought that IBM was Big Brother. Little did they know that small, innovative software companies would take on that role.

The main proponent of caution in using SpyWare is Steve Gibson (of Gibson Research Corporation, www.grc.com), formerly a columnist for InfoWorld, but perhaps better known as the author of SpinRite, the amazing disk diagnosis and repair program. He has started a Don Quixote-like quest to track down the culprits and to absolutely prove his case.

Gibson is a true craftsman of the old school. His main skill is in writing carefully handcrafted Assembler-language-based utility programs. He is so dedicated to his work that he gave up his column and miscellaneous writing in order to spend more time programming.

My only in-person encounter with Gibson happened a few years ago when I had the privilege of attending a seminar on SpinRite at a distributor show for resellers a few years ago. In explaining his product, Gibson animatedly acted out the parts of the various components of a disk drive. It was both hilarious and enlightening. This guy has an exhaustive knowledge of what makes a PC tick.

Gibson's SpyWare quest started when he heard rumors about the RealDownload download manager from RealNetworks and it's derivatives, Netscape's Smart Download and NetZip's Download Demon. As usual, he wrote his own utility programs as necessary to use in this investigation in addition to various sophisticated diagnostic tools. His account of what followed reads more like a detective novel than a technical treatise.

First, he discovered that the full addresses of all files downloaded using these managers were going out over his Internet connection. In what would become a repeating pattern, his published results were immediately met with rebuke and denial by the software companies. They portrayed him as a poor, misguided soul, tilting at windmills.

Rather than either backing down or engaging in some type of reflexive trading of denials (Yes they do - No they don't - Yes, they do, etc.), Gibson reexamined and reran his own research. He didn't want to accidentally accuse an innocent party. (We are talking about a serious ethical decision here. Bill Gates, are you listening?)

His additional research uncovered the fact that, along with the downloaded file information, a unique identifying field was also transmitted. On computers with an Ethernet card, half of the ID was the card's address while the other half was a part of Windows. On computers without an Ethernet card, that half of the ID was generated by some as yet unrevealed formula built into Windows. Now the

software companies could identify his machine along with his activity, though they couldn't yet tie him personally to his machine.

Even more research, paralleled by further denials and some threats of legal action, followed. Gibson discovered some suspicious activity during Web-based purchases from the associated software companies. He tracked that activity down and found that the same unique-to-his-PC ID was uploaded along with his personal information. This was then encoded and saved on his PC as a "cookie" so that every time he downloaded a file, critical portions of his personal information, including his name and his private (for online transactions only) e-mail address, were uploaded. So much for online privacy statements!

To this day, RealNetworks and their associated software partners deny that they are any threat to users privacy. This is in spite of the fact that Gibson has thoroughly documented and published all of his research in a format that is readily understood. (Do they have a corporate position entitled "Propaganda Minister" or something similar?) Just go to www.grc.com if you want to examine it yourself.

While you are there, you may also want to check out his free program, OptOut which allows you to detect and, optionally, remove SpyWare from your system. He also has written various other free and commercial programs (in addition to the aforementioned SpinRite), including a handy Zip Drive diagnostic tool and an Internet connection security tester.

I know that I'll be the first in line when Mr. Gibson announces the world tour of his one-man show in which he acts out the parts of the software and hardware involved in his battle against SpyWare.

REVENGE OF THE NERDS

by Wayne M. Krakau - *Chicago Computer Guide*, December, 2000

Wayne is president of Krakau Business Computer Systems, Inc., a systems integration firm that is a Novell Gold Authorized and Novell IBM Host Master Authorized Reseller. He has been working with computers for 24 years. He holds an M.B.A. in Marketing and a B.S. in Information Science and was originally educated in Computer and Information Systems Analysis Engineering. He holds CNI, MCNE and CNTE ratings and all five of the original LANDA Certifications. Wayne can be reached for questions, comments, or topic suggestions at (630)894-3360 or via E-mail to wkrakau@krakau-inc.com. For past columns and company information go to http://www.krakau-inc.com.

Call it juvenile. Call it downright stupid. Definitely call it politically incorrect. Or, do as I do and call it a guilty pleasure. I'm talking about the movie *Revenge of the Nerds*. I saw this movie when it was originally released in 1984 and have seen in a piecemeal fashion on various cable and broadcast stations many times since.

For what may be obvious reasons, I have always identified with some of the characters in this movie, especially in feeling "different" from the people in the various social cliques. I specifically remembered being stared at as I walked down the halls during a high school fire drill. The supposed affront to public decency? I was playing "blindfold" chess with another chess team member whom I happened to run into during the drill, though I suppose that you could call it "double blindfold," since neither one of us had access to a board. (I won by checkmate on the seventeenth move.)

Another example of my "outcast" status is my almost complete disinterest in competitive sports. I suspect, though that may come more from being brought up during the leanest sports years in Chicago history. The Bears and Bulls victories didn't come until I was well past my formative years. My sports motto was, and in the case of Chicago baseball, still is "Wait till next year."

I will admit to watching most of the Bears games in 1985 and some of the Bulls (or, more specifically, Michael Jordan's) games during their multiple winning seasons, but I haven't watched a Cubs game since I was a kid and I've never watched a Sox game. (Hey, I'm a North Sider. So shoot me.) However I did, of course, closely follow Bobby Fisher's chess championship match in Reykjavik, Iceland. (Like you couldn't have guessed that.)

Ah, but what's the point of this extended treatise on nerdism (if there is such a word)? The point is that I have now lived out my favorite part of the movie, *Revenge of the Nerds*. That's the part where the chief nerd ends up with the best looking cheerleader. Also, in the movie, there was a subtext that implied (at least in my opinion) that the cheerleader was smart enough that, except for the disqualifying factors of having good looks and social skills, might have qualified as a nerd herself. (All right. I'm stretching a bit on that one.) In the sequels (which, by the way, are utterly unwatchable), they even end up married.

Well, I just married an amazing woman who, though she never was an actual cheerleader, could have been (and probably should have been)! Moreover, as an added bonus she is genuinely smart enough to be a nerd, if it wasn't for her looks and social skills. Though, I suppose she does have a gap in her intelligence as far as her choice of a husband is concerned.

Her only other weakness is that as a "non-computer" person, she sees computers as potentially useful, though sometimes annoying, tools, not as they really are, holy icons to be worshipped. (The next time your system blows up, check your schedule to see if you missed your weekly burnt offering to it.) I hope to convert her as I have heard a lot about mixed marriages being difficult.

Poll Cat

In these times of multiple polls, many of which are wildly unscientific, I thought that it was my duty to provide some serious statistics on a critically important subject, the current state of Hawaiian computing. Therefore, I undertook a (completely deductible) business trip to Hawaii, taking my

executive assistant (wife) with me to assist in taking the polls. It is only by coincidence that this business trip happened to start on the day after our wedding.

In a careful, scientific study, we polled both locals and visitors at many strategic locations. These included a painstakingly selected sampling of beaches, restaurants, botanical gardens, dive shops (SCUBA), dive boats and luaus. Our results are as follows: 1. Which desktop operation system do you prefer? - Windows 98 0% - Windows 2000 0% - Linux 0% - Don't care 100%. 2. Which network server operating system do you prefer? Windows NT 0% -Windows 2000 0% - NetWare 0% - Linux 0% - Really don't care 100%. 3. What should happen to Microsoft in the antitrust case against it? Slap on the wrist 0% - Split up company 0% - Government supervision 0% - Fines 0% - Wat doing? Nuff already! Like beef? Try move, lolo haole! (Translation: What are you doing? Enough, already! Would you like to fight? Please move out of the way, you stupid Caucasian person!) - 100%. I'm sure that publishing these findings will justify our deductions with the IRS.

Wait Till Next Year

We have been so busy moving from Des Plaines to Medinah that we missed our usual Halloween party. (See the new phone number listed above and check our Web site for the new address.) This will give me time to find our *Revenge of the Nerds* tribute costumes. I've got to locate a Darth Vader costume for me, a cheerleader costume for my wife, and several fun-house style distorting mirrors. (If you've seen the movie, you'll understand.) Please hold the pompom jokes.

A WIN-WIN SITUATION

by Wayne M. Krakau - *Chicago Computer Guide*, January, 2001

Wayne is president of Krakau Business Computer Systems, a systems integration firm that is a Novell Gold Authorized and Novell IBM Host Master Authorized Reseller. He has been working with computers for 24 years. He holds an M.B.A. in Marketing and a B.S. in Information Science and was originally educated in Computer and Information Systems Analysis Engineering. He holds CNI, MCNE and CNTE ratings and all five of the original LANDA Certifications. Wayne can be reached for questions, comments, or topic suggestions at (630)894-3360 or via E-mail to wkrakau@krakau-inc.com. For past columns and company information go to http://www.krakau-inc.com.

Well, Microsoft has finally done it. They really screwed up this time. They went and released a product, Windows 2000, that, when fully patched and up to date, actually works, pretty much as advertised. This is certain to cause mass layoffs in the computer industry press as an entire niche market, publishing reports of Microsoft's errors and their fixes (if any) collapses. How unpatriotic can Microsoft get?

What's worse is that, without Microsoft handing me subject matter on a silver platter (namely bugs galore), I may have to get off my butt and do some serious research to find subject matter for future columns. Oh, the inhumanity of it all!

After issuing products that were so unreliable that even the most pro-Microsoft writers bashed them in the press, Microsoft has finally issued enough patches to make the Windows 2000's server versions work well enough that I no longer need to refrain from selling Microsoft's server software on either ethical or self-preservation grounds.

Ethics were involved in that I was unwilling to foist what I consider an incredibly buggy and inherently unreliable server operating system, Windows NT, upon them. Yes, we have worked for some time with NT servers in mixed NetWare/NT environments, but I was never the one to actually sell the NT server to the client. I have seen too many NT servers fail repeatedly even though they were set up by people with extreme expertise in NT. Even the national press constantly publishes that it is considered "standard" to constantly have to reboot NT servers due to bug-induced crashes.

I have not been persuaded by arguments of colleagues that I should push NT Server over NetWare because of the money involved in the additional (presumably expensive) hardware needed, the extra installation and setup costs, and, especially, the never-ending stream of maintenance fees involved in keeping them running.

A secondary consideration (well, maybe really primary) is that I don't want to get blamed, or even sued, when the server proves unreliable. Neither my reputation nor my wallet can afford it. I have already had enough trouble with extreme negative feedback, sometimes to the point of nonpayment, by clients unsatisfied by the reliability of the Windows 95/98 family of products.

In light of this, I have been anxiously waiting (and waiting, and waiting . . .) for the day when I can offer Microsoft's Windows 2000 as a reasonable alternative to Novell's NetWare. I fully admit that a major part of my desire to offer Windows 2000's server versions is that, first, Novell has repeatedly proven over the years that, while it made, and still makes, the best software available (despite the fact that I nitpick tem regularly), it has not been able to market its products in anything approaching an effective way. It almost seems like they are trying to hide their products from the press, and, in turn, from prospective customers.

Second, Microsoft has proven so effective in marketing their products that the companies that have traditionally been used as examples of great marketers, Procter & Gamble and IBM, could now be considered rank amateurs in comparison to Microsoft. If Firestone was run by Bill Gates, their sales would be soaring right now, and, in light of the various investigations going on, people would be

protesting the government's interference in Firestone's business. They probably would have had to hire extra "temporary" employees, and would have had to find creative new ways to underpay them.

Perhaps marriage has mellowed (or at least worn) me out, but, now that Microsoft has caught up with the most critical patches and fixes (emphatically NOT including the now-withdrawn mess known as Service Pack 2), I consider Windows 2000's server versions to be both reliable enough and, as importantly, effective enough, for my company to sell. For me, being willing to sell a product is the ultimate complement.

Reliability isn't the only issue when comparing Windows NT to Windows 2000. In fact, I am somewhat amused by the fact that reports are coming out that say that people experienced with Novell's NDS (Novell Directory System), but not with Windows NT, are having an easier time designing, installing and managing Win2K servers than experienced Win NT folks. This sounds quite similar to the turmoil created among NetWare 3.x experts when NetWare 4, with NDS, was released. To Novell's dismay, there are still thousands of 3.x sites out there run or serviced by people who have still not gotten a handle on the concept of directory services.

Now, Microsoft may be in the same boat. They are using various carrot and stick (mostly stick) methods to convince people to abandon NT, and, especially, to redirect their ongoing educational efforts away from NT and toward Win2K. They are caught in the same chicken versus egg cycle that has hampered Novell's products. Which comes first? Will Win2K be purchased by enough customers to attract both software developers and trained technicians? Or will there be enough software and trained technicians available so as to attract potential Win2K customers?

Now, as to advice on what to do, it depends upon your individual situation. If you have NetWare and are not being forced by either management demands or by application developers pressure, then stick with it. You've already got the best. In fact, if your software developer is insisting you switch just to make them happy, then you might want to reevaluate your application software vendor with an emphasis on finding one with a more customer-friendly attitude.

If you already have an NT server, your only restrictions are the availability of Win2K compatible versions of your server-based software and your budget. If your software developer hasn't yet jumped on the Win2K bandwagon, then a reevaluation might be in order. In this case, you would be looking for one that keeps their products up to date. Basically, the more you are impacted by the unreliability of NT, the more urgent your conversion becomes. If your current LAN vendor can't handle the switch, especially because they can't understand the concept or the benefits of directory services, find one who can. Of course, this advice also applies to current owners of NetWare 3.x and below.

Linux is also a choice, especially for those whose personnel have UNIX experience, but I am not yet anxious to unconditionally jump on that bandwagon. I'm waiting for a shakeout of Linux vendors, so I can see who wins the version and interface battles. I'm also waiting to see better support options for what is currently a self-supported operating system. The support companies that I liked the best have all gone out of business!

An up and coming option is to skip a traditional network server altogether and go with a stand-alone NAS (Network Attached Storage) device. I have been selling standalone NAS devices for a while and am currently gathering information for some future articles on them. Typically they are based on a hidden Linux kernel, but are operated through a browser-based interface so that no Linux knowledge is required. I think that this option may eventually overcome all NOS-based (Network Operating System) choices for many categories of users.

Now I have to get some smelling salts to revive my staff. The shock of reading this pro-Microsoft column was just too much for them.

A NASty JOB, Part 1

by Wayne M. Krakau - *Chicago Computer Guide*, February, 2001

Wayne is president of Krakau Business Computer Systems, a systems integration firm that is a Novell Gold Authorized and Novell IBM Host Master Authorized Reseller. He has been working with computers for 24 years. He holds an M.B.A. in Marketing and a B.S. in Information Science and was originally educated in Computer and Information Systems Analysis Engineering. He holds CNI, MCNE and CNTE ratings and all five of the original LANDA Certifications. Wayne can be reached for questions, comments, or topic suggestions at (630)894-3360 or via E-mail to wkrakau@krakau-inc.com. For past columns and company information go to http://www.krakau-inc.com.

As the old saying goes, somebody has got to do it (and I don't mean entering a twelve-step program for bad puns). The job in question is that of file server, and the new candidate, with the potential to revolutionize the Local Area Network market, is the independent Network Attached Storage or NAS device, with an emphasis on independence. Before we get to the details of independent NAS devices, we need to cover what Hollywood might call the "back story" leading up to my current enthusiasm for this product category.

Once upon a time in the deep, dark woods . . . oops, wrong back story. One of the most frustrating things about being a VAR (Value Added Reseller) or, more specifically, that subcategory known as Systems Integrator, is the fact that there are times when you just can't provide a solution for a prospective client. The need is obvious to both the integrator and the client, but either the actual funds, or, at least the willingness to spend them is lacking. In this case, the loss of a potential client is not to a competitor, but to complacency and the status quo.

There are different categories of existing computer systems that are eligible for this technological torpor. The first is a client with no network at all. The old term "sneakernet" is most apt in this situation. Floppies, tapes, or some form of removable drives are used to pass data between computers. In particular, I've seen a lot of old Iomega Bernoulli Boxes put to creative use. Since they are nearly indestructible, it is often quite difficult to get someone to retire them.

(Note that I still remember Iomega's old trade show demo in which they attached their drive to a standard hardware store paint can shaker and performed continuous - and flawless - read/write tests. Every few minutes, they would remove the cartridge and whack it really hard against the edge of a table before reinserting it into the drive and resuming the tests. I once encountered a 12-year-old original 10MB model that was starting to malfunction. It seemed to be overheating, so I removed the filter and found it clogged with a 12-year accumulation of fuzz and dust, totaling about a quarter of a pound. I carefully washed and dried the filter and reinserted it. The drive worked perfectly after that, and I was unable to convince the client to even consider discarding it.)

Some form of data isolation is often used to reduce the need for data transfer, typically by dedicating specific computers to specific tasks. Billing is done on computer A, mailing on computer B, and general word processing on computer C, etc. Either data switches or primitive printer sharing boxes are used to reduce the number of printers needed. Sometimes there will simply be one printer per computer. It's not pretty, but it gets the job done. In a similar vein, a bucket-brigade will successfully put out small fires.

The next category up the scale is the peer-to-peer network in which data and printers are shared by a collection of computers. Microsoft's Windows 3.x and Windows 9x networking are the two most common systems found, with Artisoft's Lantastic not far behind. I have also found a veritable potpourri of other uncommon software driving these networks, most of which would be disqualified on a computer trivia test for being just too obscure.

As a practical matter, the peer-to-peer nature of these systems was often violated by dedicating one computer to a specific file serving task. To paraphrase George Orwell in *Animal Farm* all peers are created equal, but some peers are more equal than others.

There are severe compromises in terms of performance and reliability when a workstation (or desktop) operating system is made to do double duty as a server operating system. Even if you do

dedicate a computer to a server task, you are still running a workstation-oriented operating system along with its inefficient and potentially unsafe file system.

Then there is the issue of security. Multiple crisscrossing file and other resource sharing can get very complicated, very fast. Also, remember that security isn't just an issue of watching out for those with malicious intent. It also includes keeping people from accidentally destroying or scrambling data, applications, and system settings, or accidentally spreading viruses.

Again, even though many of these potential clients have long outgrown their peer-to-peer networks, they get the job done. Now they have a fire hose, but, depending upon the vintage, it's attached to either a hand, or a steam-powered pump.

Finally, there is the system that includes a true server-oriented LAN that is either out of date, or darn right obsolete. Considering its historic LAN market share, a lot of these systems are running old, unsupported versions of Novell's NetWare. Some are running old versions of Microsoft's NT. An amazing number are running terribly obscure brands that don't even exist anymore.

(Around 1993, I had a potential client that was a huge corporation. They switched from the then-current version of NetWare to an incredibly rare Network Operating System that was actually a modified form of UNIX. I looked it up and found that it had - at the time - less that one-half of one percent market share. Shortly thereafter, the software company went under. It took that prospective client years to recover from the switch. They are just now finishing becoming a full NT shop. I doubt that they even considered going back to NetWare, as that would have been like admitting a huge mistake.)

With these true LANs, they are now up to real fire trucks along with an integrated system of water delivery (as in pipes). Next month, I'll show you how independent NAS devices can solve the problems of these and other categories of businesses. Meanwhile, I'll be reminiscing about the days before file servers, when disk servers ruled the earth and OmniNet was the latest, greatest way to connect computers. (Hmmm. I wonder if Grecian Formula will work on these gray hairs.)

A NASty JOB, Part 2

by Wayne M. Krakau - *Chicago Computer Guide*, March, 2001

Wayne is president of Krakau Business Computer Systems, a systems integration firm that is a Novell Gold Authorized and Novell IBM Host Master Authorized Reseller. He has been working with computers for 24 years. He holds an M.B.A. in Marketing and a B.S. in Information Science and was originally educated in Computer and Information Systems Analysis Engineering. He holds CNI, MCNE and CNTE ratings and all five of the original LANDA Certifications. Wayne can be reached for questions, comments, or topic suggestions at (630)894-3360 or via E-mail to wkrakau@krakau-inc.com. For past columns and company information go to http://www.krakau-inc.com.

Last month, I detailed some possible candidates for independent NAS (Network Attached Storage) devices. Now, I'll tell you what they are. This is a description of a generic NAS device, so there will be exceptions in some of the details.

In its most basic form, a NAS device is a disk drive in a little black box with a network socket, bundled with intelligence. Traditionally, it has just enough intelligence to control its own disk drive and not much else. It is dependent upon the existence of a traditional file server on the same LAN. Without that conventional file server with its conventional NOS (Network Operating System), the NAS simply wouldn't function. The file server does most of the management of data and security, while the NAS's own brain just concentrates on shoving data down the pipe (that being 10/100 ethernet connection).

In the background, it would really have more intelligence than would be readily apparent, in that it would pretend to be whatever brand of file server it found on the network. It would also have to act cooperatively with the server to allow for security. NAS devices can mimic several different NOSs. Some can even have multiple identities at once, depending upon what type of client is making the request for data. All of this does take a certain amount of computing horsepower.

While there is no law requiring it, almost all NAS devices use a modified Linux kernel to harness that computing horsepower. The kernel is the base, underlying core of an operating system. It's the raw, stripped down operating system without any extras. Each manufacturer modifies the Linux kernel to suit their particular needs. Even different models within a single manufacturer's product line may have different versions of the Linux kernel.

For those of you who have been lost in the Arctic for the last couple of years (remember, penguins are in the Antarctic - and a few other places, like the Hilton in Honolulu - but not the Arctic), Linux is a modern derivative of UNIX (itself a 30-year-old operating system) that was put into the public domain (with certain limitations) by its developer. Since it is effectively free (if you follow the rules), manufacturers love to use it as the embedded operating system for things like NAS devices. Just look for the penguins at any trade show to get more details. (Disney has its mouse and Linux has its penguin.)

To make life easier, and to avoiding requiring customers to have Linux experts on staff, a browser-based (Internet Explorer, Netscape Navigator, etc.) menu system is used to manage the NAS device. None of the systems that I've seen could really be called easy to use (like their advertising says), but they are at least useable. It's not hard to tell that they were written by ex-Unix, now Linux geeks.

These programs generally remind me of the directions to an electronic appliance that was designed overseas, where the directions have been translated into English by someone who learned English as a second language without ever visiting an English-speaking country. Nobody bothers hiring a native-born English speaker to proofread these directions, and nobody at the NAS manufacturer hires an interface design specialist to proofread their management program interface.

At least the interfaces are usually simple enough that the common day-to-day tasks can be taught to an interested nonprofessional (in computing). In many cases, a talented amateur can even get the system initially configured, depending upon how close the customer's needs are to the system defaults.

NAS devices have different models offering different levels of speed and safety, just as file servers do. The simplest and least expensive models have a single disk drive. Next up the scale are models using two mirrored drives. At the top of the scale are models offering a full RAID (Redundant Array of Independent Disks) system, usually RAID Level 5. Within a given model, you can usually configure the disks in whatever way you want, assuming you have enough disks, a minimum of two for mirroring and three for RAID 5. You can trade the extra speed and safety for additional storage.

Finally, after selling dependent NAS devices for a while, manufacturers realized that by simply reprogramming them, using the flexibility of the Linux kernel, they could make them independent of traditional file servers. That was the big breakthrough.

Since they were reprogramming them, anyway, some manufacturers added extra features. Some added SCSI ports to attach shared devices such as CD towers. Others added special functions that turned their products into Web servers, e-mail servers, asynchronous communications servers (modem sharing), and combination router/firewalls. All of this is in addition to their new roles as stand-alone file servers. In addition, they never lost the ability to work cooperatively with conventional file servers.

Next month I'll cover the benefits and limitations of using independent NAS devices. Now, I'm going to see if I can find even more obscure sets of technically-oriented acronyms to further confuse the issue.

A NASty JOB, Part 3

by Wayne M. Krakau - *Chicago Computer Guide*, April, 2001

Wayne is president of Krakau Business Computer Systems, a systems integration firm that is a Novell Gold Authorized and Novell IBM Host Master Authorized Reseller. He has been working with computers for 24 years. He holds an M.B.A. in Marketing and a B.S. in Information Science and was originally educated in Computer and Information Systems Analysis Engineering. He holds CNI, MCNE and CNTE ratings and all five of the original LANDA Certifications. Wayne can be reached for questions, comments, or topic suggestions at (630)894-3360 or via E-mail to wkrakau@krakau-inc.com. For past columns and company information go to http://www.krakau-inc.com.

This month, I'll cover the benefits and limitations of independent NAS (Network Attached Storage) devices. This will lead to the justification of their existence as alternatives to traditional file servers. As in last month's article, I am describing generic characteristics, so there will be exceptions.

The main advantage of independent NAS devices (what an awkward phrase!) is price. The category that I'm talking about ranges from $500 to $4,500, with most of the action happening in the $1,500 to $2,700 range. Right off the bat, that can be a big savings compared to a "real" file server. The extra cash buys extra disk capacity, more safety features, and sometimes, additional application-server type functions.

But, my friends, the savings don't stop there! (I think I'm watching too many infomercials.) Subtract out the cost of a Network Operating System (NOS) - it's included in the box. Then subtract out the cost of specialized, server-based backup software - you backup either from a workstation or with an included utility. You might want to add in a few bucks for workstation-based backup software. Then, subtract out the cost of specialized, server-based anti-virus software. Total this all up to get your initial hardware and software cost of acquisition.

Now (if you call our order line within the next half-hour), subtract out between 80 and 95% of the installation, configuration, and training costs. This is an educated, but admittedly wild guess. In general, the simpler the NAS device, the more you will save. If you have in-house technicians, your savings might be even greater. This will give your initial labor cost.

Finally (as a limited time offer if you order two or more via credit card), subtract out something like 95 to 99% of the maintenance costs. To update the device itself, you simply download a file from the manufacturer's Web site, select the appropriate item off the browser-based menu system, and then sit back and watch for five to ten minutes while the NAS updates itself and then reboots. Oh, you might want to make sure that everyone is logged off before you fire up the upgrade function. Unannounced reboots can be hazardous to your job security.

To update your server-based backup and anti-virus software - well, there isn't any to update. (And, to make this an irresistible deal, we'll throw in a set of steak knives strong enough to cut through a computer case!) Yes, you have to keep the workstation-based equivalents up-to-date, but that's not that big a deal.

Since these NAS devices are specifically designed to be lean, mean, file-serving machines, performance is typically quite good. Even within this relatively inexpensive category, many manufacturers rate their products as qualifying for 50-user networks, depending upon the actual network activity.

Now, for the other side of the story - the limitations. The most blatantly obvious is the inability to run either utility or applications software directly on the file server.

With the exception of those few devices that include an internal backup feature, backups must be done from another computer over a comparatively slow network cable. This means you may have to carefully plan your backup strategy to avoid running into live files, as when an overnight backup lasts until the next day. Hey, a little planning never hurt anybody.

Antiviral scans must be done from another computer. Again, planning is required to overcome the speed limitations of the LAN wire as well as to avoid neglecting scans through complacency. This is not an insurmountable obstacle.

E-mail, network-based faxing, and shared modems can be handled through little black box style specialty servers or by workstation-based solutions. These are nice tasks for old (but not that old) workstations. Alternately, you can use a relatively fast workstation for an undemanding job, such as simple word processing, and use the excess capacity to run the utility software. Check out the solutions from SpartaCom Technologies, Inc. (www.spartacom.com) for some neat workstation-based solutions.

For small organizations, or even medium-sized ones with light e-mail traffic and few e-mail management and tracking requirements, the individual workstations can be set to handle their own e-mail. This requires a direct, continuously link to the Internet, and is only practical for sites where people almost always stay at their own desks. It is also rather labor intensive to set up, but once it is set up, it simply does its thing. Note that in this situation, internal e-mail is bounced off the ISP's (Internet Service Provider's) mail server rather than an in-house mail server.

A more serious NAS software limitation is the inability to run client-server software, in which an application running on the server does most of the work of a database application. (Anything that keeps records, like an accounting system, for example, is inherently a database application, though there are ways to create a very limited database with spreadsheet or word processing software.) There are, however, mitigating circumstances to this limitation.

First, there is the trend, by software companies, to demand, upon the threat of lack of support, that you run their client-server software on a separate, dedicated server, not on a general-purpose file server, or a server running other applications or utilities. This trend seams to have stemmed from the instability of NT servers. Single task NT servers can be made much more stable than ones with multiple duties. Later the trend spread to NetWare servers. In this case, the vendors got disgusted with the quality of installations, not any inherent weakness of NetWare. There just aren't that many properly configured NetWare servers out there. (Both of these opinions - NT and NetWare related - on this trend are my own, gathered from talking to developers at various software vendors. Feel free to take them with large grains of salt.)

This means that even if you have a "normal" file server, switching to client-server software might necessitate the purchase of an additional server, anyway. In that case the decision to use an independent NAS as your file server remains separate from the client-server decision.

A second mitigating circumstance is business practicality. While I am an enthusiastic (some would say fanatic) proponent of using client-server software for both speed and safety, I still remember the stuff about justifying business costs that I learned in business school. On a small network, or one with a lightly used database, the benefits of client-server may not justify the cost. Improving throughput by using efficiently written applications, by switching from 10MB in simplex (one-way) mode to 100MB switches in duplex (two-way) mode or by reconfiguring servers or workstations might be more cost effective. You have to sit down and do the math and the system analysis on a case-by-case basis.

Again, we are back to more of that dirty word - planning! As all roads used to lead to Rome (at least poetically if not in fact) all computing solutions lead to planning. In this case we are planning a revolution in networking, so it's probably worth the effort. Now let me grab the phone before that steak knife offer expires!

COMDEX CHICAGO? - Part 1

by Wayne M. Krakau - *Chicago Computer Guide*, May, 2001

Wayne is president of Krakau Business Computer Systems, a systems integration firm that is a Novell Gold Authorized and Novell IBM Host Master Authorized Reseller. He has been working with computers for 24 years. He holds an M.B.A. in Marketing and a B.S. in Information Science and was originally educated in Computer and Information Systems Analysis Engineering. He holds CNI, MCNE and CNTE ratings and all five of the original LANDA Certifications. Wayne can be reached for questions, comments, or topic suggestions at (630)894-3360 or via E-mail to wkrakau@krakau-inc.com. For past columns and company information go to http://www.krakau-inc.com.

The question mark in the title is indicative of my own opinion of the repositioning of what was a national exposition and trade show known as Spring COMDEX, into a local show, now called COMDEX Chicago. Chicago, it seems, just got demoted.

In fact, the much larger Waste Management Expo show across the lobby dwarfed the computer show. The buzz at COMDEX was that those products that didn't generate enough booth traffic would be fed to the giant, 3-story high garbage-eating machine that was visible from the COMDEX registration area, in kind of a computer industry survival-of-the-fittest. (All right. I admit it, That's a lie. It was actually my own silly fantasy. However, I'll bet I could get a good price for the TV coverage of the event!)

Exhibitors complained and the media commented on the failure of marketing for the show. I call it doing a Novell. Just as doing a Lewinsky has taken on a meaning related to a specific action, I feel that doing a Novell represents taking a world class product and running it into the ground via major and repeated marketing misfires. Even other members of the press joked about the possibility that the COMDEX marketing folks actively recruited former Novell marketers.

My big hope is that, as the show is gradually downsized (sarcasm intended), it will reach some sort of equilibrium with its audience. After some frank talks with management, that hope was renewed. The managers that a gaggle (for want of a better term) of reporters accosted with both questions and suggestions returned not the hostility or defensiveness that one would have expected, but actually welcomed the discussion. This was a new team that acknowledged its errors and solicited and intelligently discussed suggestions. They recognized that Chicago was being under served by trade shows and potential big-name exhibitors, and they planned to remedy the situation. I think they have a serious chance of succeeding.

And now on to the show! It wasn't either awe-inspiring or even particularly surprising. It was, however, a good meat-and-potatoes trade show. The first thing you noticed upon examining maps of the trade show floor was that the largest and most centrally located booth was not occupied by a major computer company. It was inhabited by a car company, Mercedes-Benz USA!

Their exhibit gets the "I wouldn't want to be the insurance agent for that one" Award. After explaining at a press presentation that at Fall COMDEX in Las Vegas, the public was allowed to test drive Mercedes full line of vehicles on a closed test track, the Mercedes manager explained that they were unable to arrange a closed test track in Chicago. There was a palpable sigh of disappointment from the gathered press, just before he revealed that they would allow the test drives to occur on regular, unprotected streets! Geeks with very expensive wheels that they don't own - now that's a scary thought. It gave me visions of someone trying to reenact the car chase scene from *The French Connection.*

The "Close but no cigar" Award goes to an unnamed vendor. The fact that I still can't remember what the company name was and I never figured out what they were selling is indicative of the "no cigar" portion of the award. They started out with what may be the perfect marketing mix for a computer show - women in flashy dresses and the ultimate techie gadget, a Lamborghini. However, they didn't have anything displayed in their booth to indicate what they did!

I have a lot of sympathy for the pair of women at the booth. They didn't quite catch on to the dynamics of the show and seemed to be hurt by the attitude of the crowd. That is, the attendees were annoyed that the women were blocking their view of the car! The women didn't realize that the crowd at COMDEX probably makes the fans at a Star Trek convention look downright suave and sophisticated in comparison. (Just kidding guys! Live long and prosper.) I think it's something about seeing all those tight and/or revealing costumes on the female aliens.

The other odd thing at the show was a double row of employment booths, with a guard of sorts warning people to stay out unless they were really trying to find a job. That the double row of booths was a dead end aisle was somewhat disconcerting from a symbolic point of view. The fact that an employment area existed at a computer show during this major technical stock collapse was, however, a bit heartening.

While there was one booth taken by an H1-B (technical employee immigration) legal specialty firm, most of the employers were actively recruiting people who were actually U.S. residents. This is encouraging, given the recent studies showing how the increase in H1-B workers is directly depressing the salaries of resident technical workers, and, at least by implication, aiding in the discrimination against technical workers over 40. As a U.S. resident technical worker over 40 years old, with lots of friends and acquaintances (of many different nationalities, national origins, and races) in the same category, my opinion on this subject should NOT be considered unbiased. Just because I'm getting older (and slightly more decrepit every day) doesn't mean I can't learn new technologies!

Next month, I'll cover the "meat" of the show (as opposed to this little appetizer). It's a bit lean, but still quite tasty. Meanwhile, I want to put a stop to all those rumors that the only reason I go through the effort to get my press credentials for this show was because of all the free food the press gets. That is definitely not true. Well, at least not completely. Burp!

COMDEX CHICAGO? - Part 2

by Wayne M. Krakau - *Chicago Computer Guide*, June, 2001

Wayne is president of Krakau Business Computer Systems, a systems integration firm that is a Novell Gold Authorized and Novell IBM Host Master Authorized Reseller. He has been working with computers for 24 years. He holds an M.B.A. in Marketing and a B.S. in Information Science and was originally educated in Computer and Information Systems Analysis Engineering. He holds CNI, MCNE and CNTE ratings and all five of the original LANDA Certifications. Wayne can be reached for questions, comments, or topic suggestions at (630)894-3360 or via E-mail to wkrakau@krakau-inc.com. For past columns and company information go to http://www.krakau-inc.com.

Now we move from the overview level of Chicago's own preshrunk exposition, COMDEX Chicago (formerly Spring COMDEX) to the detail level, culling the tasty computing smorgasbord from the electronic road kill.

First, let's start with something that I have been searching for since I first became enamored with the concept of using independent NAS (Network Attached Storage) devices not just as adjuncts to existing servers, but as simple, stand-alone servers in their own right. One off the weaknesses of using NAS devices without conventional servers is their almost universal lack of shared serial communications (with modems) capabilities.

As primitive as plain, old modems may seem in these days of high-speed broadband communications, modems still have their uses. If your broadband line goes down (not an infrequent occurrence with DSL and cable lines), one or more modems can be used for dial-up access to the Internet and e-mail. If you want to tap into your network or remote control your desktop computer from either home or out in the field, a modem is still the lowest common denominator. That remote control function is also helpful in getting technical assistance, too.

Equinox (www.equinox.com) has come up with a simple, single-port solution, their MDS-10 Managed Device Server. It is a stand-alone device, in that it doesn't need a server to control it, so it is eligible for NAS-only networks. It has a 10Base-T connection on one end and a serial port on the other. Like NAS devices, it is controlled via a browser-based interface. With it, you can add a shared modem anywhere you have a network socket and a direct phone socket in the same vicinity. In case you need it, it will also control just about any other non-modem serial device. The MDS-10 neatly takes care of my clients' needs for dial-in, dial-out, and potentially even network-based fax applications.

Next, the folks at Sprint (www.sprint.com) (as in long distance service, cell phones, etc.) have introduced their Broadband Direct service to the area. As long as you can place an antenna in direct line of sight with their antenna on the Sears Tower (or whatever they are calling it lately), you are eligible to connect at from 512kbps to 1.5Mbps for downloads and a maximum of 256kbps for uploads.

As they say in car ads, your mileage may vary by distance, atmospheric conditions (including, I presume, sunspot activity), and, I suspect, over time. There is no bandwidth guarantee, and, the more subscribers they get in a given area will probably mean a tendency toward a lower bandwidth. Sprint also warns that there is an inherent latency in their system that makes it inappropriate for many interactive games. Because of this, I suspect that interactive video as used in video conferencing and in some types of online training might also have problems.

Even with the limitations of their technology, which Sprint seems quite willing to honestly discuss, this service could be just the thing for the many who can't get cable or DSL broadband connections.

On the software side of the show, ReadSoft (www.readsoft.com) has made a major push in the area of forms processing, a field currently just about owned by an old favorite of mine, Cardiff Software, Inc. (www.vardiff.com).While they have not specifically used the phrase, ReadSoft seems to be making a "We're number 2, so we try harder" appeal, similar to the old Avis versus Hertz ad campaign.

In a way, this campaign is working, since, prior to this show, I would have happily referred all inquiries in this field to Cardiff, while now, I would have to at least consider ReadSoft as an option. I can't say, given my limited research at and subsequent to the show, that ReadSoft would win any specific product competition, but, in providing at least a potential alternative, it will at least keep Cardiff on its toes. (No offense meant to Cardiff's fine products. It's just that I have become an anti-monopolist at heart! I get nervous when one company dominates a niche so thoroughly.)

On the extreme low-tech side of things is the QuickStudy line of products by BarCharts Inc. (www.barcharts.com and www.quickstudy.com). This is a series of large, laminated helper charts for various applications and operating systems. I know it's not as flashy as some new software or hardware product, but their charts are very well thought out and organized, and would be very handy for both beginners and experienced users, in that they separate charts for basic and advanced issues.

Finally, another old favorite of mine, The National Christina Foundation (www.christina.org) was at the show. They accept donations of old computers and peripherals and give them new, useful lives in public agencies, schools, hospitals, and other organizations. They have new links with similar state and local organizations to improve their ability to distribute these computers where they are most needed. If your computers are not complete junk (assuming there must be some age and usability limits), you can be sure that your computer will be put to good use, and maybe even get yourself a tax deduction in the process.

I have one additional note in defense of Chicago COMDEX. I hear that this year's Networld + Interop Show only had about 60% of the attendees of last year's show. I guess that Chicago COMDEX isn't the only show that is having problems.

BROADBAND - Part 1

by Wayne M. Krakau - *Chicago Computer Guide*, July, 2001

Wayne is president of Krakau Business Computer Systems, a systems integration firm that is a Novell Gold Authorized and Novell IBM Host Master Authorized Reseller. He has been working with computers for 24 years. He holds an M.B.A. in Marketing and a B.S. in Information Science and was originally educated in Computer and Information Systems Analysis Engineering. He holds CNI, MCNE and CNTE ratings and all five of the original LANDA Certifications. Wayne can be reached for questions, comments, or topic suggestions at (630)894-3360 or via E-mail to wkrakau@krakau-inc.com. For past columns and company information go to http://www.krakau-inc.com.

If you are of the Rat Pack (Sinatra, and friends and hangers-on) generation, you are probably thinking of an incredibly politically incorrect reference for an all-female orchestra. If, like me, you are a child of the Rock Era, you are thinking of revivals of old rock bands (as in the recent visit to Chicago of *Deep Purple*) populated by now-portly original band members, frequently with supplemental new members. Here, it means what is colloquially and generically called the big pipe, otherwise known as a high-speed connection, in this case, to the Internet.

Since my business deals only with businesses and institutions, I was originally going to cover only the pure business aspects of DSL (Digital Subscriber Line) service, using a title such as the gentler sounding *DSL Heck* to differentiate my articles from the thousands already written using the title *DSL Hell*. However, many businesses are examining, though not necessarily choosing, other broadband technologies, and, the individual employees of many businesses are evaluating methods of home connection. So, I have expanded the subject matter to compare technologies.

Why would anyone bother with looking for a big pipe as opposed to sticking with standard, asynchronous modems? They would bother for the same reason that a fire department doesn't use garden hoses. Fire departments use various sizes of multi-inch-diameter hoses as needed for a given fire. Businesses and individuals should use the appropriate size "pipe" for their needs, too. Please keep in mind that even when I talk about individual use of a connection, I am concentrating on use for business purposes.

The most obvious business use for broadband access to the Internet is for employees who have to do research on the Internet. There is a direct savings in terms of productivity when someone switches from using a dial-up (modem attached to a standard phone line) connection, or worse yet a dial-up connection that is dynamically shared by multiple users, to a broadband connection. Those cartoons of the skeletal remains of a computer user who died while waiting for a response apply here.

From personal experience, I can tell you about, not only the direct loss of time while waiting for a Web page or a downloading file, but also of the indirect cost of the creative loss of a train of thought. Whether it's gathering information for an article, researching a client problem, or designing a LAN, the constant annoying interruptions do take their toll.

E-mail is an even more common use for a big pipe, but one that may be useable with a slower link. There are, however, concealed costs to limiting e-mail connection speed. Large or frequent e-mails can clog the system. Also - it gets a bit subtle here - even people who have no direct need for pure research on the Internet may have a hidden need for regular access to the Internet that could justify a fast connection.

Think about a telephone salesperson in a nontechnical field as an example. On the surface, no fast Internet access is needed. However, that person may conscientiously subscribe to both industry e-mail newsletters and sales technique e-mail newsletters. That sounds perfectly justifiable, and maybe even admirable, but it has speed-related implications. E-mail newsletters these days are often in Web-style HTML (Hypertext Markup Language) format which, after adding graphics and other doodads, are a LOT larger than their simple text brethren. For good measure, newsletters often only have summaries of

articles. If you want to see the full article, you click on the link, your browser opens, and even more info comes pouring down the pipe. All of this time, the salesperson is waiting.

Also, a salesperson had better be very familiar with the company Web site, which as often as not, is hosted by an ISP (Internet Service provider) and is only accessible via the Internet. It wouldn't hurt for a salesperson to be intimately familiar with the competitions' Web sites as well. It might even be appropriate to keep track of the suppliers' and customers' sites, too. All of this adds up to a hidden need for a single person who, at first glance, needs only very limited access to the Internet. You need to make this calculation for all employees and use those calculations as the basis for deciding on your Internet connection.

Then there is remote control access to consider. If used simply for remote diagnosis by corporate technicians or an outside vendor (like my firm), the cost of the connection time is that of both the idled employee (or employees, if a server is involved) and that of the technician. If used for access by employees either traveling (via the ever-growing number of broadband links to hotels) or at home (via some form of personal broadband), only that employee's time is potentially wasted on a slow link.

But, as they say in infomercials, there's more! Next month I'll cover more business justifications and then move on to broadband choices.

BROADBAND - Part 2

by Wayne M. Krakau - *Chicago Computer Guide*, August, 2001

Wayne is president of Krakau Business Computer Systems, a systems integration firm that is a Novell Gold Authorized and Novell IBM Host Master Authorized Reseller. He has been working with computers for 24 years. He holds an M.B.A. in Marketing and a B.S. in Information Science and was originally educated in Computer and Information Systems Analysis Engineering. He holds CNI, MCNE and CNTE ratings and all five of the original LANDA Certifications. Wayne can be reached for questions, comments, or topic suggestions at (630)894-3360 or via E-mail to wkrakau@krakau-inc.com. For past columns and company information go to http://www.krakau-inc.com.

This is the second part of the coverage of broadband connections, A.K.A. big - and fast - pipes to the Internet. We are finishing business justification and moving on to the broadband choices available.

If you have a multiple branch organization or have employees who regularly work from home or out in the field, your broadband connection can be the foundation for a Virtual Private Network or VPN. Using either hardware or software, information is heavily encoded and passed between the main site and the remote site, allowing the remote users to log into the main local area network as if they were attached via a local workstation.

Obviously, the practicality of this arrangement varies both with the respective speeds of the communication connections (local and remote) and the planning that went into the underlying configuration of both system and applications software. If both ends of the connection are high-speed broadband, it will be a lot easier to get work done than if one end is a dial-up connection. Amazingly, if only a single computer is on the remote end, even the lowest speed broadband connection can be practical. Configuration and planning are the keys to getting the best out of this arrangement. Using client-server databases and loading all programs on the remote computers can make a slower-speed connection almost like being locally connected. Dial-up connections won't perform nearly as well, but might be appropriate for limited tasks.

VPN security and encoding are done either by dedicated hardware or by software running on a server at the main site and on the remote workstation. Amazingly, if you select the appropriate products, you can run hardware VPN at the main site and choose either hardware or software for the remote computers. Whenever possible, I prefer hardware VPNs because of their speed and, especially, their reliability. Individual Windows-based computers, in particular, have not proven to be either very reliable or very hacking resistant. Also, running Microsoft-designed security software is like having a target a giant target on your back. It's the most popular to hack. Software VPNs work, but, if the client can afford it, I suggest they stick with hardware at least for the broadband-equipped remote computers.

One hidden benefit of a broadband connection is the ease in which computers can be updated. It is quite easy to run Windows Update and various other semiautomatic updating procedures either built into software or using third-party utilities (I really like Big Fix from www.bigfix.com.) on multiple computers at once by simply running from machine to machine every time there is a waiting period (for a reboot, or the infamous hourglass). By the time you have started the process on the last machine in a series, the first one is waiting for your next response. I've done up to 14 computers at once using this method. (Warning - Don't try this while people are working on the computers. This is for non-work hours only. You didn't actually want to have a personal life on weekends and evenings, did you?)

There is a huge savings in both actual update time and organizing time in doing a mass semiautomatic update via a broadband connection as compared to individually downloading and tracking updates and applying them manually. (Of course, you can use your own judgment for really huge downloads such as a new version of Internet Explorer.) Even software disbursement utilities require more time and effort. Also, remember, a computer that has ANY access to e-mail or the Internet, no matter how

slow, absolutely must be kept updated to avoid security problems. (I'll have more on that later in the series.)

Okay. Let's assume you have bought into one or more of my arguments in favor of having a broadband connection. Or, perhaps you've come up with some of your own. (Playing games and downloading porn don't count!) Now, you have to figure out what type of broadband connection you want to use.

If you need the highest level of reliability, with guarantees and nonperformance penalties, and have the bucks, then select a T-1 line (1.544mbps - millions of bits per second) or a fractional T-1 line (obviously, some fraction of 1.544mbps). T-1 lines have much faster and much more expensive big brothers if you really have money to burn, but I won't be covering them in this series.

The T-1 family is also the traditional technology for direct, one-to-one connections between offices and, via a technology called frame relay, pooled access (one-to-many and many-to-many connections). T-1 or its big brothers, are preferred to a VPN for the truly security conscious, since your lines never hit the Internet. You can also use multiple T-1 lines for extra aggregate speed or for redundancy.

The T-1 family can be very expensive, depending upon your location. A few years ago, before Ameritech had switched frame relay (billed by the amount of data), I put together a one-to-many bid for a client using plain frame relay (billed at the rated speed for 24x7, whether you use it or not) for two full T-1 lines into a main office and various speeds ranging from 256kbps to 384kbps into 10 branch offices, all in the Chicago area. Not including the price for the hardware needed, the bill was estimated to be more than $72,000 - PER MONTH! They declined. We are talking about the possibility of putting in a DSL-based (Digital Subscriber Line) VPN for them later this year for noticeably less than that price per year with most of the hardware thrown in for free.

If you use T-1, watch out for the CIR (Certified Information Rate), which is, essentially, the minimum rate you might get. I've seen bids for 384kbps fractional T-1 lines with only a 9600bps CIR! That means during busy times, like during the business day, you would be better off using a modem, and an old model at that.

One sort of fake-out broadband connection uses several modems connected to a "little black box" that aggregates the combined bandwidth of multiple POTS (Plain Old Telephone System) lines to approximate a fairly low-speed broadband connection. These were popular for a while and are still sold in cases where absolutely nothing else is practical. However, true broadband connections are available for less than the cost of the multiple phone lines and multiple Internet dial-up accounts need to make this solution work. I suspect that this method is actually losing market share as few new installations are done and old installations are being replaced by true broadband connections.

Next month I will continue covering broadband options, including those that aren't quite as obscure as the aforementioned multi-modem solution.

BROADBAND - Part 3

by Wayne M. Krakau - *Chicago Computer Guide*, September, 2001

Wayne is president of Krakau Business Computer Systems, a systems integration firm that is a Novell Gold Authorized and Novell IBM Host Master Authorized Reseller. He has been working with computers for 24 years. He holds an M.B.A. in Marketing and a B.S. in Information Science and was originally educated in Computer and Information Systems Analysis Engineering. He holds CNI, MCNE and CNTE ratings and all five of the original LANDA Certifications. Wayne can be reached for questions, comments, or topic suggestions at (630)894-3360 or via E-mail to wkrakau@krakau-inc.com. For past columns and company information go to http://www.krakau-inc.com.

We are now up to the third part of my series on those fast connections known as broadband. This article continues with coverage of the "meat" of the system, the type of connections available.

One up and coming Internet access technology is wireless. There are enough limitations in wireless connections so that they are typically the high-speed link of last resort.

There are two types of wireless connection, land-based and satellite-based. In both methods, you aim a dish-style antenna at the source of the signal.

On land, that antenna is attached to either a tall building or a tower of some sort. In Chicago, Sprint Broadband Direct (www.sprint.com) uses the Sears Tower and promises to add other, alternate antennas in the future. For satellite connections, the antenna is a low-orbit geosynchronous satellite hanging over a fixed spot on the equator.

In either case, the key principle is line-of-site. Assuming you are allowed to put up a dish (not an unimportant consideration), your dish has to have an unobstructed direct view of the antenna. If any buildings, trees, or even the curve of the earth's surface get in the way, you don't get a signal. Obviously, planning must include making sure that nobody is going to build or plant anything that will block your signal in the future. Flocks of birds, bad weather, and even sunspots can seriously disrupt the signal. In addition, there is a latency issue. That means that there is a lag time, sometimes quite a large one, in getting your signal. This can really screw up games, VPNs (Virtual Private Networks), streaming audio and video, and other timing-related software.

Historically, wireless connections have been one-way, where the return signal was carried via a regular phone line. The recent availability of two-way wireless will probably put an end to those companies still offering only one-way service. However, two-way wireless is typically very asynchronous, with the upstream channel sometimes barely exceeding regular modem speeds.

To put a bit more of a crimp in your enthusiasm, even the faster downstream rates are not guaranteed. You may find that the only time you get anywhere near the rated speed is around 3:00AM or so. Peak period speeds may be nowhere near the advertised speed, with a possible mitigating factor being that they define early evening, not regular business hours, as the peak use period. As your service provider becomes more successful and gets more customers, you will find out the disadvantages of shared bandwidth, where more is definitely not merrier. Having more customers means you get a slower effective speed.

Also, everybody gets every downstream signal, but only the correct receiver unit will do something with it. At least that's the idea. Of course, that's how cell phones are supposed to work, too, and you may have noticed that people are actively selling recordings of intercepted analog cell phone conversations. If you are passing sensitive data, you might want to check and see just how much built-in encoding your provider uses. If it's not sufficient, you may want to add your own encryption to any communications with another site that you control, or at least to your e-mail software. You don't want to find the secret planning report you sent to your main office published on the Web!

An additional limitation is the proprietary nature of many wireless links. StarBand (www.starband.com) satellite links, for instance, must terminate in a PC, without an intervening router or

firewall, and must have proprietary compression/decompression running on that PC. This makes using common add-ons like firewalls, VPNs and even remote control software difficult, if not impossible. Now you know why I consider wireless connections the links of last resort.

Now, on to digital broadband cable. Yes, that's the same wire that brings you everything from the *Soprano's* to *Sesame Street*, with varying degrees of success. The first thing to remember before choosing cable is that every time your cable TV goes out, that represents a data outage you would have experienced. Also, every time your digital cable TV starts pixellating, where the picture regenerates so slowly that you see the individual dots slowly painting on the screen, that represents an extreme data slowdown - to a snail's pace - that you would have experienced. Sorry to be pessimistic, but you have to remember that cable TV consistently rates at the bottom of the scale for level of service.

The good news about cable is its potential for very high speeds, in the neighborhood of several millions of bits per second. At that speed, even having a slower upstream speed isn't that much of a disadvantage. However, at the moment, most cable systems are only offering asynchronous speeds in the range of 1.5mbps/384kbps or 768kbps/128kbps (downstream/upstream). To further annoy their customers, virtually all cable contracts allow providers to alter the rules at any time without notice. Several have arbitrarily cut some types of upstream traffic back to only 64kbps from 384kbps, giving a nasty surprise to existing customers.

While the cable industry has been quite hostile toward business use of their services, even to the point of discouraging telecommuters and home/office users, there is some indication that this may simply be a way to sell a separate business-level service by making the consumer-level service so incredibly unappetizing for business users. Some cable providers are starting this type of service, though not yet in the Chicago area.

Next month I will continue covering digital cable and move on to the most common choice of business users, DSL (Digital Subscriber Line). Now, as I watch *Shark Week* on the *Discovery Channel*, I am reminded of what an Australian shark expert once told me about how to avoid becoming lunch for a great white shark. He told me to figure out how to do the best impression of a seal that you possibly can - and then do the exact opposite.

TOLERANCE

by Wayne M. Krakau - *Chicago Computer Guide*, October, 2001

Wayne is president of Krakau Business Computer Systems, a systems integration firm and has been working with computers for 25 years. He holds an M.B.A. in Marketing and a B.S. in Information Science and was originally educated in Computer and Information Systems Analysis Engineering. He holds CNI and MCNE ratings and all five of the original LANDA Certifications. Wayne can be reached for questions, comments, or topic suggestions at 630-894-3360 or via E-mail to wkrakau@krakau-inc.com. For past columns and company info go to www.krakau-inc.com.

We interrupt this series on broadband communications with this special bulletin on diversity in the computer industry, in light of the recent terrorist attacks.

First, let's start out with the basic fact that, long before they invented H1-B status, the computer industry has included a high percentage of foreign nationals who, in many cases, have different social, cultural, and, more importantly (for the sake of this discussion) religious traditions. As a corollary, let's accept that many of these people are readily identifiable as foreign through physical characteristics, accent, mode of dress, or their actions. Finally let's take notice of those people born in the United States who could be mistaken for foreigners.

Now let's mix in the idea that most people in this country have no serious idea of their own religion's history (as a separate issue from that religion's teachings) and, frankly, haven't a clue about any other religions. Well, that's a bit of an exaggeration, but not much. Sure, people in one Christian denomination might know a little about other Christian denominations, and have at least some knowledge (not necessarily correct) about Judaism. However, since each denomination of each religion generally sees itself as THE one ultimate religion, their members don't feel compelled to learn about others.

Also, let's accept as given that they call us a Judeo-Christian society because the country is overwhelmingly ruled by Christians (mostly of the Caucasian variety) along with some Jews. It is not ingrained in our society that there are other perfectly legitimate religions, some having a common belief in the Judeo-Christian concept of God, others having multiple gods, and still others not having any concept of God or gods. Our society also lacks acknowledgment of people of good conscience who don't follow any religion - I've heard one explain "Why should I use an operator when I can dial direct?" - and those who feel no need for even a discussion of such matters.

If you buy my "facts" (or, if you prefer, my pinheaded opinions) you can see that the current national crisis could easily result in massive turmoil in the computer industry due to the possibility of actions taken against people who "look" or "act" Muslim, or more directly Arab. This is despite the fact that Islam includes different nationalities and races from around the world.

After spending much of my teens and twenties doing independent research on comparative religions, I found that all religions, no matter how well-intentioned their underlying philosophies, need to have their writings interpreted by mere humans who see things based on huge helpings of self interest, lots of potential prejudice, and a general all-around willingness to rationalize their actions. Hence, we end up with either outright vindication or at least a passive acceptance of some ethically questionable actions by religious authorities. How about the Inquisition, the Crusades, the wiping out and/or subjugation of native peoples throughout the Americas and elsewhere, the enslavement of Africans and others, the German treatment of Jews and others in World War II, the Israeli/Arab conflict, the Iran/Iraq conflict, the India/Pakistan conflict, the religious revolt in the southern Philippines, discrimination against nonwhites and Jews, and the terrorist attacks on the United States have all been either endorsed, or at least ignored, by some religious bodies. Do not start blaming all followers of Islam just because some groups are adjusting their religious interpretation to suit their own needs. Other religions, including Christianity, have had similar problems.

We now have the opportunity to either tear the computer industry (and the US) apart, or, make a decision to act according to the best principles of our respective religions or ethical standards of conduct. Do we really want to fight misguided and misused religious fundamentalism with our own version of religious fundamentalism?

I have always been annoyed when coaches invoke God's support in sporting events (as if God really wants you to maim the quarterback), but I am really disturbed when politicians start calling on God's support (as if God is really tied to political entities, like political parties or even countries). Convincing people that we are the Chosen People is a great way to justify and rationalize any action. Then next logical step is that the politician's decisions are backed by God. It's happened every time in the past in the US as well as other countries. The Founding Fathers didn't come up with the concept of the separation of church and state just for laughs. Has anybody heard about the nightmare of McCarthyism and how Americans suffered in its fight against "Godless" communists?

As a practical matter, during this crisis, the more time the police waste protecting Muslims (or alleged Muslims) and tracking down those who harm or harass them, the less time they will have to track down the terrorists and their helpers. Again, the more time they waste finding and interrogating "suspicious-looking" Muslims (or alleged Muslims) based on little more than prejudice, the less time they will have to track down real leads involving genuinely suspicious activity. Realistically, any Muslim who looks shifty these days is probably justifiably looking around due to being worried about being beaten or killed by non-Muslims.

Why, you may ask would I risk offending my fan base (both of you) on such risky subjects as religion, race, and politics. There are two reasons. One is a sort of familial guilt in being a first generation German-American (on my father's side) in a family with lots of relatives both here and back in the old country who were wholeheartedly in favor of exterminating Jews and other "undesirables".

The other reason is that I am afraid that my country will slip into some old and very bad habits, all while keeping the facts secret. Specifically, I have never seen or heard reports of the internment of Germans and Italians in concentration camps in the US starting just prior to the US involvement in World War II and ending only after the War ended. The only reason I know it happened is because almost all of my father's extended family spent the War in these camps with thousands of other people.

Since some of these people were moved from camp to camp (My relatives went from Fort Knox to somewhere in the Dakotas, and eventually to Crystal City, Texas.), they were later able to put their heads together to come up with a very rough estimate of 55,000 Germans (including Austrians) and 15,000 Italians. While they were never kept directly with the Italians, they were occasionally in adjacent camps, so that the number of Italians is only based on limited observation.

Note that a concentration camp is a place to literally "concentrate" undesirables and is not synonymous with "death camp". The camps were very much like what you would see in prisoner-of-war movies like *The Great Escape* or *Stalag 17*, though the wooden huts tended to be smaller. There were machine gun towers, armed guards, guard dogs, and lots and lots of barbed wire. I have seen drawings that the internees smuggled out and they match the verbal descriptions. At one point, there was an adjacent, identical camp that was used to hold German prisoners of war from various U-boats.

These concentration camps were not fit for people convicted of crimes in the US, much less whole families of law-abiding residents, some native-born or naturalized citizens. A few of the internees were actually big mouths who at least sounded suspicious, but most were just plain folks. Also, pretty much all of these just plain folks lost all of their possessions and property when they were interned.

Years ago somebody managed to expose the secret of the Japanese internment, but they seemed to have missed this one. The fact that this activity is still essentially unknown to those whose families were not involved scares me. Maybe it can happen, again. Now you know why I am so worried about what might happen in the wake of the terrorist attacks.

By the way, the only good thing that anyone mentioned about the camps was that my uncle met the woman he later married there. How's that for a tough dating service?

BROADBAND - Part 4

by Wayne M. Krakau - *Chicago Computer Guide*, November, 2001

Wayne is president of Krakau Business Computer Systems, a systems integration firm and has been working with computers for 25 years. He holds an M.B.A. in Marketing and a B.S. in Information Science and was originally educated in Computer and Information Systems Analysis Engineering. He holds CNI and MCNE ratings and all five of the original LANDA Certifications. Wayne can be reached for questions, comments, or topic suggestions at 630-894-3360 or via E-mail to wkrakau@krakau-inc.com. For past columns and company info go to www.krakau-inc.com.

Well, we are now back to the mundane coverage of the everyday world. This means a continuation of my series on broadband communications. I know it's not as exciting as watching the continuously updated coverage of anthrax false alarms, but we do have to get on with our lives.

DSL or Digital Subscriber Line is simultaneously the most popular and most criticized broadband technology in use today. It's popular simply because it's cheap, at least compared to its predecessor, T-1. The criticism side of the equation gets a bit more complicated. One interesting quirk is that everybody, myself included, erroneously calls it a DSL line, in spite of the fact the word "line" is already built into the acronym. (Soon, I'll be entering a twelve-step program for abusers of TLAs - Three Letter Acronyms.)

The first thing that you find out when you research DSL is that there is almost too much info. (Begin at www.dslreports.com and go on from there.) There is a whole bunch of different subtype within this technology, each with its own acronym, and lots of picky technical details. Luckily, only two subtypes are really common, SDSL (Symmetric) and ADSL (Asymmetric).

The biggest difference between them is in the way speed is given. SDSL has the same speed upstream and downstream while ADSL can be different, with the upstream speed slower, usually by a lot. While it's not usually sold that way, there is no law that says you can't pay the extra price to match speeds on ADSL.

The other difference in the two is that of target marketing, based on their underlying characteristics. SDSL is aimed directly at business, and is typically bundled with all of the extra features that a business needs. ADSL is aimed at individual users, and to a limited extent, at very small or home-based businesses. The business consumer has to be wary, however, as ads for ADSL sometimes do everything they can legally get away with to obscure the fact that they are for ADSL services.

While SDSL speeds are not absolutely guaranteed, they can be depended upon to stay at very close to the maximum advertised rate. So far, the ones that I have tested have been at 95% or higher. ADSL, on the other hand is cloaked in terms like "as fast as" or other such weasel words. During "peak" times, however that is defined, you might even get less throughput than you'd get with a modem. Even during "non-peak" times, maybe you'll get the rated speed, and maybe a lot less. Also, the contracts usually let the provider throttle back at will. That's why businesses usually don't bother with ADSL.

You may have noticed that I have, so far, neglected technical issues, especially the specifics of speed. That's because things get a bit muddled here. I'm going to play it safe (and short) by using mostly concepts, not specifics. Feel free to look up the details on your own.

First, DSL speed is dependent upon the distance from the phone company office to your site. The farther away you are, the slower the speed, with SDSL having shorter limits (usually quoted as 12,000 to 16,000 feet max) than ADSL (usually quoted as 16,000 to 18,000 max) feet. Inserting a phone company emulation device in between the real phone company and the end user can alleviate this. Phone companies in areas with high market growth possibilities are doing this. Most SDSL and ADSL lines peak out at 1.5Mbps (millions of bits per second) for short distances, but I've seen 2.3Mbps and even higher advertised. Newer types off DSL are faster, but are still quite rare. Note that the distance to your site listed in the phone company's computer might not be accurate.

Another way around the distance limitations is the DSL variant-of-last-resort (my term) IDSL (ISDN DSL, where ISDN is Integrated Services Digital Network). It can reach out beyond 21,000 feet, but can only run at 144Mbps both ways, and costs just as much as SDSL.

The distance limits are somewhat esoteric because they combine actual physical limits and the limits imposed by providers local phone companies due to either reliability worries (longer is less reliable) or marketing (hence, "usually quoted as"). Cheap DSL service can potentially replace expensive T-1 service, thereby putting a large dent in phone company profits.

Other physical limitations can also prevent DSL service. Specialized signal repeating or filtering devices placed between your site and the phone company will disqualify a given line. Segments of fiber mixed in with the copper cable can also prevent DSL, since the specification is for copper cable, but a few companies have found ways around that. Since some areas have run out of end-to-end copper lines, and many phone companies and providers can't seem to coordinate in cannibalizing existing lines, this could keep you from getting DSL. Additionally, the selected cable might not test out as being of a high enough quality to use for DSL.

The actual task of coordinating the installation of DSL among the phone company, the provider, and, where necessary, a third-party intermediary firm, can cause delays of many weeks in getting a DSL line put in. There is also the added risk of misunderstandings or disputes causing further delays.

Now here is the fun part! DSL providers and their third-party intermediary partners are dropping like flies. If one of them goes under while you are connected, you could lose your line. In many, though not all cases, phone companies and some larger providers are cherry-picking the most profitable areas from these failing companies and either reconnecting customers or keeping them from getting disconnected in the first place. If you are not in one of these areas, or the company goes under too fast, you are out of luck. Phone companies, of course, seem to be doing their best to encourage the demise of these perceived competitors.

In spite of all of this nonsense (and even more I don't have room for), I still recommend DSL as the first choice in broadband for my clients. It may not be as reliable or as well supported as T-1, but it's cheap enough to be attractive despite its limitations. All I ask is that they monitor the news about their provider and any intermediary (If I sold them the line, I do that.) and keep a filled-out application form for an alternative provider handy, just in case.

BROADBAND - Part 5

by Wayne M. Krakau - *Chicago Computer Guide*, December, 2001

Wayne is president of Krakau Business Computer Systems, a systems integration firm and has been working with computers for 25 years. He holds an M.B.A. in Marketing and a B.S. in Information Science and was originally educated in Computer and Information Systems Analysis Engineering. He holds CNI and MCNE ratings and all five of the original LANDA Certifications. Wayne can be reached for questions, comments, or topic suggestions at 630-894-3360 or via E-mail to wkrakau@krakau-inc.com. For past columns and company info go to www.krakau-inc.com.

It's the time for me to cover the security implications of having a broadband connection. But first, I'd like to apologize to anyone who might have been confused or mislead by a major mistake that I made in last month's article. I let my fingers, not my brain, do the walking and used the terms "synchronous" and "asynchronous" to describe the two major types of DSL when I should have used "symmetric" (SDSL) and "asymmetric" (ADSL). This was with a 6-inch stack of reference material in front of me that I had just reread. Obviously, old habits (and communications terms) die hard. (I'm blaming it on several years of inhaling fumes from keypunch machine lubricating oil followed by many years of exposure to radiation from monitors. That's my story and I'm sticking to it!)

The first thing to do for security when you have a broadband connection (or, for that matter, even a dial-up connection) is to put anti-virus software on every computer and religiously keep it updated. You will almost certainly have to reconfigure the software to the settings appropriate for your system, as the default settings for anti-virus software are usually both inefficient and incomplete. You might not wish to scan all files going to and from your C: drive and your server, but you certainly will want to scan everything coming down the pipe from the Internet, including e-mails, attachments, downloaded files, and Web pages. A recent option is to put in a hardware-based anti-virus product, but, so far, I have only seen that option used in the field as an addition to, but not as a replacement for, anti-virus software.

Next, if your router does not have firewall capabilities, or those capabilities are inadequate for your situation, get a separate, hardware firewall. This advice is aimed at business computers and at any home computers that can attach to a business network, but it couldn't hurt for pure home users to heed it as well.

You may optionally add software firewalls to your computers, but they have a lot of limitations. Most of the products in this category are worse than useless, as they are easily cracked by even the simplest attack, while giving the user a false sense of security. Check out Gibson Research (www.grc.com) for the latest test results. The only products that have consistently passed Steve Gibson's (and other) tests from the beginning are ZoneAlarm and ZoneAlarm Pro, from Zone Labs, Inc. (www.zonelabs.com). You can download and perform the tests yourself. Other products are just now adding improvements to tighten up and pass these tests.

I have found, in my experiments with software firewalls, they can be a real pain to use. They won't even run on some computers, and when they do run, they have the potential to conflict with lots of software. Even the one that I like the most, ZoneAlarm Pro, is probably going to bombard your computer users with confusing and potentially panic-inducing messages. Also, if, as non-techies, they respond, "Yes" to every warning message, the software will allow inappropriate actions. (I don't even want to get into the argument about basing security products on Windows.)

Step one in configuring your router or firewall is to CHANGE THE DEFAULT PASSWORDS! This is such a common mistake that it is really ridiculous. While you are changing the password, make to make it the largest, most complicated, password you can possibly tolerate, and make sure it is completely unrelated to anything in the real world (words, names, dates, places, ID numbers, etc.). Also don't cheat by writing it on a Post-It note and sticking it to the firewall!

Then you need to turn off command access to the firewall from the outside world. It can take several commands or menu selections to do this. With this access turned off, the router or firewall will ignore commands originating from its Internet-connection socket (typically called the WAN - Wide Area Network - port). Just make sure that you document the procedure so you can temporarily open up access if the manufacturer, your ISP (Internet Service Provider), or your systems integrator (like me) needs to tap into it. (Don't forget to close it up afterward.) Leaving this access open is the second most common mistake that I see.

The simplest router/firewall-based security feature is Network Address Translation, or NAT. It translates non-routable addresses that your individual computers can use into the real address of your router, which is visible to the outside world. It also solves the problem of having more computers and other network devices than you have assigned, real IP addresses. NAT is reasonably effective on its own for non-controversial, not widely known sites that have no internal Web, mail, or other servers that need to be accessed from the outside world. NAT is not perfect. It can be hacked, but only if someone is willing to make the extra effort to monitor your traffic. Using NAT alone is like buying only a basic insurance policy. You have to decide just how risk-averse you are given your particular circumstances. If you want something more than basic, you have to pay.

Adding software firewalls to your computers is one way of overcoming some of NAT's weaknesses. The most common way, however, is to activate stateful inspection (your buzzword - or really buzz phrase - for the day) on your hardware firewall, and to program the firewall to filter out inappropriate traffic, both inbound and outbound. This is where some free, or heavily discounted, router/firewall combinations may be inadequate for your system.

Firewall programming and maintenance take time and money, either for education and day-to-day log file and security bulletin tracking by in-house staff, or for outside organizations to do it for you. That is why it this technique isn't done often enough or as thoroughly as is necessary to provide adequate security.

An additional option is to have your ISP activate a security feature at their end as a replacement for or (more likely) in addition to your own firewall. Presumably, they already have the staff needed to keep their system up to date. In addition to taking on all of the maintenance overhead, they can block Denial of Service (DOS) attacks, which you can't really do from the customer end of a broadband connection.

That's it for my limited summary of the security responsibilities that you take on when you fire up your broadband connection. You need to take this subject very seriously, or you will regret it.

Not Living Up To Your XPotential

by Wayne M. Krakau - *Chicago Computer Guide*, January, 2002

Wayne is president of Krakau Business Computer Systems, a systems integration firm and has been working with computers for 25 years. He holds an M.B.A. in Marketing and a B.S. in Information Science and was originally educated in Computer and Information Systems Analysis Engineering. He holds CNI and MCNE ratings and all five of the original LANDA Certifications. Wayne can be reached for questions, comments, or topic suggestions at 630-894-3360 or via E-mail to wkrakau@krakau-inc.com. For past columns and company info go to www.krakau-inc.com.

Here is my somewhat simplistic take on the current state of Microsoft, as embodied by Windows XP. Let's start with some analogies.

For dog and cat owners, it's like finally getting your pet to stop leaving "presents" on the carpet, but still having to put up with your dog chewing up your shoes or your cat using your furniture as scratching posts.

For fish owners, it like finally getting favorite fish to stop eating its aquarium neighbors, only to realize that your favorite is getting noticeably bigger while none of the other fish seem to be having offspring, in spite of the fact that they frequently look pregnant.

For parents, it as if you finally got word from the teacher that your child has stopped beating up the other children (or at least learned to do it quietly enough so that the teachers don't notice it), but that "Little Billy (a randomly chosen, relatively common first name, of course) isn't living up to his potential." (Of course, parents do have the advantage over pet owners in knowing that their children, unlike animals, will eventually be vulnerable to common logical arguments. According to my informal poll of parents, that will occur when a child is about 35 years old - if you are lucky.)

It is at this time that I am tempted to suggest that somebody should give Mr. Gates a good, sharp rap on the snout with a rolled-up newspaper (something which I would not do to an animal or a child). He runs a company with more money than a lot of small and even medium-sized nations, and, due to lack of will more than anything else, declines to produce anywhere near the quality of software that his company has the potential of creating.

He does take incredible glee, however, in yanking around the public, the government, and especially, any potential competition. While I do have a lot of understanding of at least some reasonable level of wanting to screw up the competition, and perhaps even getting the better (legally!) of the odd government agency, I don't have much empathy for his "let the public be damned" attitude.

Who hasn't encountered individuals who, for reasons you cannot fathom, if given the opportunity to make $100 either by normal, ethical means, or by means that included lying and cheating, would choose the lying and cheating method? This decision would hold, even if the lying and cheating method took more work and entailed more risk.

One would suspect that Mr. Gates, and the people he has gathered around him, suffer, at least to some extent, from that same personality flaw. Success for these people just isn't enough if they can't somehow get away with skirting the rules at the same time.

I hold that, during the last few years, at any point since Microsoft first achieved some level of market dominance, it could have changed its strategy to that of creating the absolute best, most reliable, most secure, and most cost effective products, while marketing and distributing them in an aggressive, but totally ethical way, and made not one dime less than it has made using its existing strategy. The problem is, that would take the fun out of things. How can you feel superior if you can't feel that you have somehow gotten away with something that lesser beings can't? (The late Mayor Washington's remark about hubris comes to mind.)

As one of those "lesser beings" who has a silly habit of following the rules, I am not exactly unprejudiced in these matters. I have lost uncounted thousands of dollars (and the occasional client) while

trying to overcome deficiencies in Microsoft products that I have either sold, or, by laying hands on a PC already loaded with Microsoft products, assumed responsibility for. I would love to wholeheartedly, and without reservation, recommend Microsoft's products.

As it is, Microsoft has been making progress, though, as in the analogies above, has not quite gotten its act together. To really beat these analogies to death, let's say that it has moved from producing Yugos up to making something in the Kia or Hyundai range (ignoring recent ownership changes), depending upon the product under discussion. The problem is that they are using their market position to force prices into the Lexus range, while competitors are selling products that have both quality and prices of a Honda or a Toyota.

Just for review, you may remember that I worked on, but, on ethical grounds, declined to sell any Microsoft Server product prior to Windows 2000, though I did sell all the various versions of their workstation operating systems. Windows XP has turned out to be as reliable as Windows 2000, which means it is not as reliable as it could be, but at least it is a perfectly viable product upon which to base your business.

It has, however, turned out to be an average of about 20% slower than 2000, at least based on the independent test results that I have seen. Given that the one annoying decrease in performance between NT4 and Windows 2000 was in network throughput (how much data it could shove through one or more network cards), that additional drop in performance could really hurt your server. Given that I've been installing multiple 4-port 100 megabit network cards in servers for years, and that Gigabit Ethernet is becoming more common, this is not a trend I am happy to see.

Of course, the overall drop doesn't exactly help workstations, either, but CPU speeds are rising and memory prices are dropping so fast that you might not notice on a new machine. If, however, you upgrade an existing Windows 2000 workstation to Windows XP, be prepared to see a lot more of the ever-popular hourglass.

As to security, just what the hell is going on here? If you know that your product is going to be the most popular one to hack, why make a point of spending proportionately less time and money on security issues than any of your competitors? Making outrageously untrue statements to the public regarding security doesn't really help either.

Let's see, now. Why again, did Microsoft add a feature to XP to allow Windows to lie about its IP address, so as to make denial of service attacks and other related problems almost completely untraceable? Oh, yes, it's because system administrators of Microsoft-based networks demanded it! Yeah, right - and I'm sure there are police officers handing out free guns and bulletproof vests to kids because school administrators demanded it.

Ah, it's so sad to see that little Billy is not living up to his potential. Perhaps a military school could straighten him out (ala, *The Sopranos*).

ODDS AND ENDS

by Wayne M. Krakau - *Chicago Computer Guide*, February, 2002

Wayne is president of Krakau Business Computer Systems, a systems integration firm and has been working with computers for 25 years. He holds an M.B.A. in Marketing and a B.S. in Information Science and was originally educated in Computer and Information Systems Analysis Engineering. He holds CNI and MCNE ratings and all five of the original LANDA Certifications. Wayne can be reached for questions, comments, or topic suggestions at 630-894-3360 or via E-mail to wkrakau@krakau-inc.com. For past columns and company info go to www.krakau-inc.com.

Here are some odds and ends products that aren't appropriate for their own, individual columns. They include some of the oddest ends I have ever seen.

We will start with Levi's Mobile Dockers (www.dockers.com) The first you might have heard of these pants is their commercial, in which a woman plays voyeur using x-ray glasses, and, after snooping on various individuals, ridiculing them as she goes, finally settles on a man wearing Mobile Dockers. She is immediately very impressed and simultaneously embarrassed for getting caught staring.

There is something of a double standard at work here. If the voyeur had been a guy, the end of the commercial would have shown him being led away in handcuffs through a jeering crowd, soon to be added to his state's list of registered sex offenders.

Of course, if the voyeur had really been a guy, the only electronic gadget he would be thinking about would be a digital camera, and, specifically, how he could attach it to the x-ray specs and post the resulting photos on the Web to impress his buddies.

Then, again, there is the geek factor. While the pants look relatively normal while standing, there could be telltale bulges when sitting or crouching. That, of course, would beg the question, "Is that a Palm in your pocket, or are you just being friendly?" Worse, it could beg the same question using "palm" instead of "Palm." Either way, it would be just as bad as the classic pocket protector as an indicator of relative geekness.

On the good side, walking around looking like a refugee from a Dilbert comic with multiple devices and tools hanging from your belt is probably less attractive than having these devices neatly stowed in the various hidden pockets of the Mobile Dockers. (Hey – wait a minute! That sound like me!) In fact, just wearing Dockers instead of some of the other things that I've seen some really hard-core geeks wear is probably an improvement. They lend that nice touch of "business casual" to the ensemble.

Of course, there is the contrarian point of view to consider. If we buy into the claim by some "experts" claim that holding a cell phone near your head fries brain cells, we run into an interesting question. If Mobile Dockers get really popular, should you invest in fertility clinics (assuming the use of headsets, of course)? Perhaps brain cells aren't the only human cells eligible to fry.

If the multi-pocketed Mobile Dockers either don't have enough room or are not compartmentalized enough for you, you can really go for high-tech personal storage with the SCOTT eVEST (www.scottevest.com). (That's the way they spell and capitalize it!) This article of clothing seems to be a direct descendent of the photographers' vests that I so admired back in the days when my 35MM SLR - and its many lenses and accessories - meant more to me than my computer. (Anyone want to buy 6,000 or so slightly used, completely disorganized 35MM slides – stored mostly in their original photo processing envelopes - from a former photo-geek?)

However, they have gone the photographers' vest one better. As well as having 15 pockets, many having specialty features of their own, the SCOTT eVEST has a form of conduit combined with Velcro to provide protected paths for cables to connect your various devices to each other or to headsets and microphones. They call this system a Personal Area Network, or PAN. (FYI: They have a really interesting Web site design, too.)

Please note that if you wear both of these storage-oriented articles of clothing and actually fill the pockets with various electronic gizmos, you will probably not need to actually pass through a metal detector to set it off. Merely passing within 50 feet or so of one should be enough for your personal electromagnetic field to trigger the alarm.

Also, don't forget that although you can remove the eVEST to actually go through a metal detector, removing your Mobile Dockers while waiting in line might be a bit problematical.

Next month I will continue with some additional, though not quite so odd, odds and ends.

Meanwhile, have you heard about the e-mail message that Bill Gates sent to everyone at Microsoft telling them that security is now their number one priority. A cynical point of view (who, me?) might be to see that Bill Gate has a new title at Microsoft – that of Propaganda Minister.

On the other hand, I do have my hopeful, optimistic side, too. If Gates really means it, it could be a breakthrough. If the Microsofties can muster up enough organizational skill to eliminate the bulk of their security problems, the elimination of most of the non-security bugs could be within their grasp!

ODDS AND ENDS - Part 2

by Wayne M. Krakau - *Chicago Computer Guide*, March, 2002

Wayne is president of Krakau Business Computer Systems, a systems integration firm and has been working with computers for 25 years. He holds an M.B.A. in Marketing and a B.S. in Information Science and was originally educated in Computer and Information Systems Analysis Engineering. He holds CNI and MCNE ratings and all five of the original LANDA Certifications. Wayne can be reached for questions, comments, or topic suggestions at 630-894-3360 or via E-mail to wkrakau@krakau-inc.com. For past columns and company info go to www.krakau-inc.com.

Here are some more odds and ends products that aren't appropriate for their own, individual columns. We are getting a bit more mainstream as we go.

The first product for discussion is the Table Tote (www.pctabletote.com), a collapsible platform made to support your laptop computer while traveling. It has a plastic platform and comes with either aluminum or steel adjustable (18 to 36 inches), telescoping legs. The aluminum-legged model weighs 2.6 lbs. and can support up to 18 lbs. while the steel-legged model weighs 3.5 lbs. and can support 25 lbs. The top expands to 13" by 18".

Since the Table Tote is only 13" by 1" by 1" and weighs less than 3lbs. it is quite easy to carry. Using it is a whole lot better than worrying about your aching back or possible carpel tunnel syndrome when traveling with a laptop computer and placing it on whatever open surface you happen to find. Alternately, you can use it for a projector, instead.

My own idea for this device is a bit different. I think that system administrators of networks should buy it for use when doing repairs and/or experimentation. I have found that when working on servers, especially, once you peel the lid, or sometimes just pull the computer out of position to access the cables, you block easy access to the keyboard and mouse.

Typically, you end up using your lap or some piled up boxes for the keyboard and mouse. Often, you must use them standing up because you can't even put a chair near them. None of these techniques is conducive to good health or sanity. They don't really help the accuracy of your typing or mousing either. (That's definitely not a good thing when working on a server.)

The Table Tote could be a cure for these problems. You could set it up next to, or possibly directly over the "sick" computer and gain comfortable access to the mouse and keyboard.

Let's establish some background facts, at least as I see them, for the next product. At least half of the handheld computing devices or PDAs (Personal Digital Assistants) that I have seen in use are not kept merely in skintight leather or leather-like cases. They are carried in zippered cases that are about the size of a thicker than average software manual (as if you could find a software manual these days). Those cases typically hold a small pad of paper opposite the PDA and have room for the stylus as well as multiple cards, pens, and other doodads.

A while back Seiko Instruments (www.seikosmart.com) developed a really great add-on for PDAs, their SmartPad. It duplicated the zippered case, with the added feature of a digitizer located underneath the pad of paper so that, by using their supplied combination pen/stylus, anything you wrote or drew on the pad was reproduced in miniature on the PDA's screen through the PDA's infrared port via a built-in transceiver. That information could be saved, expanded (up to four times) and manipulated on the PDA or transferred to a PC. You could even add color if you had a color-screen PDA.

Well, now the folks at Seiko have one-upped themselves with a new version called the SmartPad2. They added a stylus-based QWERTY keyboard underneath the lower half of the notepad that is about two and a half times the size of a PDA's onscreen keyboard. You need only flip the notepad pages up to access the keyboard. Oh, and just for good measure, they added a phone holder to the left of the PDA.

Speaking on behalf of those of us who are visually impaired (without my glasses, I need to revert to bat-like echolocation - that high-pitched chirping you hear is really me - now if I could only stop nibbling on bugs), the larger surface area of the notepad for drawing and especially the larger keyboard beneath it make a huge difference in the usability of a PDA. In particular, using a PDA with bifocals is really a pain. Every little bit expansion you can get really helps.

The older, non-keyboard model of the SmartPad had two versions, the first supporting the Palm OS based PDAs from Palm, Handspring, Sony, IBM and HandEra and the second, the SmartPad for Pocket PC for Window CE based PDAs from Compaq, Hewlett-Packard, Casio, Toshiba, Audiovox, and NEC. The new SmarPad2 only has a Palm OS version for now, but unless the developers at Seiko are completely asleep at the switch, they should be following up with a Pocket PC version of the SmartPad2 sometime soon.

Now if someone will please pass the salsa, I have zeroed in on what I believe will be a very tasty fruitfly.

WHISTLEBLOWER

by Wayne M. Krakau - *Chicago Computer Guide*, April, 2002

Wayne is president of Krakau Business Computer Systems, a systems integration firm and has been working with computers for 25 years. He holds an M.B.A. in Marketing and a B.S. in Information Science and was originally educated in Computer and Information Systems Analysis Engineering. He holds CNI and MCNE ratings and all five of the original LANDA Certifications. Wayne can be reached for questions, comments, or topic suggestions at 630-894-3360 or via E-mail to wkrakau@krakau-inc.com. For past columns and company info go to www.krakau-inc.com.

OK. So he isn't as handsome as Russell Crowe in *The Insider*, even as purposely messed up in makeup, hair and clothes as Crowe was in that movie (but more on his acting ability, later). He doesn't have the financial acumen of those testifying in Congress against Enron - though I'm not sure that is a disadvantage these days. He doesn't even have the easily recognizable target for his revelations that the classic whistleblower has. There is no specific government body or company - though some would dispute this, he really is an equal opportunity whistleblower - for his disclosures to embarrass. But as with most whistleblowers, he has a virtual target firmly planted on his back, and there is no shortage of potential verbal shooters.

The whistleblower in question is Steve Gibson of Gibson Research (www.grc.com). His original claims to fame were his breakthrough disk diagnostic and repair program, SpinRite, and the column he used to write for *InfoWorld*. More recently, he has become known for security programs such as ShieldsUP! and LeakTest, and especially for his outspoken manor, as he unhesitatingly reports the results of his security research, regardless of whose toes he steps on - hence the target on his back.

Gibson made an appearance at an APCU (Association of Personal Computer Users - www.apcu.org) at the Wilmette Public Library on Saturday, March 16, 2002. Luckily the APCU invited guests for the event, so I was privileged to be able to attend the meeting.

I first encountered Gibson years ago at a distributor's trade show in which he acted out the parts of the components of a disk drive in order to explain how SpinRite works. I'm not sure that even Russell Crowe could have presented it better, and I'm sure that Crowe couldn't have written his own script.

Revealing security weaknesses in software has made him an anathema to development companies, particularly Microsoft, which has attracted security-related complaints almost as if it was asking for them, as well as the makers of personal firewall software. The extreme negative response that he has received for his revelations seems to follow the old 80/20 rule, with 80% of the criticism coming from the point of view of vested interests - from the vendors involved - and 20% sprouting from what seems like nothing less than a massive attack of damaged egos - from various security analysts and experts.

The vendors want to destroy Gibson's reputation in order to diminish what they portray as his "attacks" on them. The analysts and experts don't like someone horning in on their field who, at least on the surface, lacks their "serious" credentials in security.

The news from the dark side is even worse, though it still follows the 80/20 rule. 80% of the hackers, crackers, or whatever you want to call them, are angry because Gibson is trying to get people to take security seriously, which would certainly reduce, if not eliminate their source of fun. The other 20% of them, in different type of ego hit, are greatly offended that he revealed that there are libraries of mix and match kits that allow them to create criminal programs (viruses, trojans, etc.) without any particular skill, thereby robbing them of the "glory" of their exploits. (He calls these "Drag and Drop Viruses.)

The result of these offenses against the criminal element is an ongoing series of DOS (Denial of Service) attacks on Gibson's Web site. Even as he began his speech for the APCU, he explained that word had gotten out that he would be out of town and that his site was currently under attack. Without easy

access to his personal contacts within his ISP (Internet Service Provider), his staff back in the office couldn't effectively ward off the attack.

Gibson was gracious enough to start of the day by participating in the open forum that preceded his speech by supplying some of the answers to questions posed by the members. Naturally, some of the questions drifted toward things like SpinRite and security.

In his speech, Gibson first talked about his early academic and personal history with electronics and eventually, computers. Considering the high level of computer expertise among the APCU members, it was obvious that the audience could relate to his background.

He made his distaste for the old saw "all software has bugs" quite clear, backing it with examples of software that doesn't have bugs. Gibson obviously doesn't tolerate programming incompetence well.

He talked about his differentiation of the use of a PC as either toy or tool. As a toy, you could download and experiment with every untested and potentially buggy or dangerous program you could find. As a tool, you should lock down the PC as much as possible and only run software that you absolutely know is reliable. (The audience contributed several Microsoft jokes at this point.)

Gibson also spoke of his doubts about Microsoft's security proclamations. In particular he noted the declaration by a Microsoft representative that he knew that Windows XP was totally secure because they had run it through a very thorough automated security testing routine and that the testing program had certified the operating system as secure. This was shortly before the weekly (and seemingly unending) security hole announcements about XP started coming through. (Oops!)

He feels that Windows 95 and its descendants were explicitly designed to overwhelmingly favor ease of use (or at least Microsoft's definition of that term, which isn't necessarily what actual users think) over security. While the two characteristics aren't absolutely diametrically opposed, one often precludes the other.

Gibson's final comment was about Microsoft's misguided attempt to defray what amounts to invalid criticism about Windows lack of raw sockets capability (an obscure diagnostic tool that turns out to be absolutely heaven-sent for hackers). Microsoft was told that Windows was wimpy because "real" operating systems have that capability. That's nonsense! Microsoft reacted by adding raw sockets to Windows XP, effectively, just to look cool - only they added it unprotected, while other operating systems restrict access to administrators. We'll see how many millions of dollars companies and organizations lose over the next few years in nearly untraceable DOS attacks facilitated by this added feature.

After that, Gibson headed home so he could get his Web site back up. Such is the life of a whistleblower.

COMDEX IMPRESSIONS

by Wayne M. Krakau - *Chicago Computer Guide*, May, 2002

Wayne is president of Krakau Business Computer Systems, a systems integration firm and has been working with computers for 25 years. He holds an M.B.A. in Marketing and a B.S. in Information Science and was originally educated in Computer and Information Systems Analysis Engineering. He holds CNI and MCNE ratings and all five of the original LANDA Certifications. Wayne can be reached for questions, comments, or topic suggestions at 630-894-3360 or via E-mail to wkrakau@krakau-inc.com. For past columns and company info go to www.krakau-inc.com.

COMDEX Chicago 2002 is a story of good news versus bad news, with the bad news being not really so very bad. This is an overview of my impressions, not a thorough review off the event.

First, note the title change - from Chicago COMDEX in 2001. The organizers are no longer trying to disavow (as the say in *Mission Impossible*) this show as they did last year. Best of all, since they didn't over-hype it this time, they didn't incur the wrath of the vendors, the attendees and the media.

The fact the show was "downsized" and had lower attendance was tolerable in light of the fact that its big brother, Fall COMDEX in Las Vegas, suffered much the same fate. This is not to say that vendors were satisfied with attendance. They just didn't have the same level of anger.

There were some interesting trends to watch at the show. The most apparent was the increase in security-related products, especially those using biometrics. This entails the use of some physical characteristic to establish identity. The fingerprint was the characteristic of choice, though retinas and even voices were represented.

The new twist was the use of smart cards to hold the information that would be matched against the person, such as having the card hold an image of a fingerprint. If you say you own an ID or credit card, you would have to put your finger and the card in a scanner, which would confirm (or deny) a match between your finger and the stored image on the card.

I have become less enthusiastic about this technology since word has come out of its limitations and implications. In one case a bank installed thumb scanners on its ATMs and had to almost immediately yank them because of a rash of muggings that included thumb amputations. Think about what would have happened if they used retina scans!

In another situation, a security testing firm was able to take a cast of a person's thumb, create an artificial thumb, pump a warm liquid rhythmically through it to simulate life, and use it to crack security at a government installation.

I don't think I have to tell you the number of ways a person could hack a system that uses biometrics for identification across a wire (as in remote locations). Such a system could be hacked at the base image, by replacing or redirecting it, or by replacing the bitstream of info coming out of the scanning device.

I love this technology in principle, but, for the moment, I have been disillusioned regarding its ultimate security. It's like a simple alarm system for cars. It discourages the impulse crime by nonprofessional, but won't stop anyone who really knows what they are doing and is willing to make a serious effort.

Actually, the most eye-popping (if you'll excuse the expression) security exhibit wasn't one specifically dedicated to biometrics. It was for the National Safe Skies Alliance. The alarming part of this exhibit wasn't the threat of some draconian measures that someone might suggest for airport security. It was the fact that almost all of the security enhancements that are suggested in this organization's literature are NOT currently implemented at our airports today. (Keep in mind, however, that we are thoroughly protected against World War II veterans awarded Congressional Medals of Honor. One got stopped at an airport because the medal has a pin for attaching to a uniform. You wouldn't want a dangerous character like that to get on a plane with such a deadly weapon, would you?)

Two contradictory trends fit in the good news versus bad news category. Perhaps they might be more accurately described as being in the yin versus yang category. On the west side of the exhibit area there was a booth occupied by Women in Technology International (WITI - www.witi.org). This booth was quite popular - among women, amazingly enough (who would have thought). In fact I never managed to get near enough to get this organization's literature. Only a few steps away, WITI had a dedicated seminar area with a revolving series of speakers. Again, this area was inhabited by - you guessed it - women (how shocking).

On the east side of the exhibit, seemingly as a counterpoint to WITI, was an open booth that initially was a bit of a mystery, since the people exhibiting there arrived late. I happened to be there when a woman and a man arrived, carrying many exhibition-style cases. Shortly thereafter, the booth was absolutely swamped with potential customers, with a long line quickly forming. The product was a type of electric massaging device. You might wonder why the sudden overwhelming interest in such a seemingly mundane product.

Oh, did I forget to tell you that the female half of the pair I mentioned stripped out of her exercise clothing and was wearing an itsy bitsy, teeny weenie, (though not, for you music fans out there who are old enough to remember, either yellow or polka-dotted) bikini? Of course, she was the one offering demonstrations of the massagers. Oh, did I also neglect to inform you that the crowd around this booth was exclusively male? (Will wonders never cease?)

I'll let you decide if, in the grand scheme of things - call it Karma, if you will - these two exhibitors balanced each other out of the equation. (Two steps forward and another two steps back?) Please note that there is no way that I am going to publicize the massager company by either mentioning its name or describing its products in detail.

Next month I'll cover PDAs (Personal Digital Assistants), cell phones as well as some not so trendy stuff.

COMDEX IMPRESSIONS - Part 2

by Wayne M. Krakau - *Chicago Computer Guide*, June, 2002

Wayne is president of Krakau Business Computer Systems, a systems integration firm and has been working with computers for 25 years. He holds an M.B.A. in Marketing and a B.S. in Information Science and was originally educated in Computer and Information Systems Analysis Engineering. He holds CNI and MCNE ratings and all five of the original LANDA Certifications. Wayne can be reached for questions, comments, or topic suggestions at 630-894-3360 or via E-mail to wkrakau@krakau-inc.com. For past columns and company info go to www.krakau-inc.com.

This is the continuation of my impressions of COMDEX Chicago 2002, and I'll cover the biggest and flashiest trend, convergence in wireless communications, of the mobile kind.

Basically, everybody wants to do everything. The people who make mobile phones, usually represented by the various mobile phone service providers, are adding e-mail, Web browsing and text messaging to their mobile phones. The really ambitious ones are going one (major) step further and adding PDA-style (Personal Digital Assistant) computing functions to their phones.

Pretty much all of the phones on display at the show now support text messaging. With the recent standardization of SMS (Short Message Service), you can send these messages to anyone whose phone and provider support SMS. In the past, you could only send messages to other customers of your own provider, and even that feature could be location dependent.

The more sophisticated phones added Web browsing and e-mail. Depending upon the model, you might be limited to Web pages and e-mail that support WAP (Wireless Applications Protocol), a specially abbreviated, text only method of accessing pages and mail. If the original document doesn't directly support WAP, a gateway computer can try to do a translation by stripping off all extraneous material.

WAP harks back to my good news, bad news theme from past month. The good news is, with only text to worry about, transmission is effectively much faster, and less scrolling is needed to view things. If you have tried to see either Web pages or e-mail on any normal-sized mobile phone, you know how important that scrolling factor is.

On the other hand, with WAP some pages and e-mails just don't translate well and their meaning is garbled or partially lost. Also, there are times when you really want to see graphics or tabular info in its original format. With WAP, you don't have that choice.

Optimum use of the inherently limited connection speed of a mobile phone isn't exactly an unimportant consideration, regardless of exactly what type of signaling method is used. If you use a simple dial-up connection via your standard mobile phone transmitter, you will probably get only about 9.6kbps (thousands of bits per second). If you use some type of reserved, out-of-band communications method, such as GPRS (General Packet Radio Service) you might get something in the 30-40kbps range, similar to what you would get on a land line. Note that it is important to state that these are the real, in-the-field, numbers. The oft-quoted theoretical values, especially for the out-of-band methods, are a lot higher than what you really get. (Your mileage may differ!)

The really fancy mobile phones also have PDA functions, ranging from simple address/calendar programs on up to running one of the two most common PDA operating environments, Palm and Pocket PC. Kyocera displayed quite a nice Palm-based unit at the show. It looks like a standard mobile phone on steroids. The keypad flips down to access the Palm keys.

Coming from the other side of the industry, PDAs are adding mobile phone communications. They already have PDAs that access the Web and e-mail, either via internal receivers or through add-on modems. The newest models have added mobile phone functionality.

The flashiest presence at the show - for any product, not just PDAs - was the RIM BlackBerry. (For some reason, I keep telling people about the RaspBerry PDA. It's starting to get embarrassing.) RIM

has turned their popular product into a mobile phone. Just plug in the earphone and start dialing. Naturally you can access the Web and get your e-mail, too.

A big advantage of the BlackBerry is that it has an actual keyboard. Even for SMS messages, a keyboard is a lot easier to use, and a lot more accurate, than either scribbling with a stylus or triple tapping a phone keyboard.

An even better implementation of a wireless convergence device at the show was the HandSpring Treo 180. HandSpring started out with a Palm-based PDA with enhanced applications and added full mobile phone features. Unlike the BlackBerry, this includes a normal phone-type microphone and speaker, positioned as you would expect in a flip-top phone - in this case a really big flip-top phone, though a smaller than average PDA. It's also a speakerphone. Of course, the Palm features integrate with the phone features.

For good measure, the keyboard has a much better feel than the BlackBerry, which is vital to maintain speed when your thumbs block most of the keyboard at any given time. As a Palm-based computer, it can also run any Palm-compatible application - and there are a lot of them out there.

Please note that I am not completely unprejudiced regarding the Treo, as I bought a pair of them for my wife and myself a few weeks after the show. The keyboard is so fast and easy to use that we are both hopelessly addicted to blasting SMS messages to each other and just about anyone else we encounter who has an SMS-capable phone. Even my chunky thumbs really fly over the keyboard.

One really odd sub-trend at the show had corporate sales reps of various phone companies hawking the games that you can play on their phones. I still haven't figured out why, if you specialize in selling mobile phone service to large corporations, you would openly promote a way for your customers' employees to both waste their time and waste valuable mobile phone minutes as well.

What's next? Perhaps they will come up with a line of PDAs that include a dual-books expense account system, one for the real amounts and one for the amounts you report to your company. For maximum marketing effect, perhaps they could license the ENRON brand name.

THE MAGIC CONNECTION - Part 1

by Wayne M. Krakau - *Chicago Computer Guide*, July, 2002

Wayne is president of Krakau Business Computer Systems, a systems integration firm and has been working with computers for 25 years. He holds an M.B.A. in Marketing and a B.S. in Information Science and was originally educated in Computer and Information Systems Analysis Engineering. He holds CNI and MCNE ratings and all five of the original LANDA Certifications. Wayne can be reached for questions, comments, or topic suggestions at 630-894-3360 or via E-mail to wkrakau@krakau-inc.com. For past columns and company info go to www.krakau-inc.com.

Viewing the British TV series *Connections* (and its successors, *Connections 2* and *Connections 3*), originally broadcast in this country on PBS and currently being rerun on *The Discovery Science Channel*, you could learn such arcane information as how the man who figured out how to feed Napoleon's army was responsible for launching the Space Shuttle (depending upon how you and I and the Bills - Gates and Clinton - define "responsible"). In a similar vein, I'm going to show you how a simple magic trick I learned as a child was responsible for solving an ever more common systems integration problem.

As in the TV series, I'll start with the historical background of the issue at hand. In the old days, when dinosaurs (like IBM) ruled the earth, the circuit boards you plugged into the early IBM and compatible PCs were studded with jumpers, those little rectangles that connect two bare pins, completing the circuit between them. The placement of the jumpers determined various characteristics of the board, including what resources the board would use for communications with the motherboard and the CPU (Central Processing Unit). You had to be careful not to bend or, heaven forbid, break a jumper pin.

Early versions of jumpers required the use of hexadecimal (base 16), binary (base 2) or sometimes octal (base 8) arithmetic to figure out the correct settings. The Proteon ProNet 10 card stands out in my memory as having a particularly nasty collection of jumpers. Later, manufacturers broke down and actually documented the settings in comprehensive tabular lists.

Another advance was the replacement of jumpers, at least partially, by DIP (Dual Inline Package) switches. Some of them flip up and down like a light switch and others slide back and forth. Just to make sure that things weren't too easy, for a long time the DIP switches were labeled with a zero or a one or with Open or Closed, assuming they were labeled at all.

Eventually, manufacturers figured out that if they silk-screened the documentation for the jumpers or switches directly onto their boards, they could save a bundle on preventable tech support calls.

All of this jumper and switch nonsense was done under the assumption that mostly hard core computer geeks and computer professionals were the only ones configuring systems and inserting boards into PCs. As computers were being sold to a much wider audience, the incidence of problems rose alarmingly. Boards just won't work properly if the resources that they use conflict with other boards or some resource used by the motherboard.

This could get complicated as, depending upon the board involved, you had to worry about high memory regions (4-digit hexadecimal addresses), low memory region or port allocation (3-digit hexadecimal addresses), DMA (Direct Memory Access) channels and IRQs (hardware Interrupt Requests). Some boards, like sound cards, required multiple, non-conflicting resources.

The IRQs could be particularly troubling. Each IRQ number, from 1 through 16 (only 1-8 before the IBM AT was introduced) corresponds to an actual direct path to one of the legs on the CPU chip. When the CPU receives a signal via one of the IRQ-related legs, it stops what it is doing and immediately gives its full attention to the task associated with that IRQ, as defined by software loaded into the processor. That software defines what port the CPU will look at when it gets an interrupt signal. Under

certain circumstances, IRQs could be shared, but that only made things trickier to figure out. Back then, there was only limited support for IRQ sharing, so running out of available IRQs was always a possibility.

Even with all of these different resources to consider, for an experienced technician or a talented amateur, the task of organizing this mess was definitely doable, if some what tedious (assuming the documentation was up to snuff, of course). An inexperienced user, however, would have to depend upon luck or multiple tech support calls. The problem is, as computer sales grew, the ratio of experienced to inexperienced users tipped overwhelmingly in favor of the inexperienced, nontechnical user, giving computer manufacturers support department fits. The daunting job of manually allocating resources was also seen as slowing sales among the non-geeks.

The answer was a neat trick called Plug and Play, which is a standard for automatically allocating system resources so that software could figure out how to talk to hardware. The idea is to make boards smart enough to supply their resource needs when asked. Then some additional brainpower added to the motherboard figures out who gets what.

That automation is the good news. Next month you'll get the bad news, as well as the rest of the connections.

THE MAGIC CONNECTION - Part 2

by Wayne M. Krakau - *Chicago Computer Guide*, August, 2002

Wayne is president of Krakau Business Computer Systems, a systems integration firm and has been working with computers for 25 years. He holds an M.B.A. in Marketing and a B.S. in Information Science and was originally educated in Computer and Information Systems Analysis Engineering. He holds CNI and MCNE ratings and all five of the original LANDA Certifications. Wayne can be reached for questions, comments, or topic suggestions at 630-894-3360 or via E-mail to wkrakau@krakau-inc.com. For past columns and company info go to www.krakau-inc.com.

We're back now to continue with my own miniature version of the British television series, *Connections*, hopefully minus incidents like the one in which the host, James Burke, was nearly eaten alive by fire ants as he waited, on camera, for a Space Shuttle to launch in the background. Since there was no potential for a second take, he had to just stand there and tolerate the six-legged invaders.

We ended last month with the introduction of a system that was, at least theoretically, going to solve all system resource allocation problems, Plug and Play. This rapidly became known as Plug and Pray, since it didn't always work, and was somehow programmed to know when timing was critical (as with scheduled implementation dates) so as to fail to work properly in just those situations (at least it seems that way to me).

One of the key concepts of Plug and Play is that IRQs (Interrupt Requests) can be shared. This means that you could end up with more than one board assigned to the same IRQ. In practice, this often, though not always worked. Boards do not always play nice and share their toys.

Note that features imbedded into motherboards could use system resources, not just added boards in expansion slots. IDE, floppy and sometimes SCSI controllers, sound and video circuits, and I/O ports (serial, parallel, and USB) may be imbedded in modern motherboards, all demanding the previously discussed system resources.

Now we fast forward to the present to see a nasty trend that is making life difficult for system integrators. That is the tendency by motherboard manufacturers to get so paternalistic about things like system resources that we are insulated from the process so that there are no easy ways for us to manually assign them. Plug and Play isn't so bad to work with if there is some logical way to override IRQs as necessary. Without that ability, things can get very messy. It's like driving a car with mandatory cruise control and no real accelerator.

So, there I was trying to configure a NetWare 6 server with a high-end Intel motherboard specifically designed to be in a specialized server. I go to the CMOS Setup page associated with assigning IRQs and find a row of fixed "logical" IRQ numbers on the left side of the screen and a row of menus for assigning "real" IRQ numbers on the right. There are no instructions. The manual contains only a blurred screen shot of this page with no additional information. Intel's Web site (www.intel.com) has tons of documents related to this motherboard, but nothing on IRQ assignments.

I contacted the computer manufacturer and they didn't have any information on IRQ usage with this motherboard. They were kind enough to contact Intel directly and came back with the following "facts" from Intel tech support:

1. IRQs, in general, don't matter, so don't bother yourself worrying about them.
2. Shared IRQs don't exact any performance penalty, even in a server.
3. IRQ order doesn't matter since they are all the same priority.
4. NetWare has never supported Plug and Play, so IRQs really don't matter with NetWare, anyway.
6. With Windows-based servers, IRQs don't matter.

5. There is no way on this motherboard to assign IRQs to specific boards and/or embedded circuits.

As to item 1, I can only think, could they be even more paternalistic and condescending? It reminds me of the old line that was, (and probably still occasionally is) used on women, "Don't worry your pretty little head about such things." Note that there was no excuse offered for not documenting this page.

For items 2 and 3, I direct you to the Adaptec Web site, www.adaptec.com, or, for that matter, to any add-on board manufacturer's Web site. You will find a plethora of documents relating to the criticality of performance optimization in servers by the proper allocation of IRQs, specifically by avoiding shared IRQs for critical devices like disk controllers and network cards, and by assigning devices based on the numerical priority inherent in IRQs. (The priority order is usually listed as 0, 1, 8, 9, 10, 11, 12, 13, 14, 15, 3, 4, 5, 6, 7. 2 is for cascading to 9 and is not available and 8, the real-time clock, also considered not really available, is sometimes listed as last.)

For items 4 and 5, I direct you to the Novell and Microsoft Web sites, www.novell.com and www.microsoft.com, for even more info on the importance of manually assigning server IRQs using the setup routines in motherboards to avoid both compatibility and performance issues. Additionally, I will note that Novell uses Plug and Play information from individual devices while not actually supporting the motherboard-centric aspects of the standard. It has done this at least as far back as NetWare 4.1. That's how its semiautomatic configuration and setup routines work.

Item 6 is where the Connection comes in. The IRQ setup screen reminded me of a magic trick that I learned as a child. In it, you spread out cards in columns. Then you have the subject mentally select a card and tell you which column it is in. You then gather the cards and spread them out in rows and ask what row the chosen card is in. Depending upon the number of cards involved, the pattern continues until you "magically" announce the subject's chosen card (to thundering applause, no doubt). Actually, this is a simple mathematical game of elimination. By noting the series of answers from the subject, you can figure out the chosen card.

Back to the present, again, I used a version of this magic trick on the motherboard. To prepare, I first discarded any IRQ that I knew I didn't want to use, and freed up any that I could spare. NetWare likes to keep IRQ 15 to itself, so I disabled the unneeded secondary IDE controller that uses it. I wasn't using COM1, COM2, or LPT1, so I disabled them and freed up IRQs 4, 3, and 5, respectively. Then I assigned all of the desirable IRQs in numeric order from top to bottom repeating as necessary to cover all of the logical IRQs and rebooted.

Some cards will tell you their IRQs, but you may have to rely on either the operating system or a utility program to determine other IRQ assignments. I frequently use a DOS-based utility, EMDIAG, from Adaptec. It is available as part of early versions of the driver files for their 4-port network cards. There are other utilities available if you look on the Web.

I wrote down the possible logical IRQs for each of the cards or circuits I was analyzing, based on their reported real IRQs. Then I reassigned the real IRQs on the setup page in a different order. Using logic here can reduce the number of iterations needed to solve this puzzle. I rebooted and recorded the results. This led me to reduce the possibilities on my written list of boards and circuits. Based on this I reassigned the IRQs and again rebooted. This time I was able to reduce the possible IRQs associated with each board or circuit to one, thereby solving the problem.

Now that I knew which board or circuit was associated with which logical IRQ, I could assign real IRQs as necessary for maximum performance and compatibility. It's magic! (Wait here for thundering applause.) And I did it without encountering those pesky fire ants.

BEAM ME UP, HANDSPRING

by Wayne M. Krakau - *Chicago Computer Guide*, September, 2002

Wayne is president of Krakau Business Computer Systems, a systems integration firm and has been working with computers for 25 years. He holds an M.B.A. in Marketing and a B.S. in Information Science and was originally educated in Computer and Information Systems Analysis Engineering. He holds CNI and MCNE ratings and all five of the original LANDA Certifications. Wayne can be reached for questions, comments, or topic suggestions at 630-894-3360 or via E-mail to wkrakau@krakau-inc.com. For past columns and company info go to www.krakau-inc.com.

With their Treo 180, HandSpring, Inc. (www.handspring.com) has introduced one of the first viable convergence (I love that buzzword) devices, combining a cellphone with a PDA (Personal Digital Assistant) and adding Internet access and e-mail to top it off.

I will be the first to admit that I didn't really notice one major characteristic of the Treo until about two months after I bought a pair for my wife and myself. It had to be pointed out to me that, if you activate the speakerphone function and hold the Treo away from your body, it looks like you are using a communicator from the original *Star Trek* series. Aha! I think I've found the ultimate geek badge. I'm sure that Dilbert would be proud. (Actually, a lot of people do stop me and ask to see it.)

In addition to its high-fashion geek appeal, the Treo 180 has a lot going for it. As a cellphone, it comes with an earbud, and AC adapter, a car power adapter, and a built-in speaker phone that is loud enough to almost completely make up for its lack of a full-featured car kit.

The cellphone has the usual quick list of frequently dialed numbers and a history of both outgoing and, via Caller ID, incoming calls. It can also access the underlying address book from its PDA half, even to the point of displaying names instead of numbers on its call history page. You can also copy and paste numbers into it from the notepad, calendar, or other programs running on the PDA.

The main phone screen shows a dial pad with buttons large enough to allow (very firm) finger presses without resorting to the attached stylus. Another screen shows the phonebook and gives that ability to search on strings of characters or numbers. All four screens (quick list, dialing pad, phonebook and history) can be operated using the keyboard optionally in conjunction with the rocker switch on the upper left side of the Treo. The numbers on the keypad are available during a call so you don't have to switch to the number pad screen to negotiate phone systems. It is equipped with a silent mode, vibration, and a screen backlight.

To use the phone in the normal, non-speakerphone mode, you just flip open the lid and hold it up to your head as if it was the big brother of one the smaller flip-top style phones. The microphone is in a tiny hole just below the keyboard. The speaker is near the top of the opened lid.

In this area, your only phone provider choice is VoiceStream. In other areas you must accept Cingular, and in a few places, you can choose one of them. Their rates are quite competitive, though, from my experience, they have a bit less dense coverage than some of their competitors. When comparing notes with people using other providers, I find that I run into more local weak and dead spots than they do. It's usually not that big a deal. It's just somewhat inconvenient.

My favorite, and I will admit, quite addictive, cellphone feature is SMS (Short Message Service) text messaging. There are two important factors at play here. The first is the relatively new protocol standardization and cross-communication among different phone companies. This means that you can send and receive text messages between any phones that follow the SMS standard and are connected to a provider that supports the standard, which includes just about every provider in the U.S. and a good number in other countries, too. For good measure, VoiceStream supports sending and receiving SMS messages directly to and from regular e-mail accounts.

The second factor is one of the major breakthroughs included in the Treo, its full keyboard. Perhaps thumboard would be a more appropriate term, since that's what you use most of the time. You would be amazed at how fast you can type with it. This means that, unlike people sending text messages with normal cellphone keyboards, you don't really have to use cryptic abbreviations and multiple

keystrokes - you just type normally. However, as a courtesy, you probably would like to keep your messages extra short if you are sending to a regular cellphone with a tiny cellphone-sized screen, as opposed the comparatively giant Treo screen. You wouldn't want people to think that, just because you are a Geek (at least in my case) you were rude or something.

When you work with the Palm Operating System-based software, the keyboard really hits its stride. You can enter names, notes, appointments, etc. at a very rapid rate, further enhancing this PDA's ability to replace paper and pen.

HandSpring's standard Palm applications are typically enhanced versions, which will appear in later versions of Palm's own devices. There is the usual Date Book (calendar), the Phonebook, the To Do List, the Calculator and the Memo Pad. As with other PDAs, the Phonebook and Date Book can be synchronized with Microsoft's Outlook, and other software on your desktop computer. Of course, there are lots of additional applications available for downloading.

Next month I will continue covering the Treo 180, including a valuable undocumented feature.

Tolerance Follow-up

A gentleman who was interned with my family during WWII was kind enough to provide me with additional information on the internment of Germans and also a pointer to a Web site with lots more info, www.foitimes.com/internment. Here is a quick rundown of facts that were new to me:

1. My family was held for two years after and many others were held for four years after the war ended.
2. Japanese forcibly deported from Peru were interned in Crystal City, Texas with the Germans.
3. There is a reunion for internees from the Crystal City camp, organized by Richard Santos (richardgsantos@yahoo.com), Chairman of the Zavala County Historical Commission, on November 9, 2002.
4. Many who weren't actually interned were subject to life-destroying government harassment.
5. The most common cause of qualification for internment or harassment was that someone who had a grudge started rumors.
6. For only the second time in U.S. history (the first being during the Civil War) a charge of Sedition was used, for some reason only against Germans.
7. Court proceedings to authorize internment, or to convict people of Sedition or other charges, were held in secret, mostly without benefit of counsel, and sometimes without the accused even being present.
8. People were often taken away in secret so that even their families didn't initially know their whereabouts.
9. A huge number of law enforcement and the military personnel were tied up interning and harassing people when they could have been put to much better use finding real spies or in the War effort, all without stomping on the Constitution.

Are there any lessons to be learned here?

INTEGRATOR'S BLUES

by Wayne M. Krakau - *Chicago Computer Guide*, October, 2002

Wayne is president of Krakau Business Computer Systems, a systems integration firm and has been working with computers for 25 years. He holds an M.B.A. in Marketing and a B.S. in Information Science and was originally educated in Computer and Information Systems Analysis Engineering. He holds CNI and MCNE ratings and all five of the original LANDA Certifications. Wayne can be reached for questions, comments, or topic suggestions at 630-894-3360 or via E-mail to wkrakau@krakau-inc.com. For past columns and company info go to www.krakau-inc.com.

The subtitle of this column could easily be "Mothers, don't let your kids grow up to be systems integrators." To paraphrase an old cliché, "Sometimes you get the computer bear and sometimes the computer bear gets you." I still have bite (byte?) marks from this misadventure which should have taken a portion of one weekend but actually took four full weekends spread out over several months, and included my single longest tech support call ever (eleven and a half continuous hours).

The system administrator at a client of mine and I were going kill a "fun" weekend migrating a NetWare 5.0 server over to a new NetWare 6 computer. I had taken the original "Breakout" NetWare 6 class for instructors, but she had only a quickie sample NetWare 6 seminar. I knew things weren't going to go well when I found out that she hadn't done any of the preparation work that I had suggested.

It turns out that during a tech support call to Novell earlier that week, the technician contradicted all of my advice and declared my suggestions superfluous. Given a choice of applying potentially disruptive general patches and eDirectory (formerly called NDS or Novell Directory Services) updates to a production server, and just leaving things be, she made the obvious choice. I couldn't blame her for taking a Novell technician's word over mine. I could feel bear's hot breath on the back of my neck as he approached.

We ran Novell's NWDEPLOY program to prepare the old system for the migration. It automatically updated the eDirectory; something that I had been warned was better done manually. We then ran the Migration Wizard to actually move data and properties from the old server to the new server. The data transfer went well, though it took many hours.

The movement of properties, which included the eDirectory, seemed to go well, but left us unable to run any tools to access, alter, or even check the validity of the eDirectory. Every program we ran, such as DSREPAIR (the Directory Services repair program) would fail and give us a mismatching version error. This error was quite similar to the old DOS error that you would get from running say, the DOS 5.0 CHKDSK command on a DOS 6.0 system.

After a couple of hours of experimentation, we gave up and called Novell, paying for a tech support incident. The first thing they asked was the eDirectory version on the new server. When we told them, they didn't believe us! They said there was no such version and talked to us like we were either completely inexperienced or perhaps just a few transistors short of a CPU. This argument went on for quite some time until one of their technicians linked up and remote-controlled the new server. That finally convinced them that we had a nonstandard, never-released-to-the-public (at least in theory) version of eDirectory.

As far as we could tell, though based totally on conjecture, some developer had accidentally left in a line of code that stated, in essence, "If the following conditions are true, then this must be my special test server, and it should be loaded with this special, experimental version of eDirectory instead of the standard version." Somehow we met those conditions, though, from subsequent inquiries on my part, we may have been the only ones to ever run into this since-fixed bug.

After spending several hours trying to get the new server into a usable state we finally had to give up since it was already Monday morning and employees would need a live server within the hour. The

technician assisted us in reviving the old NetWare 5 server and apologized for the problem. Chomp! That bear took quite a piece out of me.

A few weeks later, we tried again. This time we were at the client's new location where the air conditioning was not yet installed in the computer room, so conditions were quite miserable. This time we decided to get creative (always a dangerous idea) and back up some of the auxiliary volumes onto tape so that the migration would go faster. We planned to restore those volumes while we did the post-migration installation of updated applications and the reconfiguration of the many printers this client used.

This time the deployment and migration went so well that we didn't have to call tech support at all. We were so proud of ourselves. We cranked up Veritas Backup Exec, inserted the appropriate tape, and watched while the server promptly froze. The bear already had one paw on my shoulder.

After reconfiguring the system every way we could, including even swapping out parts where possible, we finally had to give up again. We had hours of updating, reconfiguration, and subsequent testing to do and it was already late Sunday night. We couldn't risk not having a viable server in the morning.

The administrator was convinced that I had sold her a defective server, and I was worried about another mysterious NetWare bug. It wasn't until several weeks and many hours of research later that I found a brand-new message in the forums on the Veritas Web site that indicated that there was a bug in Backup Exec which causes a complete server lock-up if you try to back up information from a NetWare 5 server using traditional volumes to a NetWare 6 server using the newer NSS (Novell Storage Services) volumes, a bug which still exists. Chomp! There went another hunk of my backside. Boy, that bear is hungry.

Next month I'll cover yet another lost battle with the bear as well as our final victory over him. Meanwhile, could somebody please call a Medic?

INTEGRATOR'S BLUES - Part 2

by Wayne M. Krakau - *Chicago Computer Guide*, November, 2002

Wayne is president of Krakau Business Computer Systems, a systems integration firm and has been working with computers for 25 years. He holds an M.B.A. in Marketing and a B.S. in Information Science and was originally educated in Computer and Information Systems Analysis Engineering. He holds CNI and MCNE ratings and all five of the original LANDA Certifications. Wayne can be reached for questions, comments, or topic suggestions at 630-894-3360 or via E-mail to wkrakau@krakau-inc.com. For past columns and company info go to www.krakau-inc.com.

This is the continuation of my adventures in migrating from a Novell NetWare 5 server to a new NetWare 6 server. At this point we have suffered through two marathon weekends, several weeks apart, thwarted by a combination of human errors and multiple software errors in both NetWare and Veritas Backup Exec. This exercise has descended well into the "Sometimes the bear gets you" category.

Now that we knew that the attempt to back up a NetWare 5 traditional volume and restore it to a NetWare 6 NSS (Novell Storage Services) volume (which I unaccountably misnamed Novell Storage System last month - oops) caused our latest failure, we attempted another migration a few months later. This time we would migrate all of the volumes, no matter how long it took.

I had the idea to duplicate the existing production NetWare 5 server and then use the duplicate for the migration, thereby relieving us of the necessity of completing the job during one weekend. We could start the migration earlier in the week while the production server was still running, and then update the new server with any changed data later

After spending many hours in repeated attempts at duplicating the production server, the administrator ran into insurmountable bugs in Backup Exec that were confirmed by Veritas. So much for sneaky ideas. The bear took another swipe at me.

On our third migration attempt, we ran into a problem with server-to-server communications. The two servers became confused as to whether to talk IP or the older IPX. After consulting with Novell, this time using one of my dealer technical incidents, we elected to set both servers to talk only IPX. The technician reminded us the while IP was appropriate for larger networks, IPX could easily outrun it on a small network like ours, which included only the two servers and one workstation. The prediction turned out to be true, as we migrated all of the volumes using IPX in less time that it took for just the first two using IP. We were quite pleased with ourselves.

Then the bear growled again. The eDirectory (formerly NDS) had obsolete information on users, printers, and queues, and the file system had obsolete rights assignments, all left over from the previous migration attempt. This time the folks at Novell were sympathetic and let us continue using the same support incident.

With Novell's help, we figured out that the Migration Wizard, the workstation-based program that controls the various steps involved in the migration process, had failed to automatically delete the files that it uses to temporarily store eDirectory and trustee (file rights). That's why we got the old information. This bug has since been fixed.

By the time the Novell technicians had figured out how to manually erase the obsolete information and transfer up-to-date information to the new server, eleven and a half hours later, it was again very late on Sunday night, so we had to give up and go back to the old server again. Thus, ended my longest support call ever. The bear had taken another bite out of me.

As we reconvened a few weeks later for yet another migration attempt, I had the distinct feeling that, if this try failed, the system administrator would see to it that I wouldn't be leaving the building by either the elevator or the stairs. I knew that I had the opportunity for only one more battle with the bear.

Finally, the migration ran cleanly. Then we started the post migration tasks, including testing printers, installing and testing new software versions, and testing the preexisting software.

In his last gasp for life, the bear launched one final attack. This client owned Tobit's David, an e-mail system that also included Tobit's network fax system, FaxWare, but they only used the FaxWare half of the product. Now, FaxWare didn't work.

We inspected the new server and realized that some undetermined critical files were copied to an archival directory, not to the one where they started. Rather than experimenting to try to fix the problem, we saw this as an opportunity to implement our plan to move the fax system to a communications server. After installing David (along with FaxWare) on the communications server and moving over all of its data, we ran into another problem. While FaxWare actually sent and received faxes, all of its phonebooks and fax archives were inaccessible. After many attempts at fixing the problem, we finally had to give up.

The system administrator called Tobit the next day and they walked her through reinstalling David on the new NetWare 6 (the migration target) server. It worked, but that didn't solve our long-term problem of moving the fax system to the communications server.

That's when they told her that they had purposely designed FaxWare so that it could NEVER be moved to another server. The original server name is embedded in each and every page of every incoming and outgoing fax as well as every individual phonebook record. FaxWare will not acknowledge the existence of any file with a different server name, so moving the system to a different server is not practical, unless you don't mind giving up your accumulated fax history and your phonebooks. I do hope that whoever thought of that bright idea is no longer in the computer industry and, therefore, can't do any more damage.

Well, the final score was the bear 3 and us 1, but the final victory was ours. I do, however, think that the bear died with a smile on his face, knowing that he left behind one final, and ultimately unsolvable, problem. Meanwhile, maybe I can find a job in a less stressful field, like lion taming.

INTERNET SLAMMING

by Wayne M. Krakau - *Chicago Computer Guide*, December, 2002 (unpublished due to demise of paper)

Wayne is president of Krakau Business Computer Systems, a systems integration firm and has been working with computers for 26 years. He holds an M.B.A. in Marketing and a B.S. in Information Science and was originally educated in Computer and Information Systems Analysis Engineering. He holds CNI and MCNE ratings and all five of the original LANDA Certifications. Wayne can be reached for questions, comments, or topic suggestions at 630-894-3360 or via E-mail to wkrakau@krakau-inc.com. For past columns and company info go to www.krakau-inc.com.

Internet Slamming is my term for a scam in which people or businesses are charged a fraudulent fee on their phone bill for some combination dial-up Internet access, e-mail, Web site hosing, and online phonebook-style listings on a vendor's Web site.

This new scam is closely related to the better-known fraud of Slamming, the switching of long distance services either completely without permission or through the use of some type of trickery or deceit. It is one member of a group of scams the telephone industry calls Cramming that involves sneaking charges into the miscellaneous section of a phone bill.

I was recently Internet Scammed. The first sign was an unusually high total on my business phone bill. It included a charge of $33.10 from ILD Teleservices, Inc. that I traced it back to the miscellaneous section of the bill where I found a $29.95 charge from National Online Services plus various taxes.

SBC reps told me that they are not allowed to question this type of bill (not true). They also stated that any action they could take on a customer's behalf would be considered anti-competitive behavior by legislators and regulators (also not true). All of this was in spite of the fact that they accepted my assertion that the charge was fraudulent. They referred me to ILD.

ILD claimed that they are only an intermediate billing company and have no real association with National Online Services. They represented themselves an innocent third party and referred me to National.

I got enough information from National to figure out how they got my name. A couple of months earlier, I received a call offering a company listing on their Web site (which I am purposely NOT going to mention - let them get their own publicity) bundled with Web site hosting and dial-up Internet access. I explained to the salesman that my company actually sold telecommunications and Web hosting services so that the listing was the only thing that applied to me. He offered me a free 30-day trial for only the listing. I talked to three different representatives of National and they all assured me that the offer was self-expiring and that I would only have to take action if I wanted to keep - and pay for - the listing. Since already I have a listing at a (legitimate) competitor's site, I didn't suspect anything. When I later inspected National's site and found it to be technically deficient and that my company listing was inaccurate, I declined to sign up for their services.

After a lot of arguing, I managed to get National to cancel my contract, but they wouldn't do it retroactively, so I as warned by both ILD and SBC that I would be out another two months of bogus charges.

I started searching the Web for info on this scam and was absolutely overwhelmed by the number of complaints against ILD. I tracked down a law firm, Horwitz & Associates (coincidentally in Chicago) that was gathering information on ILD for a potential class action suit and got additional clues as to the nature of this scam.

A company sets itself up as a billing agent for many small service companies. It then accepts all bills without question. The lawyers try to prove that, despite protestations to the contrary, the billing company is actually in collusion with if not the outright owner of the service companies. If they can prove it (which is no easy task), then they can file a civil RICO (Racketeer-Influenced and Corrupt

Organization) suit. Since the law has fallen so far behind the technology, the lawyers are having difficulty finding consumer law violations, especially in third party relationships (involving customer, billing company, and service company).

As further evidence of intent, many of the people who are being hit with Internet access fees don't even have computers, including some who are told they signed up via the Web! A good number of victims are elderly or disabled. There are many cases of people who can't even contact the vendors or the billing companies to try to cancel. Their bills just keep coming, month after month.

Amazingly, my next SBC bill not only lacked additional bogus charges, it listed a credit for the previous false charge! I found out that, because phone companies give a bit more flexibility and autonomy to their business line (as opposed to home line) representatives, a few business owners who complained enough (like me) received credits. However, as in the case of having a credit card company assist you in protesting a bill, the underlying vendor can protest right back and reinstate the dropped charge. This means that I still could get stung additional charges (something like the nearly unkillable murderer in slasher films).

What can you do to avoid Internet Slamming? Just as in protecting yourself from viruses, many of the suggestions are somewhat obvious, but only if you are warned of your vulnerability. Here are my suggestions:

1. Consider blocking the miscellaneous section of your phone bill. It is similar to blocking the long distance section. Your bill can't be changed except by a call from you on your phone. There are some small disadvantages to this technique, so check with your phone company first.
1. Read your phone bill carefully every month and immediately report anything suspicious.
2. Do not engage telemarketers in any conversation. If they have a recording of you saying "Yes," or a synonym thereof, they can splice in a question in front asking if you want to accept their services, and then use this as "evidence" to trick your phone company into believing in the validity of their charges. Typically they ask you to confirm some innocuous information, such as your address, to get your response, without necessarily warning you that you are being recorded. Unless you are willing to hire attorneys and forensic audio experts, you won't be able to counter their manufactured evidence. Warn anyone with access to your phones, including employees and family members, about this.
3. Don't enter any contests unless you are willing to hunt down and read all of the fine print. Contests are often fronts to gather information for Internet Slamming and other scams. This includes contests via e-mail, faxes, Web pages (especially pop-up ads), and even at shopping malls.
4. If someone contacts about winning a contest, assume they are lying scum (see above) and cautiously move forward from there (assuming, of course, you don't see your own front door on TV). If it is a legitimate contest, they won't mind restricting all further contact to your attorney.
5. If you become a victim, use the Internet to gather information to use in fighting back. Whois is a particularly useful tool. (Search for "whois" using your favorite search engine.) When a company refuses to give out their info (as National Online Services did with me), whois will find the Web site owning company and its main contact, phone number and street address. You will need this if you want to try sending a registered letter of cancellation or even a subpoena. General searches for a company might also yield additional consumer or legal information.

The irony in this is that ILD is a founding member of CERB, the Coalition to Ensure Responsible Billing! This association of third-party billing companies has a list of rules designed to avoid fraudulent charges, which, based on my Web searches, are all violated by ILD. The real nature of CERB is revealed

when you read its filings with various government bodies politicking for less “restricting” rules for its own industry.

In the end, the best advice I can give is “Let the non-buyer beware.”