InfoWorld

GET TECHNOLOGY RIGHT

IT Strategy Guide

Managing Enterprise Architecture

INSIDE

Introduction

IT BUZZWORDS MAY COME AND GO, BUT BEST practices built into extensible frameworks endure on all fronts, be it security, networking, or application development. Methodologies evolve, improve, and expand or contract to meet the demands of ever-changing business climates.

Deciding which technologies will best support your IT infrastructure can be daunting, but a solid foundation that rests on reliable platforms and products is the cornerstone of every successful IT initiative. Increasingly, the components that make up a sound IT architecture are interoperable, which makes for easier access to information across a wide set of applications and data repositories.

Responsive, flexible, secure

This guide is designed to help you build a sound architecture with airtight security. It also aims to help ensure that what you implement today can be expanded and repurposed as your enterprise grows. This approach will help you avoid wasteful spending as a result of "rip-and-replace" tactics, which can also disrupt essential workflows.

There's no substitute for experience. *InfoWorld* is honored to share with you a wealth of resources and lessons learned as you pursue your enterprise goals.

ITIL: IT by the Book

In the 1970s, when the American auto industry found itself under attack by leaner, hungrier Japanese competitors, it fought back by adopting some of the very production processes the Japanese had pioneered. Using techniques such as statistical process control, quality circles, just-in-time inventory management, total quality management, lean manufacturing, and Six Sigma, the industry focused on improving how its people worked and how its processes operated. For example, workers were encouraged to stop the assembly line when anything went wrong so the process could be fixed permanently, rather than simply scrapping rejects at the end of the line.

Today, American IT organizations are at a similar crossroads, facing challenges from offshore outsourcers and from internal financial pressures. In response, they're stealing a page from their global competitors' playbooks — a process framework developed in the United Kingdom called ITIL, or IT Infrastructure Library.

Like the CMM (Capabilities Maturity Model) for application development, ITIL is a set of best practices and standard methodologies for core IT operational processes such as change, release, and configuration management; incident and problem management; capacity and availability management; and financial management for IT. Although the datacenter is ITIL's primary target, its best-practices templates apply across almost every IT environment, from the service desk to the corporate desktop.

ITIL adoption is growing like a weed. Four years ago, ITIL was already in high gear in Europe, but almost no one in the United States had heard of it. Today a rapidly growing North American industry of consultants, conferences, and training resources is spreading the ITIL gospel and helping customers implement it (see "ITIL Resources," right). "We can't keep up with the demand from organizations like the big airlines, government departments, and banks and insurance companies," says David Ratcliffe, CEO of Pink Elephant, an ITIL consultancy based in Toronto.

"It's spreading like wildfire across large U.S. companies," says Kathryn Pizzo, a group program manager at Microsoft's consulting division. That growth raises some interesting big-picture questions about the future of IT: Could ITIL, with its concept of a CMDB (centralized configuration management database), be the catalyst for the widespread realization of utility computing? And will success with ITIL hinge more on automating processes, as vendors would like us to believe, or on getting human beings to work more efficiently?

Playing Catch-Up

In a sense, ITIL is nothing more than a reincarnation of the stringent management processes that evolved in the mainframe world before the proliferation of PCs, client-server, and Web-based architectures made those operational disciplines seem anachronistic.

So why is ITIL taking hold in the U.S. now after being virtually ignored for years? For one thing, its two biggest benefits — improving service and reducing costs — are right in line with the new marching orders IT organizations got when the economy slowed down in 2001. "It's really about moving from being technology-centric to being services-centric," says Regina Kershner, director of IT service management at HP Services.

Second, corporate mergers and outsourcing have made it more important for IT shops to speak the same process language. "ITIL provides a common language so you can work more effectively with your outsourcers for end-to-end service delivery," Gartner Research Director Steve Bittinger says. Bittinger thinks that ITIL was slower to take off in the U.S. because American corporate culture is more entrepreneurial and less process-oriented than Europe's.

Finally, the sheer scale and complexity of today's IT operations — including the need to have a better handle on IT processes in order to conform to standards such as COBIT (Control Objectives for Information and Related Technologies) — demand the replacement of cobbled-together, homegrown processes with standardized, disciplined ones based on ITIL.

Nationwide Mutual Insurance Company, one of the early U.S. adopters of ITIL, is a case in point. In 2001, Nationwide revaluated its IT operations processes and realized they desperately needed overhauling. Incident and change management were major pain points. Poor communication and fragmented "tribal knowledge" were widespread. "It was borderline crisis," explains Doug LeMaster, director of IT program management. "Processes were ad hoc [and] customer expectations weren't being met."

After re-engineering key processes based on the ITIL framework, Nationwide saw major improvements in systems availability, Nationwide's IT Process Officer Jack Probst says, estimating that downtime decreased by 50,000 user minutes. What ITIL did, he explains, was help standardize the language, process, and workflow of key operations. "ITIL provided the behavioral disciplines necessary to make it happen," Probst says.

Getting Started With ITIL

At its core, the ITIL framework is a set of 44 books originally published by the British government's Stationery Office between 1989 and 1992 — available on the IT Service Management Forum's Web site (infoworld.com/1935) — each dealing with a different operational process. The framework can be implemented in stages and most experts recommend a phased deployment.

Many companies have also turned to larger consulting organizations such as HP Services and IBM Global Services to provide training and packaged ITIL offerings, including suggested workflows, to help customers quickly get up to speed. These consultants also provide needs assessments and benchmarking to help customers determine how they're doing. "There are no metrics police in ITIL, no independent British government-sponsored metrics for how good you are," notes Pink Elephant's Ratcliffe. But there is an official, all-or-nothing ITIL certification process called BS 15000, overseen by the British Standards Institute.

Once committed to the ITIL framework, companies must decide which technologies will best support their process re-engineering road map. Despite ubiquitous vendor claims of "ITIL compliance" and "ITIL compatibility," the framework is technologically agnostic — it stops short of prescribing technology standards. In fact, a basic level of ITIL can be implemented with almost any technology, even spreadsheets and — believe it or not — paper. "There's no rocket science here," says Ram Duraiswamy, vice president of IT Governance Strategies at Mercury Interactive.

Microsoft's Pizzo concurs: "It's about the processes, not the products. For example, whether they have any communication mechanisms in place in the datacenter to alert all the stakeholders that a change is in the pipeline, what the approval stage is, etc.," she explains, emphasizing that you "don't necessarily need to buy anything."

Nonetheless, vendors have plenty to sell. In an attempt to ride the wave, they are building the ITIL taxonomies, workflows, and language into their products and reference models, while touting pre-built integration as an ITIL enabler. Microsoft, for example, has created the MOF (Microsoft Operations Framework), its own ITIL adaptation, and IBM and HP have similar schemes. "Our angle is to bundle these processes in with our tools and products, instead of expecting people to do stand-alone process improvement projects," Pizzo says. In theory, an MOF deployment could automatically catch an error in Exchange Server, generate a trouble ticket, and automatically transfer that ticket to a Remedy help desk system.

But customers say they're far from achieving full ITIL automation. "There isn't just one tool that does it all," Don McGinnis, an IT staff manager at State Farm Insurance says. State Farm implemented several ITIL processes using a variety of tools, including HP OpenView Service Desk, HP Network Node Manager, and a legacy mainframe automation tool called CA-OPS/MVS.

But McGinnis wishes there was a single tool, citing incident management as an example. "With incidents, you try to recover whatever's not running, then go to a review group that looks at it from a problem management point of view and digs down into the root cause," he explains. "Then you go into change once you've identified the problem, schedule it, get agreement, and then pass it on to release. That's how they all work together, that's why you need the tool. ... The more ITIL processes you can put under one tool, the better off you're going to be."

Nationwide's Probst recommends assessing your current and desired processes before selecting tools, as some do better than others at supporting the nuts and bolts of industry-specific, highly customized business processes. "If you select the tool before you have your processes down, you're hosed," he says.

ITIL Resources

To learn more about ITIL, try these links to public education courses and events.

InteQ offers professional services classroom training (infoworld.com/1906).

ITIL has a global homepage (infoworld.com/1924).

IT Services Management Forum is a nonprofit organization dedicated to advancing best practices (itsmf.com).

Pink Elephant is an ITIL and business strategy consultancy (pinkelephant.com).

U.K. Office of Government Commerce provides further ITIL information (infoworld.com/1925).

CMDB: The Black Belt of ITIL

Unlike Six Sigma, ITIL doesn't have a "black belt"-level designation for its most advanced practitioners. But if it did, it might very well go to those organizations that have successfully implemented ITIL's configuration process utilizing a tool called the CMDB (configuration management database).

The CMDB, which the ITIL framework describes only conceptually, is a comprehensive master database describing all IT infrastructure components in a given environment and how they relate to each other, who owns them, what incidents are related to them, and so on. In its most sophisticated incarnation, a CMDB is similar to what a nerve center might look like for truly autonomous utility computing.

"It's a big undertaking," says John Long, Technical Strategist at IBM Tivoli's Technical Strategy Organization. "Companies have to start small and then begin to link in lots of related data that they may already have." Numerous vendors offer pieces of the puzzle, including traditional systems management vendors (CA, HP, IBM), traditional service desk vendors (Remedy, Peregrine), and smaller, niche companies, including Austin, Texas-based Troux Technologies and London-based Tideway Systems.

Vendors are leading the push to automate pieces of CMDB functionality, including asset discovery and inventory management. "Once you identify an event that spawns an incident, you want your infrastructure to begin reacting to that; there's an automated workflow that should be there," IBM's Long says. "Most customers focus on the human workflow ... but increasingly they're seeing the need for a more autonomic approach for solving incidents."

Tim Howes, CTO of Opsware, agrees. "If you try to build a CMDB by hand, you end up with some pretty serious problems in terms of information accuracy. ITIL does a good job of flushing out each of the processes, but it doesn't define a way to enforce the process, to ensure that the systems and information on which those processes depend are accurate and up to date," Howes says.

Peregrine Systems Vice President of Product Marketing Craig Macdonald thinks ITIL configuration and change management processes should be built around a single integrated technology platform to avoid ad hoc workflow customization and to enable closed-loop capabilities that would go beyond the ITIL specifications, such as auditing of changes after they've been made.

"Change management tends to be very cross-functional in nature, requiring very complex workflows [and] approval processes," Macdonald says. "When you start requiring customization to the workflows that come out of the box with many vendor products including our own, you're adding complexity. What companies tried to do in the late '90s and early 2000s was take out-of-the-box change management solutions and integrate them with other technologies to create a more robust process. But many such implementations failed," he adds.

Larger vendors worry that ITIL will create a standard framework that would encourage shopping for best of breed vendor components instead of bundled solutions. "ITIL tends to help the smaller guys," says Opsware's Howes. "You take away the proprietary advantage that IBM has, all [its] different technologies that implement the processes you need. If you have a standardized process it's easier to mix and match."

Some vendors are hoping that the standards authorities who control ITIL — namely the British Office of Government Commerce — will get more specific about how to implement it. "We often wonder about how the ITIL standards are going to evolve," says Mercury Interactive's Duraiswamy. "Most of it is still at 10,000 feet and above. They leave a lot to interpretation."

— *David L. Margulius*

The New Security

In today’s era of perimeter-invading worms, malicious e-mails that don’t rely on attachments, and tenacious spyware, safeguarding the enterprise demands a security framework that marshals a more sophisticated combination of technologies. A traditional firewall and an up-to-date virus scanner may no longer be enough.

But what is enough, exactly? Getting a handle on which solutions to deploy — and where — has become increasingly difficult, as emerging technologies have begun to overlap and functionalities have merged.

For example, if your new firewall can block application-layer threats, do you need an intrusion detection system? Should you choose a rules-based IDS, or one that uses anomaly detection to flag zero-day attacks? And when should you consider host-based security measures, or a specialized application security solution?

Naturally, the answers to these questions depend largely on the value of the assets you’re trying to protect. In any case, it’s critical to keep an eye on the changing landscape of point solutions. By keeping abreast of security advances, you’ll be better positioned to capitalize them before newly evolving threats infiltrate your enterprise.

Firewalls and IDSes

Firewall vendors such as Check Point Software Technologies and Juniper Netscreen are touting new application-layer filtering capabilities, and these are important advances. After all, if your firewall is intelligent enough to block a DoS attack or a NetBus Trojan probe, you can rest so much easier.

Nevertheless, compared to a well-tuned IDS, even the most modern firewall is a blunt instrument — and necessarily so. A stateful inspection firewall is an effective way to block unauthorized port traffic, defend against IP address spoofing, and thwart other, more recent types of attacks. Proxy firewalls, which prevent direct connections to hosts inside the network, provide yet another layer of protection.

But all firewalls have holes, if only because they must remain open to legitimate traffic. They can’t inspect the contents of point-to-point VPN traffic, and even those that do make application-layer decisions can identify only a narrow range of threats that ride almost universally welcome protocols such as UDP and HTTP.

More and more malicious attackers are using port 80, which is almost always open between segments. In fact, if I were a malicious coder, I’d look first to port 80 — or another commonly opened firewall port — in order to gain entry to a network. To counter this, you need the data-level inspection that only an IDS or IPS can provide.

Detect or Prevent?

Because they can prevent malicious exploits, IPSes are

Location Is Everything

Company size is an essential factor in deciding where best to deploy your security solution.

Solution	Small business	Midsize business	Large enterprise
Firewall	Network edge/host	Network edge/host	Network edge/host
IDS/IPS	N/A	DMZ/internal	DMZ/internal
Anti-virus	Host	E-mail/host	E-mail/gateway/host
Anti-spyware	Host	Host/gateway	Host/gateway
Network quarantine	N/A	Remote/VPN	Remote/VPN/internal
E-mail filtering	Host	Gateway/host/ASP	Gateway/host/ASP

outpacing IDSes as the preferred security systems of choice. After all, if an IPS can prevent an attack, why would you ever choose an IDS instead?

The problem is that many, if not most, IDSes and IPSes suffer from high percentages of false positives. And, whereas an IDS will only log a false positive, an IPS will block traffic marked as potentially dangerous, thereby preventing a significant amount of legitimate traffic from entering your network. Although vendors are working on improving accuracy, accidentally denying legitimate traffic can be even more catastrophic to your business than failing to block a malicious attack.

IDSes and IPSes are the best solutions for preventing buffer overflows — second to patching, that is — and for recognizing abnormally constructed data. IDSes and IPSes excel at inspecting packet data and lower-level packet information, which is what makes them so effective at identifying threats. But whereas IPSes will block identified threats, IDSes simply alert administrators after identifying malicious traffic.

IDSes and IPSes are located either at a network filtering point — to identify threats passing between networks — or on a host computer. Host-based IDSes and IPSes are designed to protect only the host on which they are located. While network-based IDSes and IPSes will identify — and in the case of an IPS, stop — general threats, host-based solutions are configured to protect systems against malicious attacks targeting specific operating systems or application software. For example, an IPS for Microsoft SQL Server will be designed to prevent SQL injections and guesses at database passwords.

IDSes and IPSes use two types of technology, the most commonly employed of which is fingerprinting. Fingerprinting, aka pattern-detection, solutions work much the same way anti-virus scanners do, that is, making use of databases that store predefined malicious byte patterns to identify specific threats. Perhaps because this approach is the most popular, fingerprinting databases must be constantly updated; they can be defeated by new and slightly modified threats.

The second technology, anomaly detection, uses baseline profiling to recognize statistically deviant traffic patterns. For example, an anomaly-detection solution would flag high levels of sustained network traffic originating from a low-traffic host, or it would notice unauthorized manipulation of system files. Anomaly-detection engines are useful for detecting zero-day or slightly modified exploits.

Which technology should you choose in an IDS or IPS? Ideally, as *InfoWorld*'s reviewers discovered in a recent

Rx for Security Firewalls and IDS and IPS devices cover a lot of terrain, but plenty of stand-alone solutions can fill in some of the gaps.

Yes ■ No □

Threat	Anti-spam	Anti-spyware	Anti-virus	Firewall	IDS/IPS	Network quarantine
Application-layer attacks	□	□	■	□	■	□
ARP poisoning	□	□	□	□	□	□
Buffer overflows	□	□	□	■	■	□
Cross-site scripting	□	□	□	□	■	□
Dedicated attacker	□	□	□	■	■	■
Directory transversal	□	□	□	□	■	□
DoS/DDoS attacks	□	□	□	■	■	□
E-mail malware	□	□	■	■	■	■
Man-in-the-middle attacks	□	□	□	□	□	□
Network-layer attacks	□	□	□	■	■	□
P-to-P attacks	□	□	■	■	■	□
Packet sniffing	□	□	□	□	□	□
Password cracking	□	□	□	□	■	□
Phishing	■	□	□	□	□	□
Physical compromise	□	□	□	□	□	□
Port scan	□	□	□	■	■	□
Protocol anomalies	□	□	□	■	■	□
Session hijacking	□	□	□	■	□	□
Spam	■	□	□	□	□	□
SQL injection	□	□	□	□	■	□
TCP/IP spoofing	□	□	□	■	■	■
Web/HTML attacks	■	■	■	■	■	□
Wireless attack	□	□	□	□	□	■
Worm, virus, or Trojan	■	■	■	■	■	■

product roundup (infoworld.com/1826), the solution should contain both components. Most threats have an easily recognizable byte pattern, but anomaly detection should be layered on top of this pattern-detection capability in order to discover threats that do not have a specific signature and to stop zero-day exploits.

Early adopters placed IDSes and IPSes outside the firewall or on the DMZ to complement external security defenses. Unfortunately, reported events quickly overwhelmed administrators. Today, IDSes are deployed inside the trusted network as an early warning system to notify administrators when the perimeter has been compromised. Host-based IPSes, such as Sana's Primary Response (infoworld.com/1827), may be the most effective way to lock down specific Web and database servers.

Anti-Virus

Anti-virus technology isn't a panacea. Traditional anti-virus scanning solutions that use pattern databases have been great at detecting already-known threats but have proved terrible at dealing with zero-day exploits, slightly modified one-off malicious programs, and buffer overflows. The SQL Slammer worm infected tens of thousands of computers in less than 10 minutes. If anti-virus scanners can't defeat zero-day attacks, how can they be expected to deal with zero-minute attacks?

Vendors have responded by incorporating heuristic scanning tools and by submitting more frequent database updates. Similar to anomaly-detection technology, heuristic scanners analyze files by looking for coding actions often related to malware, such as modified executables, self-contained SMTP engines, and writing to sensitive registry areas. Heuristic technology isn't new, but vendors are increasing their efforts to make it more accurate by minimizing false positives. Vendors have also recoded anti-virus programs to check for and download pattern databases more frequently. Whereas updating databases weekly used to be often enough, today's anti-virus tools need to be checked daily or at least have the updates pushed to them as soon as a new threat is identified. Unfortunately, anti-virus scanners will never be 100 percent accurate, and all it takes is one unpatched system or one computer lacking an up-to-date anti-virus scanner to infect the whole enterprise.

Security Source List

Anti-spam
Barracuda barracudanetworks.com
Brightmail brightmail.com
McAfee mcafee.com
MessageLabs messagelabs.com
Mirapoint mirapoint.com
Postini postini.com
Proofpoint proofpoint.com
SpamAssassin
spamassassin.apache.org
Tumbleweed tumbleweed.com

Anti-virus
Computer Associates ca.com
McAfee mcafee.com
Symantec symantec.com
Trend Micro trendmicro.com

Firewalls
Check Point Software Technologies
checkpoint.com
Internet Security Systems iss.net
Netfilter/IPTables netfilter.org
Sana Security sanasecurity.com
Symantec symantec.com

IDS and IPS
Internet Security Systems iss.net
Lancope lancope.com
Snort snort.org
Sourcefire sourcefire.com
StillSecure stillsecure.com
NFR Security nfr.com

Network quarantine
Check Point checkpoint.com
McAfee mcafeesecurity.com
Microsoft microsoft.com
Trend Micro trendmicro.com
Zone Labs zonelabs.com

As a result, vendors are developing ways to quarantine infected computers and those that don't meet corporate security policy. Several anti-virus vendors offer solutions that will cut off network traffic to and from computers that don't meet predefined criteria. Trend Micro's Network VirusWall appliance will check computers for patch status, enforce the use of up-to-date anti-virus software, and isolate infected machines.

Because malware can infect a computer from dozens of different vector — the Internet, removable media, and p-to-p channels — the best location for anti-virus software is on the desktop, given that. But no matter how malware arrives, it must execute on the desktop to infect the computer. By placing defenses on the desktop, you can detect malware regardless of how it arrives. Defending on e-mail servers is another good strategy, because most worms and viruses arrive via e-mail — although this trend won't last forever. And many entities are placing anti-virus solutions on gateway devices in order to inspect network traffic. Although in theory the gateway would be an opportune position from which to catch malware, scanning each network packet against a large signature database will significantly slow down network throughput. In practice, most gateway solutions scan only a few popular protocols such as SMTP, HTTP, and FTP. That leaves a whole lot of other ports and protocols for malware to exploit.

Network Quarantining

This year's hottest security technology is network quarantining. No matter how strong your network security defenses are, if one misconfigured computer connects to your network, then it's game over — a lesson that was driven home by Slammer and Blaster.

Network quarantining solutions prevent computers that are not properly configured, not patched, or not running updated anti-virus software from connecting to the network. New arrivals are pushed onto a restricted network and inspected. If the computer meets security policy, it is approved and is allowed to connect to the regular network.

Dozens of vendors are developing network-quarantine solutions, including Check Point, McAfee, Microsoft, Trend Micro, and Zone Labs. Some solutions such as Microsoft's Network Access Quarantine Control require specific server and client software. More frequently, quarantining requires special network appliances or software that interfaces with existing routers or switches to handle blocking on the network layer or below. Unfortunately, quarantining is not easy to implement, and the bugs are still being worked out.

Battling Spam

Anti-spam filters are becoming increasingly sophisticated, with accuracy rates in the high 90s being the norm. The best solutions combine Bayesian filtering and content inspection. Most use some combination of Bayesian filtering and content analysis along with whitelists and blacklists.

As a general rule, accuracy improves the farther away you get from the desktop. In test after test, desktop solutions such as those from McAfee, Microsoft, and Symantec fare the poorest. ASP solutions such as MessageLabs and Postini are among the most accurate. *InfoWorld* has also found Brightmail, Mirapoint, and Proofpoint to be very good at blocking spam and avoiding false positives.

Many anti-spam solutions also contain anti-virus mechanisms. Some perform simple file-attachment blocking, and others contain anti-virus scanning functionality. File-attachment blocking is easy to beat, so products using proven anti-virus solutions fare better in removing legitimate threats. As long as you pick an accurate product, anti-spam solutions don't have too many disadvantages beyond the initial expense and setup. The biggest worry is that false positives might block legitimate e-mails, but with training and adjustable scoring, false positives can be minimized.

Security solutions are improving as major vendors combine various technologies into single-offering packages. Unfortunately, software is becoming more complex, and malicious attackers are getting bolder. Until the root causes of computer exploits — poor programming practices and lack of persuasive authentication — get resolved, your enterprise will need multiple defenses.

At a minimum, that means a late-model firewall, network anti-virus, an anti-spam gateway, and e-mail filtering. Midsize and larger companies should also deploy an IDS/IPS and consider a network-quarantine solution.

Even the best solutions are undermined by poor user practices and untrained administrators. Spend as much time on these issues as you do investing in technology.

— Roger A. Grimes

Network Change Management: Responsive, Flexible and Secure

IT COULD BE CALLED THE "IGNOMOMENT;" THE split second following a definitive action when you realize you've just made a tragic mistake. For network administrators, this means the difference between going home at 5 p.m. or 5 a.m. The truth is, despite incidents of careless backhoe drivers pulling up fiber bundles or hurricanes bringing down the power lines, administrator error is the most common reason that a network fails.

"There's always a reason why something doesn't work," says Richard Willmott, market manager at IBM Tivoli. "Finding that reason is the hard part." Too often network administrators are up against a wall, lacking the budget, lab, and time necessary to determine fully the ramifications of modifications made to an internal routing protocol configuration, or accurately determine the impact of a large-scale access list modification on live traffic. Although there is no way to eliminate human error, there are certainly ways to abridge its effects.

A solid change-management process, along with proper training and sufficient IT resources, can turn that sinking feeling brought on by disparate systems and outdated tools into guarded confidence. Then there's ITIL (IT Infrastructure Library), which is a collection of best practices for IT management. It describes in detail the steps necessary to institute various management practices to reduce problems and gain visibility into network infrastructures. Lastly, there are plenty of vendors whose products aim to streamline and automate the change-management process. Nothing is failsafe, but that's no excuse for not trying.

Software developers have the edge when it comes to testing and implementing changes. It's all but unheard of to find developers writing and distributing code without any form of testing. A lab for a developer can be a laptop, and a full-scale software development lab infrastructure can be had for the cost of a few servers.

Yet, for the devices delivering the signals, changes of any scale are typically undertaken without the benefit of prior testing. Why? Because it's nearly impossible to test every aspect of proposed network configuration changes thoroughly. Rather than simply requiring a few servers for a development environment, a network lab requires a wide variety of expensive network hardware to truly mimic the production environment. This means simulating TDM circuits and frame-relay networks, ISDN lines, and any other link types in use on the production network. Simple tests can be accomplished with a subset of the production gear but the costs are high and confidence that the proposed change will function as expected can waver.

For many infrastructures, there are two paths available to deal with this problem. One is a lab environment that can simulate portions of the network; the other is strong change-management policies and change-management software to back up those policies. It's one thing to inadvertently cause network disruptions, it's quite another to realize that you have no backup of the functioning configuration and must replicate detailed parameters from human memory or outdated configurations.

Commercial products are available to help and, fortunately for buyers, this space is hotly contested. Several vendors offer product suites that claim to assist in maintaining policies across disparate network devices and performing automated configuration backup, searching, and restoration. AlterPoint's Device Authority Suite (infoworld.com/1321) offers a complete network development environment patterned on the Eclipse IDE that features extensive automatic scripting tools to develop configuration changes and push them to selected network devices.

Such tools form the core of any change-management initiative. Without proper methods to develop, deploy, maintain, and verify configuration policies across dozens or hundreds of devices, all the procedures in the world will not make a difference. Enterprising carriers and datacenter operators have even integrated help desk software and change window identification to speed the process of linking problems to recent changes and to assist in network troubleshooting. If the framework for thorough change management is available, capitalizing on integrations such as this will definitely help admins sleep at night.

Now This Won't Hurt a Bit

For many network administrators, initiating network change management is like a trip to the dentist — necessary but dreadful. Generally, network configuration changes are slight (an addition to an access list or a change to an SNMP community, for example) and require only seconds to implement. Navigating through onerous change-management guidelines can sometimes seem to complicate seemingly straightforward tasks. Abiding by the guidelines pays off, however. If you don't summon the courage to see the dentist, problems only get worse; it's a similar situation with configuring networks.

Managing Network Changes Step by Step

A simple approach to developing a change management database with common tools is the key to success.

- Focus your initial CM implementation on a single CI (critical infrastructure) such as edge security or a WAN.
- Define a list of stakeholders who will be affected by the initiative.
- Identify the level of detail necessary for implementing CM on your CI. Don't dig too deep. Identify key aspects that can be affected by changes, not minutiae.
- Develop your initial CM database with common tools such as Microsoft Access. Start simply, keep an eye on your larger CM goals, and plan a way to migrate data.
- Institute a project freeze on your CI. Moving targets are harder to work with. Limit CI changes to emergency fixes.
- Begin the documentation process. Bring previously identified data into focus, and work within your CM database until the data makes sense to you and the stakeholders.
- Initiate change. Unfreeze the CI and roll out a change of any magnitude, working with your newly developed CM framework to track changes that arise.

Enterprises can learn much about managing network configurations from service providers, whose very livelihood depends on handling changes smoothly, quickly, and accurately. Nearly all large ISPs have rigorous change-management procedures in place and back those up with thorough configuration management tools. Some ISPs aren't as dutiful and it can show.

Recently, I assisted during a network outage of a multistate MPLS (multiprotocol label switching) network. Although I was not privy to the carrier's network, I was on the call with the NOC (network operating center) administrator looking into the problem. Due to the high-level nature of private MPLS networks, carriers have a much greater impact on the performance and reliability of the service. Where a traditional frame-relay network functions at layer 2, MPLS networks function at layer 3, and the carrier is responsible for maintaining valid routes across all POPs (points of presence). Thus, when a network failure occurs and all network links are active, the problem may lie within the carrier's routing tables. Such was the case here. The problem was eventually traced to a change made in a router thousands of miles from the furthest point of this network, where routes were erroneously injected into the routing tables for my client's MPLS network. The failure was triggered by a seemingly innocuous change to a routing table with no relation to the unintentionally affected network and, until someone contacted the tech who made the change, no one knew it had been made. It took three hours to identify and fix the problem. If an effective change-management policy had been in place, the problem could probably have been averted.

Following the ITIL Framework

Increasingly, effective change-management policies follow the ITIL framework.

In terms of change management, ITIL lays out a foundation based on a CMDB (change management database). The CMDB can take any form, from a simple Microsoft Access database to a fully fleshed-out SQL-driven solution. A CMDB can also be sourced from a vendor, such as Troux Technologies' Troux CMDB for ITIL product. There are also a few hosted CMDB services that are best suited to small businesses, such as myCMDB.com.

The CMDB database contains documentation on all the moving parts of any infrastructure and provides a framework for the modification of this data when changes are instituted. The process of creating and maintaining a CMDB starts quite simply: define and document the critical infrastructures within your network. This can be a slippery slope and anyone with a stake in a particular application will claim that it is highly critical, so some discretion is required. Good examples of critical infrastructures would be security systems, payroll databases, shipping and inventory systems, and interfaces to external partners. These systems need to be documented from the ground up, with all the data residing in the CMDB.

One of the big concerns with the CMDB approach is determining a suitable level of detail, which is why identifying critical infrastructures is so important. It's certainly feasible to document every aspect of the network in excruciating detail, but that may be counterproductive. Limiting the extent of data within the CMDB can prevent drowning under a sea of meaningless data while searching for the necessary elements that affect an existing problem. For instance, the network-management tools described above can participate in the CMDB, but it may not be necessary for device configurations to find their way to the CMDB. External resources may best handle such intricate low-level detail, with the CMDB providing detail on the overall function and purpose of that infrastructure.

Most configuration management vendors support the concept of change management. Although there is a distinction between the two, they are mutually inclusive. One aspect of configuration management that can clearly assist the change-management effort is policy management and adherence verification. Rendition Network's TrueControl, for instance, can generate a report verifying that every Cisco 2950 switch on the network has an identical access list in place, and has TCP small servers disabled. Further, policies can be created and configuration changes pushed to like devices simply, reducing the chance of human error affecting network resources. Of course, this also introduces the potential for the amplification of a single mistake across the network.

Baby Steps First

Of course, best practices guidelines can only help if the infrastructure is stable; there's no sense in erecting scaffolding on a burning building. Instituting a short freeze on any new projects is a good way to achieve stability. This tactic will undoubtedly cause some short-term problems, but the long-term ROI is well worth it. Once the staff is no longer tasked with rush implementations of new systems, stability will increase. Then begin the change-management process by identifying the critical infrastructures, developing the database, and investigating the tools that can automatically update the database. Most IT shops keep track of various network elements in small ways, such as an Excel spreadsheet of VLAN locations or external IP numbering assignments. Migrating this data to a central repository is a simple first step. As with any systemic change, starting small and gaining early victories will pave the way for success when the project becomes more challenging.

After a CMDB is in place, several tools can assist in maintaining it. Help desk software that interfaces with a CMDB can provide valuable information to help desk engineers and simplify the movement of information but may be too complex to implement immediately. Ensuring that the chosen software is accurately populating and retrieving data from the CMDB is far more valuable than a quick implementation. The CMDB is only worthwhile if it is accurate.

Thus, the age-old IT mantra of "plan, plan, plan, and implement" is the best course of action when introducing any form of change management. Small steps get the best results.

Uniting minds over the concept of change management is an important precursor to a change management implementation. Until everyone believes in the benefits, there's little incentive for them to use the system. If the policies are given a positive spin and those in positions of network responsibility see positive results, change management will catch on quickly.

— *Paul Venezia*

Measuring Success Through Metrics

CYNICS MAY SAY THAT THE FIELD OF CHANGE MANAGEment is so vast that the term is practically useless. Yet, if the sun is always shining behind the clouds, one can credit effective change management with improved operational efficiencies and higher customer satisfaction — at least that's the boardroom consensus.

But how do CIOs and IT shops assess whether specific changes implemented in Q2 are paying off in Q3? How do they know if they have applied the right rigor to manage workflow changes or systems integration projects?

"In the manufacturing sector, it's easier — you count the beans, right?" notes Roger Dunn, CEO of SourceIQ, a software and services company. "But for companies rich in BI, it's more difficult."

When dealing with network change management practices, the impact of changes is generally binary and easy to detect. Workflow changes can also be easy to quantify in terms of success or failure because any solid BPM solution measures process efficiencies, according to Hank Barnes, vice president of marketing at Ultimus, producers of Web-based workflow automation software. For example, it can get very expensive for a company that requires the head of manufacturing to sign off on RFQs of $5,000 or more because her time is money — lots of it. If that company were to raise the threshold to $50,000, the Ultimus BPM Suite will capture and display the metrics of approval response times and cost savings, reporting them in a standard module, all online, Barnes says.

"In the end you're still asking, 'Are customers getting what they want and are we getting the money?' " Barnes concludes.

David Rowlands, vice president of Lean Six Sigma at Xerox, emphasizes simplification and automation "from quoting to supply chain to fulfillment to collection of cash." He cites two distinct measurables: business metrics (inventory, cost, level of service, and customer satisfaction) and process metrics (time-related and quality-related events). Both are crucial to assessing change management.

"It all links to how you form the process," Rowlands says. "By measuring in-process, we gauge and predict the outcome to how it will impact either customer satisfaction or financials."

Richard Willmott, market manager at IBM Tivoli, takes a straightforward view when gauging the success of system configuration changes. Using the example of a patch management deployment, he says, "I know when I send a transaction out I have to guarantee it gets from one end to the other, that it's delivered to the end point, and then validate that the install process proceeded." In the end, he says, "You're never going to get change 100 percent right."

At the network layer, though, it's always something to aim for.

— *Richard Gincel*

The Top 20 IT Mistakes

WE ALL LIKE TO THINK WE LEARN FROM MISTAKES, whether our own or others'. So in theory, the more serious bloopers you know about, the less likely you are to be under the bright light of interrogation, explaining how you managed to screw up big-time. That's why we put out an all-points bulletin to IT managers and vendors everywhere: For the good of humanity, tell us about the gotchas that have gotten you, so others can avoid them.

As it turns out, our many contributors to this article had a lot to say — but precious little to say on record. Names may be withheld, but the lessons are still potent. We've distilled this glut of information down to the top 20 mistakes — instances in which wrong decisions can lead to costly project overruns, business disasters, and in the worst cases, lost jobs. Read on, takes notes, and avoid.

1. Botching Your Outsourcing Strategy

Mistakes relating to outsourcing could easily fill our top 20 list on their own. There are two different flavors. The first is the sin of commission: outsourcing important IT functions to avoid the hard work of understanding them. Relinquishing those functions can make it hard to get simple things done.

The other mistake is to hold on to functions that could easily and effectively be outsourced, such as running your own messaging environment. IT organizations with an overt bias against outsourcing could be courting disaster. For example, one CTO we interviewed took over operations for a Manhattan-based online services company, only to discover that the Web-hosting infrastructure for all mission-critical and revenue-producing applications was in-house because the IT staff didn't trust third-party operations. When the great blackout of August 2003 darkened parts of Manhattan for as long as 28 hours, the company's UPS systems kept everything running for only a relatively short time — while competitors at well-provisioned Web-hosting companies experienced no downtime.

2. Dismissing Open Source — or Bowing Before It

For better or worse, many IT shops are susceptible to "religious" behavior — a blind, unyielding devotion to a particular technology or platform. Nowhere is that more true than with open source.

On the one hand, the most conservative IT shops dismiss open source solutions as a matter of policy. That's a big mistake: Taking an indefinite wait-and-see attitude toward open source means passing up proven, stable, and scalable low-cost solutions such as Linux, Apache, MySQL, and PHP. On the other hand, insisting on open source purity in your IT operation can delay progress, as developers are forced to cobble together inferior or unwieldy open source solutions when more appropriate commercial software solutions already exist.

Open source software is not inherently better than commercial software; it all depends on the problem to be solved and the maturity of the solution being considered.

3. Offshoring With Blinders On

Any list of IT mistakes would be incomplete without a mention of offshoring. The experience of one vice president of operations provides an instructive cautionary tale. At his previous employer, the vice president opened a branch office in India for software development and encountered numerous surprises, many counter to conventional offshoring wisdom.

At the time, India had been experiencing an IT employment boom similar to that of Silicon Valley in the late '90s. According to the vice president, the workforce was not stable as a result. Transportation difficulties and the importance of time with family in Indian culture meant that employees generally worked eight-hour days — the concept of the Silicon Valley engineer who goes sleepless at release time was, well, foreign.

In the end, the cost of offshoring the branch office was only 20 percent less than the going rate in the United States, and for cultural reasons, far more face time than initially expected was needed to ensure the commitment U.S. management demanded — which resulted in trips to India at least once per quarter. The vice president emphasized that offshoring can indeed work but said it's a mistake to assume that managing offshore IT is in any way equivalent to managing local IT or that cost savings will be as dramatic as you might expect.

4. Discounting Internal Security Threats

IT managers focusing on external threats can easily lull themselves into a sense of false security. According to Gartner, 70 percent of security incidents that incur actual losses are inside jobs, making the insider threat arguably the most critical one facing the enterprise.

Of course, not all insider threats are born of malicious intent. In September 2004, HFC Bank, one of the United Kingdom's largest banks, sent to 2,600 customers an e-mail that, due to an internal operator error, made recipients' e-mail addresses visible to everyone else on the list. The problem was compounded when customers' out-of-office messages — containing home and mobile phone numbers — responded to the mailing.

Even malicious acts are often carried out using very little technical sophistication. In a joint study released this year by CERT and the Secret Service, 87 percent of insider security breaches were found to have been achieved using simple, legitimate user commands, suggesting that IT needs to be vigilant about granting only necessary privileges to end-users. Identity management with specific permissions can help.

5. Failing to Secure a Fluid Perimeter

IT's responsibility now extends to Starbucks and beyond. The increasing mobility of workers, combined with the proliferation of public wireless hotspots and broadband in the home, means that IT is now responsible for securing systems on networks it does not control. In this environment, solid security means implementing host-based firewalls that will provide some level of protection on an unsecured broadband connection at home or at sites with public Wi-Fi access.

If you're an experienced IT manager, you might feel comfortable with the top-of-the-line firewall you purchased three years ago. You configure it to block all incoming traffic except port 25 for inbound e-mail, and your employees generally make outbound WAN connections to the Web via ports 80 and 443. This is a common approach, but in a more decentralized IT environment, centralized approaches to network security are no longer sufficient. By encrypting traffic on your internal LAN, you will better protect your network from insider threats and from intruders who might have hopped onto your network via rogue wireless access points.

6. Ignoring Security for Handhelds

Although even inexperienced IT managers recognize the need for username/password authentication on network resources and desktop and laptop PCs, most IT shops still seem to be in a "wild West" phase when it comes to handheld devices.

A CTO of a wireless software company tells us about a venture capitalist who lost his BlackBerry on a business trip while he was in the middle of closing a highly sensitive, confidential deal. The BlackBerry wasn't password-protected, so even after the panicked venture capitalist contacted his IT department to have e-mail delivery to the device stopped, anyone who happened to pick up the lost BlackBerry could read e-mails already received.

In this case, the minor convenience of not requiring a password had major implications. Ignoring the security of easily lost devices, particularly those belonging to key executives that traffic in confidential information, is a recipe for disaster.

7. Promoting the Wrong People

As CTO or CIO, rewarding your top technologist with a promotion to a management position might seem like the right thing to do. But when a technologist is not ready to give up constant, hands-on technology work in favor of more people-oriented management duties, it could be a mistake you'll regret on many levels.

One vice president of IT painted a grim picture of such a decision: The promoted employee could be resented by former peers and might not like the new management duties, which could lead to poor performance. Even worse, the new manager might feel compelled to cling to the ill-fitting position because the old position might no longer be available.

Just such an experience put this particular vice president in the tough position of having to deal with a new manager's performance problems, which led to a double whammy: A top technologist left the company, and the new manager still had to be fired.

Management training can help avoid such disasters. But use your gut. Either the aptitude is there, or it isn't.

8. Mishandling Change Management

The former CTO of a computer equipment manufacturer describes one situation in which a talented — but perhaps overly ambitious — systems administrator decided to make seemingly simple changes to a set of critical servers during routine maintenance.

While this individual was making the changes, all of which had been agreed on and planned in advance, he decided on his own to upgrade BIND (Berkeley Internet Name Domain), the open source server software that powers mission-critical local DNS for many companies.

A few hours later, the entire business was at a standstill, as all DNS functions failed. Reversing the "one small change" took hours, and millions of dollars in revenue were likely lost as a result. The lesson is that even talented employees can cause major problems when they don't follow change management procedures.

Remember, change management is cultural. It all starts at the top: If IT management cuts corners, so will IT staff.

9. Mismanaging Software Development

In his seminal book *The Mythical Man-Month*, Frederick Brooks posited that planning software-development projects based on per-unit "man-months" ultimately does not work due to the unique nature of software development.

Even if the building of software could be broken into easily managed, interchangeable time units, the vast productivity difference between the best coders and merely average ones means IT managers might get their best work out of fewer, but more talented, programmers doing their work in less time.

Henri Asseily, CTO of BizRate, tells us via e-mail, "The right individual will always create better and faster core software than a group of people [will]. Everyone in every industry talks the usual talk of, 'We invest in people,' or, 'Our people are our greatest asset,' but nowhere is it more important than in IT. Simply put, a great programmer is 100 times more valuable than a regular programmer."

The mythical man-month has been part of software lore since Brooks' book came out 30 years ago, but many IT managers still plan projects and staff them based on this disproved paradigm. Holding on to this method might lead a naïve IT manager to staff a project with the right number of people for a defined amount of work, but CTOs such as Asseily insist that getting quality people is most important.

"IT managers should devote most of their free time to [finding] the best people. Almost nothing else matters, really," Asseily says.

10. Letting Engineers Do Their Own QA

Not allowing engineers to do their own QA is an axiom of software development, but for small software development teams, there is always the temptation to cut corners. In fact, sometimes management colludes with developers to enable the practice. One CTO relates a situation in which a software development project was running significantly behind schedule and the lead developer had begun to do his own QA to try to speed up the release process. To make matters worse, the lead developer had planned a vacation that was approaching rapidly. A day before the vacation commenced, the developer pronounced all outstanding bugs resolved, and the system was released into production. By the time the developer arrived at his tropical destination, the system was crashing to the point of being unusable. Many of the existing bugs had not been corrected because the developer had not tested thoroughly or formally. Allowing engineers to perform their own QA is akin to allowing defendants to be the judges and juries for their own trials.

11. Developing Web Apps for IE Only

Despite the fact that mission-critical applications continue their march onto the Web browser and that Windows continues to dominate the corporate desktop, Web developers should avoid the temptation to develop applications only for bug-ridden IE. IT shops that insist on using IE for Web applications should be prepared to deal with malicious code attacks such as JS.Scob.

First discovered in June 2004, JS.Scob was distributed via compromised IIS Web servers. The code itself quietly redirects customers of compromised sites to sites controlled by a Russian hacking group. There, unwitting IE users download a Trojan horse program that captures keystrokes and personal data. Although this might not sound like a threat to corporate IT, keep in mind that employees often use the same passwords across corporate and personal assets.

Many enterprises may not be able to avoid using IE. But if you make sure your key Web applications don't depend on IE-only functionality, you'll have an easier time switching to an alternative, such as Mozilla Firefox, if ongoing IE security holes become too burdensome and risky for your IT environment.

12. Relying on a Single Network Performance Indicator

When it comes to network performance, there's no single metric by which to judge network health. Douglas Smith, president of network analysis vendor Network Instruments, points out that it's a mistake to think that network utilization can be quantified in a single way. When management asks for a single network utilization report, IT is typically sent scurrying for a single metric for network health that is ultimately impossible to define.

That said, certain aspects of a network, such as port utilization, link utilization, and client utilization, can and should be measured. In any scenario, successful network analysis means taking a step back and looking at the data in the context of your enterprise.

Network utilization requires judgment calls. If two ports on a switch are 90 percent utilized and the others are not utilized, do you consider your switch utilization to be 90 percent? It might be more appropriate to ask which application is causing those particular ports to reach 90 percent utilization. Understanding the big picture and analyzing utilization levels in context are the keys to getting a sense of your network's health.

13. Throwing Bandwidth at a Network Problem

One of the most common complaints addressed by IT is simple: The network is running slower than normal. The knee-jerk reaction is to add more capacity. This is the right solution in some cases but dead wrong in others. Without the proper analysis, upgrading capacity can be a costly, unwise decision. Network Instruments' Smith likens this approach to saying, "I'm running low on closet space, and therefore I need a new house."

Capacity aside, common root causes of slowdowns include unwanted traffic broadcasting over the network from old systems or apps, such as IPX traffic, or misconfigured or inefficient applications that spew streams of packets onto the network at inconvenient times.

According to Smith, one of Network Instruments' banking customers was considering upgrading its WAN links due to complaints from tellers that systems were running slow. The IT team used a network analyzer to determine that increased traffic levels were being caused by a security app that ran a daily update at 3 p.m. When the IT team reconfigured this application to make updates at 3 a.m. instead, they were able to quickly improve traffic levels without making the costly WAN upgrade.

14. Permitting Weak Passwords

In the Internet age, new threats such as worms and phishing tend to garner all the security attention, but the SANS Institute's Top 20 Vulnerabilities list released in October points to a basic IT mistake: weak authentication or bad passwords (infoworld.com/2193). The most common password vulnerabilities include weak or nonexistent passwords; user accounts with widely known or physically displayed passwords (think Post-it Notes); administrative accounts with weak or widely known passwords; and weak or well-known password-hashing algorithms that are not well secured or are visible to anyone. Avoiding the weak authentication mistake boils down to simple IT blocking and tackling — a clear, detailed, and consistently enforced password policy that proactively deals with the most exploited authentication weaknesses detailed in the SANS report.

15. Never Sweating the Small Stuff

CTOs and CIOs like to talk about the strategic application of technology, but ignoring basic tactical issues can lead to simple but extremely costly mistakes. Missing a $30 domain name registration payment can be enough to grind your business to a halt. In one notorious example, last February a missed payment by *The Washington Post* knocked out employee e-mail for hours until the renewal was paid.

As datacenter environments become denser, even low-level facilities issues may demand scrutiny. On his Weblog, Sun Microsystems President Jonathan Schwartz quoted a CIO who responded to a "what keeps you up at night" question with, "I can no longer supply enough power to, or exhaust heat from [our datacenter]. I feel like I'm running hot plates, not computers." A CIO who overlooks burning — but not necessarily obvious — issues such as these may soon be in search of another job.

16. Clinging to Prior Solutions

A common mistake for IT managers moving into a new position at a new company is to try to force solutions and approaches that worked at a prior job into a new environment with different business and technology considerations.

One current vice president of operations describes a new, low-cost open source environment he had to manage after working in a more traditional shop that relied on high-end Sun hardware and Oracle and Veritas software. The new startup company couldn't afford the upfront cash required to set up a rock-solid environment based on commercial software, so they ran a LAMP (Linux, Apache, MySQL, PHP) architecture with an especially aggressive Linux implementation on 64-bit AMD Opteron machines. Gradually, the vice president realized that his old solutions wouldn't work in the new environment from a technology or cost angle, so he changed his approach to fit the new reality, using none of the technologies from his prior job.

17. Falling Behind on Emerging Technologies

Staying current can prevent a disaster. For instance, the emergence of inexpensive consumer wireless access points during the past few years has meant that anyone can create a wireless network — a real problem for any reasonably structured corporate IT environment. A Network Instruments retail client, for example, was installing a WLAN to serve the needs of employees who measured warehouse inventory levels. Soon enough, management wanted access to the WLAN, and without asking for approval, some employees installed wireless access points at their desks.

Fortunately, the IT staff had implemented ways to check for rogue access points, and a WLAN channel scan with a network analyzer quickly showed there were more access points on the network than the administrator knew had been deployed. In this case, the IT staff recognized an emerging technology that might be stealthily introduced by employees and developed procedures to inventory the threat, thereby controlling it.

18. Underestimating PHP

IT managers who look only as far as J2EE and .Net when developing scalable Web apps are making a mistake by not taking a second look at scripting languages — particularly PHP. This scripting language has been around for a decade now, and millions of Yahoo pages are served by PHP each day.

Discussion of PHP scalability reached a high-water mark in June, when the popular social-networking site Friendster finally beat nagging performance woes by migrating from J2EE to PHP. In a comment to a Weblog post about Friendster's switch to PHP, Rasmus Lerdorf, inventor of PHP, explained the architectural secret of PHP's capability of scaling: "Scalability is gained by using a shared-nothing architecture where you can scale horizontally infinitely."

The stateless "shared-nothing" architecture of PHP means that each request is handled independently of all others, and simple horizontal scaling means adding more boxes. Any bottlenecks are limited to scaling a back-end database. Languages such as PHP might not be the right solution for everyone, but pre-emptively pushing scripting languages aside when there are proven scalability successes is a mistake.

19. Violating the KISS Principle

Doug Pierce, technical architect at Datavantage, says that violating the KISS (keep it simple, stupid) principle is a systemic problem for IT. Pierce says he has seen "hundreds of millions" of dollars wasted on implementing, failing to implement, or supporting solutions that are too complex for the problem at hand. According to Pierce, although complex technologies such as CORBA and EJB are right for some organizations, many of the organizations using such technologies are introducing unnecessary complexity.

This violation of the KISS principle directly contributes to many instances of project failures, high IT costs, unmaintainable systems, and bloated, low-quality, or insecure software. Pierce offers a quote from Antoine de Saint-Exupery as a philosophical guide for rooting out complexity in IT systems: "You know you've achieved perfection in design, not when you have nothing more to add, but when you have nothing more to take away."

20. Being a Slave to Vendor Marketing Strategies

When it comes to network devices, databases, servers, and many other IT products, terms such as "enterprise" and "workgroup" are bandied about to distinguish products, but often those terms mean little when it comes to performance characteristics.

Quite often a product labeled as a "workgroup" product has more than enough capacity for enterprise use. The low cost of commodity hardware — particularly when it comes to Intel-based servers — means that clustering arrays of cheap, workgroup hardware into an enterprise configuration is often more redundant and scalable than buying more expensive enterprise servers, especially when it comes to Web apps.

— Chad Dickerson

Delivering the Goods

Meeting business managers’ expectations means knowing where to apply technology to best effect

ORIGINALLY, I THOUGHT I’D END THE YEAR WITH MY predictions of what next year held for IT. Apparently, predicting the future direction of technology is easy; just read the pronouncements from vendors, bloggers, and the press.

But in the real world, the problem is deciding which technology you should be focusing your IT dollars, a process that becomes a balancing act between business goals and IT capabilities. The push to normalize IT — in other words, to treat it like any other business unit — means IT must also prove its short-term business value. Unfortunately, this effort limits what IT can really do, as my friend Dan Kusnetzky at IDC points out.

With that constraint in mind, I thought that for the last column of this year I would look at what IT’s business partners want and the technology that delivers it. Wallet Share This phrase is being heard in the executive washroom with increasing frequency. It means, “It costs too much to find new customers, so let’s milk the old ones for every cent they’ve got.” Two IT trends are emerging to help companies gain wallet share: customer data hubs and the integration of BI with CRM.

The customer hub concept means transitioning from separate silos of customer information to a single repository. The other trend is to take that single source of data and run it against predictive analytics, for example. Predicting customer behavior is a sure way to sell a lot more products.

Sarbanes-Oxley will push accountability down the org chart. In turn, far more employees will need dashboards and business activity monitoring tools to keep track of key performance indicators for their parts of the world.

Minimizing the number of platforms companies have to support has always been the pot of gold at the end of the rainbow. TCP/IP all the time, every time, will someday be the answer. VoIP is an interesting step in that direction.

Companies are looking carefully at everything they are doing and are asking, “Could someone else do it better than we can?” Hosted and managed applications, as well as offshoring, are becoming more enticing than ever.

There’s a real fight between businesspeople who are eyeing wireless as a way to gain extra productivity and those in IT who are reluctant to support yet another new technology.

The benefits of wireless are real. It is, in the words of my friend Tony Meadow at Bear River Associates, “the last frontier of computerization.”

IDC’s Kusnetzky counters: “Don’t do anything until people are screaming at you.” Unfortunately, when it comes to wireless, next year they may well be doing that. One Version of the Truth In logistics, this means total visibility into your supply chain. In CRM, it’s about being a customer-centric company. In ERP, it’s about automating the reconciliation process.

Data integration is the perennial solution to all these issues, and it cuts down on errors and redundancy while reducing the cost of data entry and validation. Today you can integrate without moving any data using portals and middleware.

So as you can see, I have no fearless forecast for the future of technology. In fact, I don’t see much point in making any forecasts at all. The bigger question you’ll have to answer for yourself is, What’s in your IT wallet?

— *Ephraim Schwartz*