AUTUMN 2011, VOLUME 1, NO. 3

IT in Europe

INFORMATION SECURITY EDITION

Special European edition of Information Security magazine | *www.SearchSecurity.co.UK*

Forecast: Cloudy

While UK companies seem eager to move to the cloud, industry bodies try to raise awareness of the risks

also

PREPARING NETWORKS FOR THE CLOUD

WHITELISTING FOR MALWARE DEFENCE

THWARTING SOCIAL ENGINEERING

TechTarget

Blue Coat

Find The Dangers That Lie Within Social Networking

Blue Coat web security solutions analyze more than 3 billion web requests per week to provide businesses with a real-time defense against malware and other threats, whether you need appliance-based or cloud-based web security.

ForSaferSocialNetworking.com

AUTUMN 2011, VOLUME 1, NO. 3

IT in Europe

INFORMATION SECURITY EDITION

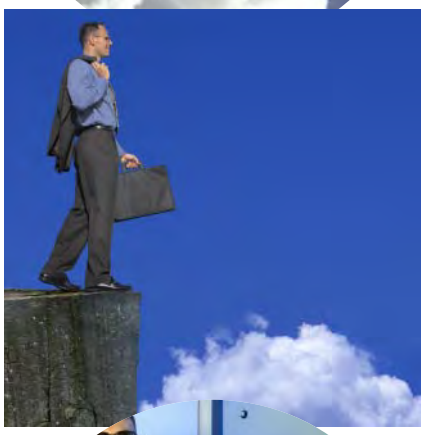

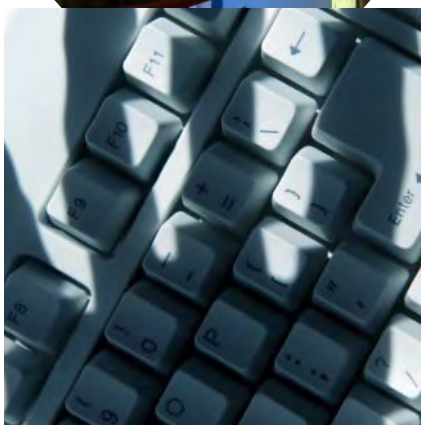

FEATURES

DEPARTMENTS

Your One Stop Shop for All Things Security

Nowhere else will you find such a highly targeted combination of resources specifically dedicated to the success of today's IT-security professional. Free.

IT security pro's turn to the TechTarget Security Media Group for the information they require to keep their corporate data, systems and assets secure. We're the only information resource that provides immediate access to breaking industry news, virus alerts, new hacker threats and attacks, security standard compliance, videos, webcasts, white papers, podcasts, a selection of highly focused security newsletters and more — **all at no cost.**

Feature stories and analysis designed to meet the ever-changing need for information on security technologies and best practices.

INFORMATION SECURITY®

www.SearchSecurity.com

Breaking news, technical tips, security schools and more for enterprise IT professionals.

SearchSecurity.com

www.SearchSecurity.com

Learning materials geared towards ensuring security in high-risk financial environments.

SearchFinancialSecurity.com

www.SearchFinancialSecurity.com

UK-focused case studies and technical advice on the hottest topics in the UK Security industry.

SearchSecurity.co.UK

www.SearchSecurity.co.UK

Information Security strategies for the Midmarket IT professional.

SearchMidmarketSecurity.com

www.SearchMidmarketSecurity.com

Technical guidance AND business advice specialized for VARs, IT resellers and systems integrators.

SearchSecurityChannel.com

www.SearchSecurityChannel.com

TechTarget®
The IT Media ROI Experts

Saying 'Yes' to Cloud Computing

UK Bureau Chief Ron Condon says by promoting the use of cloud computing in IT, infosec pros have a chance to finally say 'yes.'

IT'S NO SURPRISE that the use of cloud computing in IT is receiving a lot of attention. Government and the private sector are both desperate to save money, and cloud providers, with their pay-as-you-go pricing models, promise a quick way to cut capital expenditures.

But what about security? That seems to be the main point of contention stopping many organisations from taking full advantage of the new utility computing model. They just don't know if it's safe.

This presents the security professional with a great opportunity. So often in the past, security people have acquired a reputation for saying 'no' to everything and blocking progress; now is their chance to make things happen and become a business enabler.

Security people should evangelise the benefits of the cloud, while showing that any major transformation of business processes needs to be based on sound risk-based principles. That way, organisations can avoid embarrassing mistakes and also take full advantage of the benefits cloud can bestow.

The recent Public Administration Select Committee report, which severely criticised government IT procurement practices, made this very point and warned against "gold-plating" security requirements. "Over classifying routine administrative and operational information causes unnecessary technology and operational costs, and prevents the public sector [from] taking advantage of the economies and efficiencies of commodity software and new opportunities," it concluded.

Security is, of course, a major consideration in any kind of outsourcing deal, and customers need to understand the dangers, and learn to ask the right questions. How will data be handled, what controls does the outsourcer have in place and what happens when the contract ends?

The security professional is well placed to guide the process, to flag up the dangers, as well as encourage use of the most economic services where risks are lower.

Security professionals also need to stay up to date in the fast-moving industry of cloud-service provisioning. Technologies that will enable easier encryption of data in the cloud are in early development, as are the necessary identity management tools that will allow companies to integrate their cloud usage with their on-premise systems.

In addition, many of the big cloud service providers are now building huge server farms in Europe that satisfy the compliance concerns of companies here, thereby removing another big obstacle to widespread adoption.

The cloud, in all its manifestations, will change the face of IT over the coming years,

bringing down costs and helping organisations to be more flexible and agile. The sooner organisations come to this realisation, the sooner they will reap the benefits and stay competitive against later adopters. By enabling this to happen safely, security professionals have a real chance to help their companies stay profitable in difficult times.›

Ron Condon is UK bureau chief for SearchSecurity.co.UK. Send comments on this column to feedback@infosecuritymag.com.

Cloud Legal Issues

Learn to avoid cloud computing legal issues and SLA conflicts. BY STEWART ROOM

ENTERPRISE CLOUD COMPUTING has been a hot topic for several years now. While it may seem like there's been plenty of time of think through cloud computing legal issues and complications, a number of quandaries in this realm continue to hinder enterprise adoption and usage of cloud computing.

To me, cloud computing legal issues split down into three key domains. First, there is the contractual framework that is required to ensure good service levels. Second, there is the regulatory environment that applies to the processing of data. Third, there is the issue of applicable law, which is about the jurisdictional confusion that flows from the way that (a) cloud services are organised and (b) the legitimate aim of nation states to have access to intelligence for national security, law enforcement and other high-level public concerns.

Ensuring good service levels

Creating a suitable contractual framework for ensuring good service levels should never be a problem, because a contract outlining a cloud service agreement is like any other outsourcing contract, and outsourcing law has developed a massive degree of sophistication over the past 20 years. Thus, one would expect the typical cloud computing contract to include consideration of issues such as service location, sub-contracting, technology refresh, downtime, dispute handling, price, exit and so on.

As with any contractual relationship, the greater the clarity, the lower the risk of satellite disputes around points of detail.

Where problems will arise is if there is inequality in the bargaining power of the parties, which can manifest itself with a take-it-or-leave-it attitude, or where the cloud service provider is unable to answer fundamental questions, such as where data is located, or who will be accessing it. Be wary of these sorts of interactions, because if this is the profile of the pre-contractual relationship, it's easy to see why this may not augur well for the post-contractual relationship.

In most cases, though, negotiation is possible, and so the resulting service-level agreements (SLAs) become critical documents. As with any contractual relationship, the greater the clarity, the lower the risk of satellite disputes around points of detail. One common dispute area concerns the compensation mechanisms for downtime; parties should try to be clear on whether cash rebates are paid, or service credits issued, or some other mechanism put in place.

The regulatory environment

The law has traditionally held the view that commercial organisations are free to agree to the contractual framework of their relationships. However, this traditional approach breaks down when regulation or consumer protection law becomes applicable.

Most organisations deploying a public or hybrid cloud service will face regulatory obligations of some sort, even if these arise as a result of company law, the need for good record-keeping and the need to limit unnecessary operational risk. Additionally, some sectors of the economy are specifically regulated (financial services, pharma, telecoms, etc.), and where organisations process personal data, there is general regulation through data protection law.

It is vital that organisations contemplating cloud computing understand their regulatory obligations. Taking data protection law as an example, there are regulatory expectations about reporting security breaches as well as concrete statutory obligations to manage personal data to limit the scope for unnecessary processing. If the contractual framework for cloud services does not "map" fully to the regulatory environment, the organisation will be put in breach of law, with potentially harsh consequences (in the UK, the data protection regulator can now issue fines of up to £500,000 for breach of the Data Protection Act).

Applicable law

One common vision of cloud involves international data flows, where information moves from country to country, across many geographical borders. The extent to which data flows across such borders is case-specific, but regardless the applicable law issue needs to be fully factored in to pre-contractual planning. For example, US legislation has a long-arm approach for national security purposes, as represented by the Patriot Act, which means some forms of data can be subject to compulsory, secret export to the US. Yet, EU data protection law is against surreptitious transfers of personal data out of Europe, and some countries, like France, even have specific "blocking statutes" to prevent certain forms of international data flow. The key point to remember is international laws can often be contradictory in scope and purpose, which can cause operational difficulties for multi-nationals who are left to resolve these tensions. So, if the law requires data to reside only in particular countries, this needs to be spelled out in the contract to avoid potentially serious problems downstream.

The key point to remember is international laws can often be contradictory in scope and purpose, which can cause operational difficulties for multi-nationals who are left to resolve these tensions.

In conclusion, cloud computing legal issues can pose complex problems, but the comforting message is these problems are manageable. What is needed is a clear under-

standing of the operational goals of the proposed cloud computing arrangement, a specific effort to transfer those goals into contractual mandates, and a consultation with your organisation's counsel to not only eliminate any potential contractual ambiguity, but also ensure the contract's legal language marries up with the business strategy.›

Stewart Room is a partner at Field Fisher Waterhouse LLP. Follow him on Twitter @stewartroom, or visit www.stewartroom.com.

ANALYSIS | CLOUD SECURITY RISKS

Public Sector IT at the Crossroads

Can UK public sector IT overcome its fear of cloud computing risks in order to reap badly needed savings? BY RON CONDON

THE REPUTATION OF UK public sector IT has never been lower. A July report from the Public Administration Select Committee (PASC), the Parliamentary group that monitors the quality and standards of administration within the Civil Service, ominously called "Government and IT—A Recipe for Rip-Offs", painted a grim picture of how IT projects are currently run.

Its main charge was that a small group of major systems integrators had cornered the market on big public-sector IT projects, and the civil service has lost most of its IT skills, becoming a passive customer that blindly accepts what it is given by its suppliers.

The committee uncovered some astonishing examples of over-charging as well, such as the government department that paid £3,500 each for its PCs.

But the report's primary theme is the declaration that the government's whole process for procuring IT projects was fatally flawed. It listed several big-ticket IT projects that have run late, under-performed or failed over the last 20 years:

- The Child Support Agency's IT system
- The IT system that would have underpinned the National ID Card scheme
- The Defence Information Infrastructure Programme
- The Single Payments Scheme by the Rural Payments Agency
- The National Offender Management System

The authors added: "During the course of our inquiry there was evidence of continuing IT mismanagement: the Department for Work and Pensions (DWP) chose to cancel a contract with Fujitsu for desktop computers; one of the NHS partners involved in the electronic patient record system pulled out after the suppliers failed to meet a deadline; and the flagship Universal Credit (UC) programme was reported to

be behind schedule due to problems meeting the deadline for building the new IT system." Also worthy of mention is the multi-billion pound NHS National Programme for IT, which now faces cancellation nine years after it started.

So how can this cycle of failure be broken? The government has already declared a willingness to adopt cloud computing in some form or another. If implemented properly, there are a variety of cloud services that would provide lower costs and greater flexibility than the current huge siloed data centres run by government.

A research study by the Centre for Economics and Business Research (Cebr), published in February 2011 called "The Cloud Dividend", assessed how much money could be saved across Europe in different market sectors. The study, commissioned by storage vendor EMC, found that the macroeconomic benefits of cloud computing for the government, education and health sector for Europe's biggest economies (UK, France, Germany, Italy and Spain) will amount to 112 billion euros between 2010 and 2015. For the UK, the accumulated benefits in the same period will be 19.445 billion euros (£16.601 billion).

The study, commissioned by storage vendor EMC, found that the macroeconomic benefits of cloud computing for the government, education and health sector for Europe's biggest economies (UK, France, Germany, Italy and Spain) will amount to 112 billion euros between 2010 and 2015.

However, the figures assume the government's plan for G-Cloud, a wide-ranging project to create a secure cloud infrastructure, hosting services and applications to be used by public sector bodies around the UK, would come to fruition as planned.

Yet in recent months, doubts have been cast in regard to G-Cloud. Some sources say the plan has been canned altogether, with the focus now shifting to data centre consolidation instead. A recent publication from the Cabinet Office confirms that change.

Furthermore, a recent Freedom of Information request sent to 25 government ministries revealed little interest in cloud computing. The research, commissioned by Indian IT services company HCL Technologies, found six respondents, including the Treasury, said they had no plans to adopt cloud computing. Their main objections included being tied to one service provider, system integration and Web security. Although several departments said they would follow the government's IT strategy and guidelines, it was only the Department for Work and Pensions that mentioned G-Cloud in its responses.

Security is often cited as a reason for not adopting cloud services in government. But as the PASC report noted, departments must guard against "gold-plating" their security requirements. "Over-classifying routine administrative and operational infor-

mation causes unnecessary technology and operational costs, and prevents the public sector from taking advantage of the economies and efficiencies of commodity software and new opportunities," the report said. "It also acts as a further barrier to more effective use of SMEs in the supply of IT goods and services. Government must do more to demonstrate how a risk-based approach is helping achieve a better balance in information assurance."

However, government departments are highly regulated and need to be audited against a whole range of standards by different authorities, said Stephen Simpson, UK cloud lead at IT services company Logica. "Today, these authorities look at the organisations independently of each other and this must change in order to achieve large-scale consolidation," he said. "It only makes sense for ICT cloud consolidation to take place between similarly regulated organisations; so the focus of the authorities can then move towards the general protective monitoring of common services rather than those of the constituent organisations."

Simpson said the public sector will remain reluctant to embrace public cloud while the majority of the available infrastructure is not located inside the UK. Logica estimates 70-80% of candidate cloud services need to be run at impact level 2 (IL2) or impact level 3 (IL3) (as defined by CESG, the government security body). "IL2 can in principle be managed outside of the UK, but the two levels tend to be mixed-up together in planning," Simpson said. "Redaction—the filtering of sensitive information—is also a major issue, particularly for unstructured data; organisations are reluctant to put information into a shared service centre when some of the content might be sensitive."

"Redaction—the filtering of sensitive information—is also a major issue, particularly for unstructured data; organisations are reluctant to put information into a shared service centre when some of the content might be sensitive."

–STEPHEN SIMPSON, UK Cloud Lead, Logica

However, with the general economy struggling to recover, the government knows it must save money by improving its use of IT. There is much to play for, and many opposing forces are jockeying for position. The big systems integrators, under criticism for over-charging and under-performing, have the most to lose from any changes, while smaller service suppliers see an opportunity to step in and grab the business.

Andy Burton, chairman of the Cloud Industry Forum, a body representing cloud service providers, not surprisingly falls into the second camp and wants the government market opened up to smaller players. "The UK government controls over 100 data centres across the UK, and this infrastructure is inefficient, capital intensive and provides a poor return both economically and in the restriction of agile solutions," he said. "The Cloud Industry Forum believes the government can save at least £2-4 billion per annum

within three years if it mandates a new process for specifying and procuring IT solutions. This saving[s] would not restrict the government IT agenda and arguably would stimulate creativity and collaboration as is seen on the private sector," Burton said.

While data security, privacy and sovereignty remain key cloud computing security risks among some parts of government, Burton insists: "As with any technology, solutions exist to these issues and practical precautions need to be taken, as with any solution, to ensure all data is safe, and this applies to using a cloud-based service."

He also makes a point shared by all parties: that the people buying IT services on behalf of the government need to reclaim many of the skills they have lost over the years of outsourcing. "Government IT executives must reposition themselves as leaders who can bring their organisations to new levels of performance and efficiency through IT. Cloud computing is definitely a plausible option for improving performance and efficiency and the government should be considering this."

There is, however, some good news. Several providers report that local authorities, NHS Trusts, police authorities and others are already taking matters into their own hands to lower costs and respond to demand for better efficiency. By consolidating their own data centres using virtualisation, or pooling resources with neighbouring bodies, they are developing private clouds and avoiding the security concerns they might have while working with an outside supplier.

Ron Condon is UK bureau chief for SearchSecurity.co.UK. Send comments on this column to feedback@infosecuritymag.com.

Forecast: Cloudy

WHILE UK COMPANIES SEEM EAGER TO MOVE TO THE CLOUD, INDUSTRY BODIES ARE TRYING TO RAISE AWARENESS OF THE RISKS. BY RON CONDON

CLOUD COMPUTING APPEARS to be taking off fast in the UK, and while data suggests the vast majority of enterprise customers are becoming increasingly comfortable with cloud-based services, many are uncomfortable with the security ramifications of data in the cloud.

Not only are businesses saving capital expenditure on in-house IT equipment thanks to cloud computing, but they are also benefitting from the extra agility and flexibility that cloud services can deliver. It means they can pay for services when they need them, and save money when they don't.

The evidence of this comes from an attitude survey conducted in January and February by market research firm Vanson Bourne on behalf of the Cloud Industry Forum, a trade body for the emerging industry. The researchers asked 450 senior managers who have responsibility for making IT purchasing decisions what they thought about cloud services.

The survey sample included a broad spectrum of organisations from both the public and private sectors, and ranged from small companies to large enterprises. It found that 48% of organisations already use some cloud-based services, with the greatest use coming from private companies with more than 200 employees, where the figure was 53%. Of those not currently using cloud, 31% said they were planning to do so within the coming year.

The initial driver for most cloud customers had been the flexibility offered, often to cover a short-term lack of resources or meet a tight deadline. But, once companies got a taste for the cloud, the survey results suggest they liked it and started to crave the longer-term economic benefits. The survey found 94% of cloud users were satisfied with the experience, and 85% said they expected to increase their use of cloud services in the coming year. Growth focused on three core applications: email, disaster recovery and data storage.

But, when it comes to extensive implementation of cloud computing, UK adoption is still patchy, limited mainly to a few specific applications. And it seems many companies are still hesitant about entrusting any valuable information or mission-critical systems to a cloud service provider (CSP).

When asked about their main concerns, 64% cited data security, followed by data privacy (62%), dependence on Internet access (50%), doubts over supplier reliability (38%) and contract lock-in (35%).

When asked about their main concerns, 64% cited data security, followed by data privacy (62%), dependence on Internet access (50%), doubts over supplier reliability (38%) and contract lock-in (35%).

The physical location of data held in the cloud also proved to be an important factor. Some insisted it should be kept in the UK, while others were happy as long as it remained in Europe. This seemed to be of most concern to smaller companies with fewer than 20 employees, and also the public sector. Larger enterprises, which often have greater resources to help them manage the risk, seemed to be more relaxed about the issue.

Despite these stated concerns, the survey responses revealed a general willingness among customers to sign up for services without question and to accept whatever the provider had to offer. Barely half (52%) of the companies using cloud services said they had negotiated the legal terms of their contract, rather than simply signing the contracts the providers handed them.

The report comments: "Some of the most striking results from the research show users are often in the dark over questions of liability, indemnity, insurance and ownership of content stored in the cloud; and that while users have certain expectations in these areas, they often do not know if they are being met in their contracts with CSPs."

It also reveals only 45% said their provider offered them a chance to agree to changes to their contract (38% answered 'Don't know'), and 46% of customers allowed their

supply contract to be renewed automatically.

Thus, while the market is clearly expanding fast, and some customers are still holding back because of doubts over security, others are committing themselves to standard contracts with little thought given to how any future problems might be resolved.

This mixed picture is to be expected in an immature marketplace, but, as SearchSecurity.co.UK reported recently, even large companies can overlook basic security provisions when committing to a cloud service provider. In that instance, a major company had signed up for a service with no real up-time guarantee, no controls over which of the provider's users could access the service, and little idea of where its data was being held. Furthermore, the provider had insisted on transferring the company's own Active Directory servers to its own data centre, thereby exposing its non-cloud applications to the risk of snooping.

Can today's enterprise customers, in order to avoid that kind of situation and prepare for problems and disputes that will inevitably arise in any contract relationship, learn what kinds of cloud computing security questions to ask in the first place, and negotiate contract terms that can be properly enforced?

The Jericho Forum, a think-tank devoted to information security matters, has produced a simple model to help companies decide which systems are best suited to the different flavours of the cloud: public, private, community or hybrid.

Fortunately, a lot of helpful information is being produced at a rapid rate by a variety of security-focused organisations, in order to raise the level of professionalism and confidence in this nascent industry.

The Jericho Forum, a think-tank devoted to information security matters, has produced a simple model (.pdf) to help companies decide which systems are best suited to the different flavours of the cloud: public, private, community or hybrid.

Jericho has also joined forces with the Cloud Security Alliance (CSA), which has produced the freely downloadable Consensus Assessments Initiative Questionnaire (CAIQ), a list of key questions any customer should consider when adopting a cloud service.

In addition, the professional body ISACA has just published a book on the subject for its members, called *IT Control Objectives for Cloud Computing: Controls and Assurance in the Cloud*, which provides readers with extensive checklists and advice about managing a cloud service contract. (The first two chapters are freely downloadable.)

For its part, the Cloud Industry Forum (CIF) has created a code of practice for its members and has produced a series of white papers offering help and support for customers, including how to write a good contract. According to CIF Chairman Andy Burton, the aim is to get providers to be open about how they work, and to provide clear and relevant information about their services, such as whether they aggregate services

from other companies, whether they can guarantee round-the-clock operation, and also where they store data.

As part of the code of practice, providers are encouraged to complete the CSA's CAIQ document, thus providing the prospective customer with all the answers they are likely to need. If the supplier does not complete the questionnaire, customers can draw their own conclusions. "It's essential to provide transparency so customers can compare vendors and make a rational decision," Burton said. "That is why we are seeing a lot of private cloud adoption rather than public, because customers are still taking a cautious approach."

"It's essential to provide transparency so customers can compare vendors and make a rational decision. That is why we are seeing a lot of private cloud adoption rather than public, because customers are still taking a cautious approach."

–ANDY BURTON, chairman, Cloud Industry Forum

In many ways, the current cloud market is reminiscent of the early days of IT outsourcing, when many organisations rushed in thinking they could find a quick fix to a problem. Many of those early contracts later turned sour, and customers found it hard to back out of arrangements and switch to alternative providers. If the CIF survey is to be believed, some customers may have already laid themselves open to similar experiences later on with their CSPs.

Nevertheless, if the cloud trend is as unstoppable as the march of the PC in the mid-80s, well-defined contracts and service-level agreements may help to mitigate some of the risks.

But to exploit the full economies of scale the cloud has to offer, companies still need to do more, says Paul Simmonds, a founding member of the Jericho Forum. "A lot of companies are going into the private cloud because they cannot guarantee the security of the public cloud," he said. Simmonds sees two barriers blocking more extensive cloud use: the lack of viable encryption for data residing in the cloud, and the problems of identity management.

The other barrier, identity management, is one of convenience, he explained. "At Astra Zeneca [where Simmonds was CISO until recently] I had 66,000 users. The last thing I wanted to do was give users another username and password. It's hard enough in any corporation to keep users synchronised, and to keep up with joiners, leavers and movers, even when you own both ends of the problem," he said.

"The cloud service should not be holding usernames and passwords" Simmonds said. "They should be leveraging the appropriate SAML assertions from your existing identity system to apply a set of rules that say: If this is an Astra Zeneca password with an Astra Zeneca certificate backing it, and it is an assertion I can validate, then let them into the account."

He said a handful of providers understand the requirement, but "99% of cloud service providers, such as Salesforce.com, require a unique username and password for their service. They don't want something different for every customer, which is why we need some new standards for how you do this stuff."

In the meantime, it is down to the vendors and standards bodies to educate the market and help companies understand what they are buying. For example, many customers still have unrealistic expectations, according to Paul Lightfoot, managed services director for The Bunker, a hosting company running two data centres in former nuclear bunkers in Kent and Berkshire. "They come to us and say, 'We see you're PCI DSS compliant, so if we store all our credit cards with you and there's a problem, it's your responsibility'," he said. "We have to explain that it's a shared responsibility, and depends as much on their applications and practices as on what we do."

Even so, he said today's customers are better informed and many insist on talking to other customers before committing to a service. The best solutions come when customers are well informed and carefully consider the relative risks associated with different types of systems and deploy them according to their level of sensitivity in an organised way, Simmonds said.

Lightfoot recommends putting non-sensitive data on low-cost shared systems, more crucial data on a dedicated blade in a shared system, and mission-critical systems on a dedicated server. Customers of the Bunker, like those of many large cloud providers, can also choose to pay more for a completely redundant service, to ensure a failure at one data centre will not affect them.

Regardless, say the experts, the customer must come to the process from a position of knowledge rather than blindly hoping to save money. That means doing some research and using the free advice offered by groups, such as the CSA and CIF, and learning to ask the right questions before signing any contract.

Ron Condon is UK bureau chief for SearchSecurity.co.UK. Send comments on this column to feedback@infosecuritymag.com.

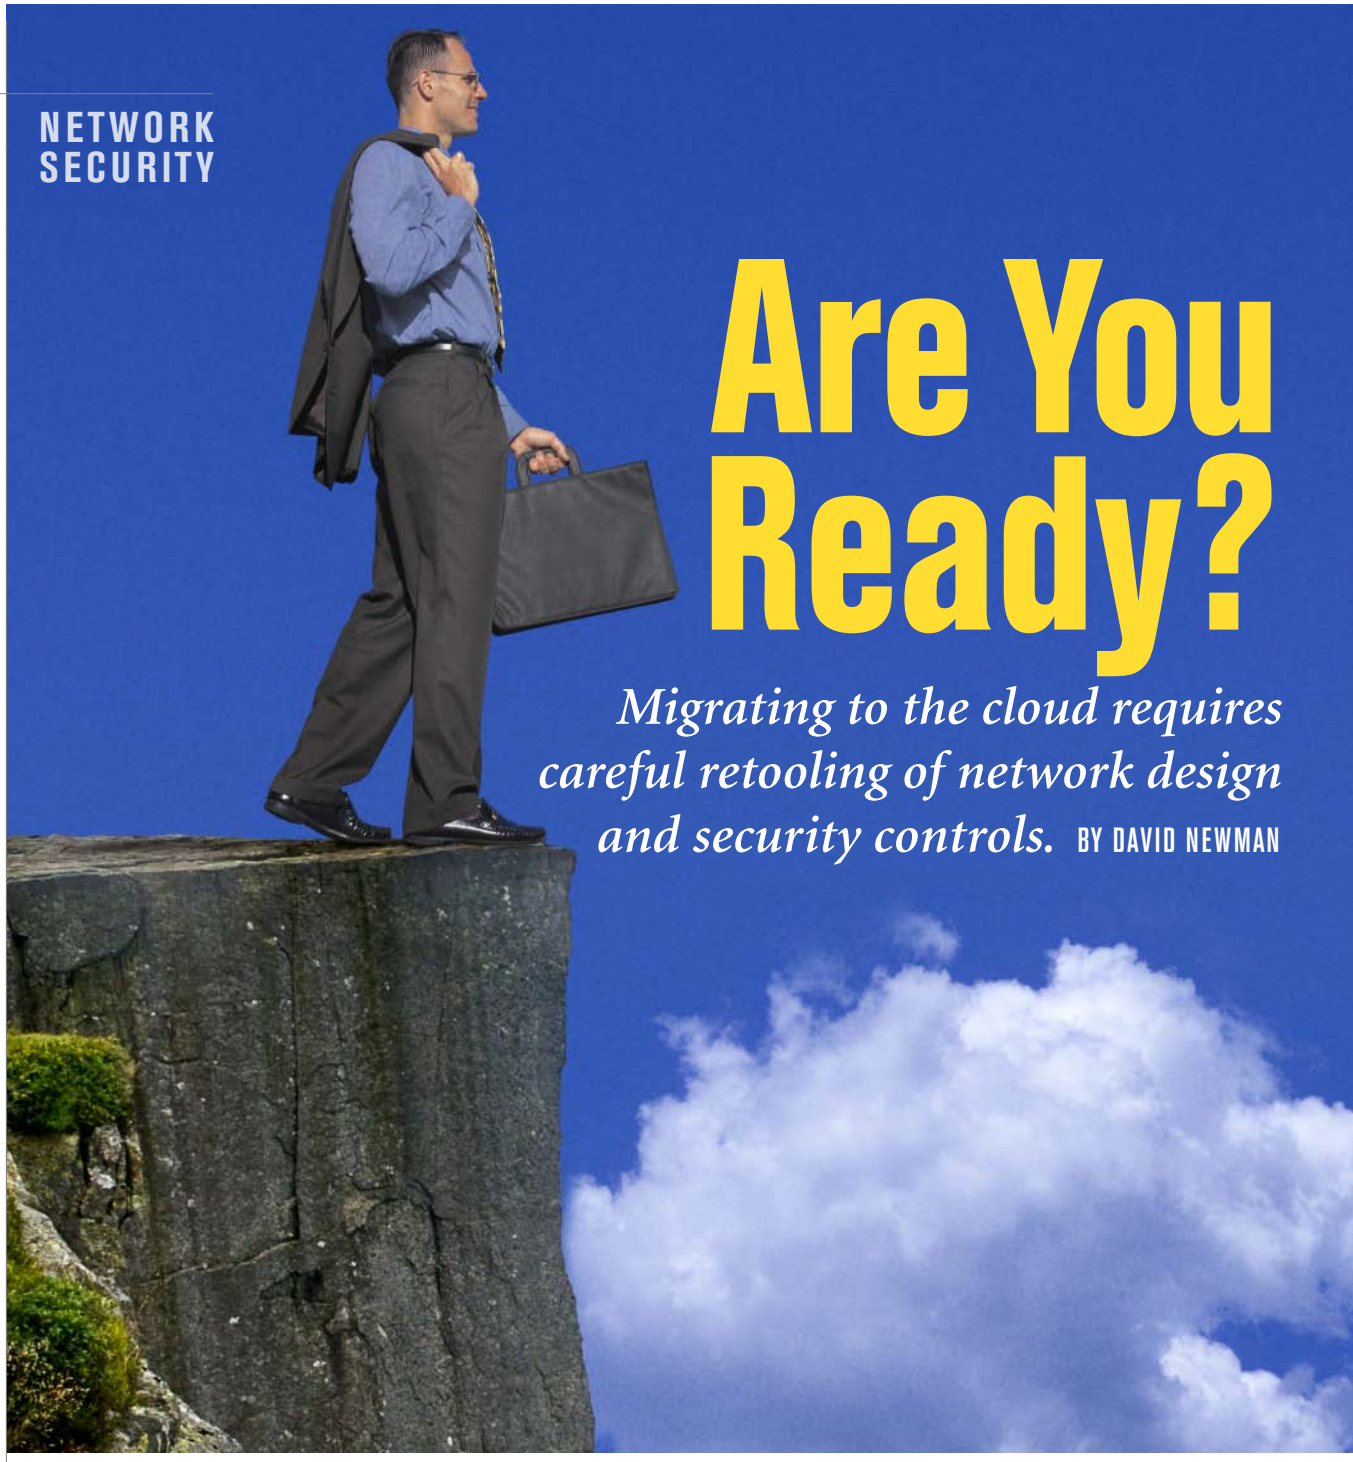

Are You Ready?

Migrating to the cloud requires careful retooling of network design and security controls. BY DAVID NEWMAN

CLOUD COMPUTING PROMISES many benefits: It can reduce IT costs and downtime while vastly increasing storage, mobility, and provisioning options. It's also a potential security nightmare: perimeters disappear; clients and servers move around at will; and old models of access control, authentication, and auditing no longer apply.

All these challenges can be met, but any migration to the cloud requires careful planning. Cloud computing fundamentally changes long-standing best practices in network design, encryption and data loss prevention, access control, authentication, and auditing and regulatory compliance. To prepare their network for the cloud, organisations need to take stock of their infrastructure and adjust their practices and processes accordingly.

START WITH THE PLUMBING

A common misperception about cloud computing is that moving services to an off-site provider will reduce bandwidth requirements. In fact, the reverse is often true: Cloud computing can increase bandwidth requirements due to increased Internet connectivity. A move to cloud computing also has implications for virtualisation and the suitability of existing security infrastructure and security policies.

To understand how cloud computing can radically shift network and security requirements, consider a common hub-and-spoke network design *(see Figure 1)*. Here, branch offices connect with one or more enterprise data centres where key applications reside. There's a well-defined perimeter to the public Internet, and the bandwidth, latency, and packet loss characteristics between sites are easy to measure.

Figure 1

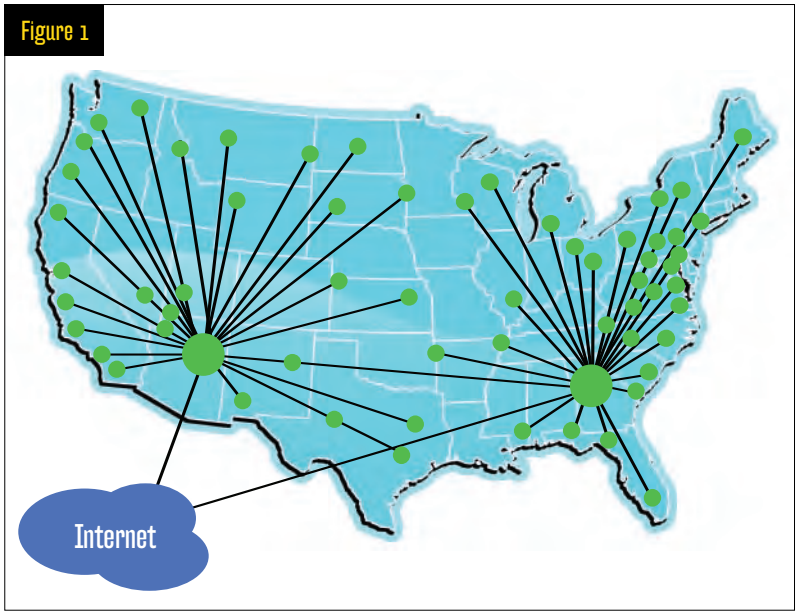

In contrast, cloud computing involves Internet connectivity for every site in the enterprise *(see Figure 2)*. Here, given that applications now reside in the cloud, there is no clearly defined perimeter. Further, the traffic characteristics of every site's Internet connection may affect application performance. As a result, some organisations find a move to the cloud results in increased requirements for bandwidth and security monitoring.

Figure 2

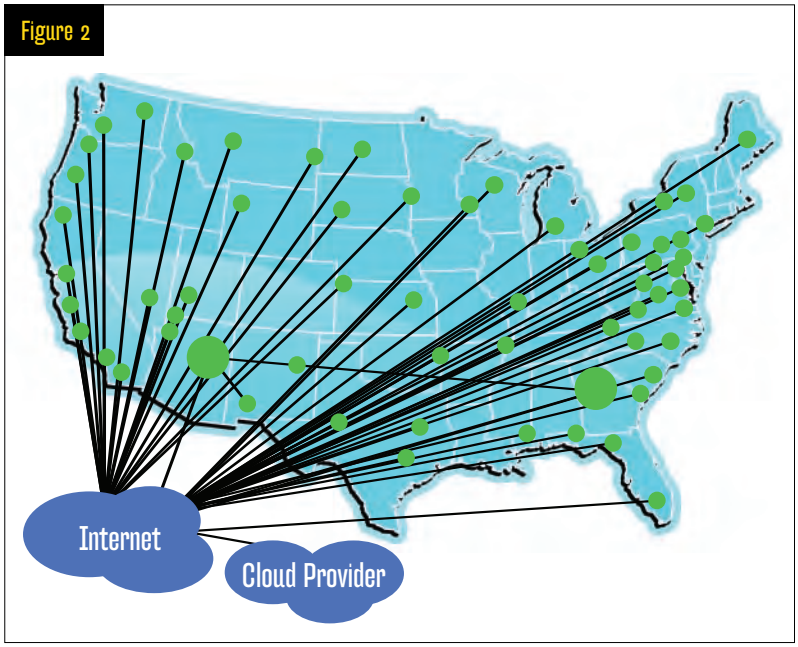

Beyond basic network characteristics, there's also the question of what kind of traffic leaves the enterprise as it moves to cloud computing. Understanding what kind of traffic you have is just as important

as knowing how much traffic you have. If network flow analysis—which uses existing flow-reporting tools in routers and some switches to provide an in-depth view of application traffic—isn't already deployed, this would be an excellent time to consider implementing it.

To be fair, this is an extreme, strawman example of cloud network design. Hybrid designs are more likely, with branch-office Internet connectivity still channeled through one or more internal data centres. Even so, cloud computing means key applications are reached via new connections outside the enterprise. Testing the network characteristics of these new connections is critical.

NETWORKING AND VIRTUALISATION

Virtualisation is a key enabling technology for cloud computing and data centre consolidation. Well before moving to the cloud, many enterprises have adopted virtual servers as a means of saving on hardware, increasing uptime, or both. For these organisations, migrating a virtual infrastructure to the cloud can have a significant impact on application performance.

Consider vMotion from VMware, which moves virtual machines (VMs) between host servers with virtually no downtime perceived by users or applications. This is truly the "killer app" for virtualisation; network managers like vMotion because it's such an easy, hitless way to move VMs around.

It's possible to use vMotion to move VMs across slower wide-area network links, but not with its zero-downtime benefit.

For all its benefits, though, implementing vMotion into the cloud can affect application performance. First, there's the issue of bandwidth: vMotion requires lots of it, and assumes a high-throughput, low-latency network. It's possible to use vMotion to move VMs across slower wide-area network links, but not with its zero-downtime benefit. This could be an issue when using vMotion between an enterprise staging site and the cloud provider, or even within the cloud provider's network if that encompasses multiple physical sites. Either way, if network managers want to avoid VM downtime, ensuring close proximity of VMware hosts is a must.

Second, vMotion generally requires source and destination VMware host servers to reside within the same layer-2 network (that is, within the same broadcast domain). This isn't a problem even in large data centres, which deliberately create very large broadcast domains to accommodate virtualisation. However, it could be an issue in moving VMs across different IP subnets, for example between an enterprise and the cloud provider. Suitability for vMotion should be a part of any network design review. The same caveats apply for vApps, which does for applications what vMotion does for VMs.

IMPACT ON SECURITY DEVICES

If Internet traffic increases with cloud computing, then so too will the load on security devices such as firewalls, VPN concentrators, and IDS/IPS appliances. This has implications both for pure performance and for security policy. The performance piece is simple: Increased Internet connectivity means a heavier workload for security devices. It's great to upgrade to, say, a 100-Mbit/s Internet connection as part of the move to cloud computing, but if existing security devices are rated only to 10 Mbit/s, they will quickly become a bottleneck.

Depending on security policy, a move to the cloud may require enabling additional IDS/IPS signatures, and this too can have a negative performance impact.

Depending on security policy, a move to the cloud may require enabling additional IDS/IPS signatures, and this too can have a negative performance impact. Network Test has conducted multiple performance assessments of multifunction security devices where forwarding rates drop by a factor of 20 times or more when IDS/IPS signatures are enabled. VPN devices such as IPSec or SSL concentrators also can degrade throughput and increase latency.

Other policy issues to consider include interoperability and changes to existing firewall rule sets. Cloud providers have their own security devices, but long experience with IPSec and SSL VPN troubleshooting suggests interoperability isn't a given. Even though both IPSec and SSL are based on open standards and may work flawlessly inside a multivendor enterprise network, there's no guarantee of interoperability with a cloud provider's equipment. Similarly, firewall and IDS/IPS rule sets will change as enterprises move more applications into the cloud, possibly affecting other parts of the firewall rule set in unexpected ways.

Testing can help validate a move to the cloud, provided it's done with a meaningful workload. When it comes to performance measurement, some security appliance vendors perform tests using overly simple workloads. It's possible, for example, to test a firewall the same way as an Ethernet switch, and then only with large packets. However, this isn't a very stressful load; it will produce impressive numbers for a data sheet, but it's not representative of enterprise traffic.

A better practice is to model the particular mix of applications that will reside in the cloud, paying particular attention to transaction sizes, transaction durations, concurrent connection counts, overall bandwidth utilisation, and network characteristics such as latency, jitter, and packet loss. With these key metrics in hand, it's possible to craft a synthetic workload that will yield meaningful predictions about security device performance for a given enterprise.

ENCRYPTION AND DLP IMPLICATIONS

Figure 3

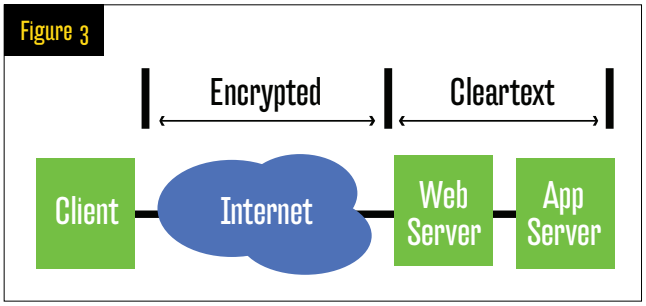

As noted, cloud computing changes or eliminates the concept of a perimeter, and that has profound implications for encryption and data loss prevention (DLP).

Prior to cloud computing, network managers were mainly concerned with a single set of encryption endpoints between customers and Internet-facing servers *(see Figure 3)*. That changes with cloud computing, where there are now three sets of encryption endpoints to consider: from customer to Internet; within the cloud; and from cloud to enterprise *(see Figure 4)*. Encryption within the cloud may be necessary for regulatory compliance, or because a cloud provider's network may span multiple physical locations.

There's no one right approach as to where to apply encryption in the cloud model. The simplest approach of encrypting everything from end to end sounds appealing, but also has the unintended consequence of "blinding" some key security and network management tools such as application-aware firewalls and deep-packet inspection devices. Encryption everywhere also can complicate DLP, where the imperative is to maintain visibility of where data is sent and stored.

A revised policy should banish, once and for all, insecure protocols such as FTP that allow cleartext transmission of passwords and other sensitive data.

As usual, policy is the right place to begin in redesigning encryption and DLP for the cloud. At a minimum, a cloud-aware security policy should specify that traffic never leaves the enterprise unencrypted. Security policies should be revised to add requirements for detection of any breach of the encryption policy, including within the cloud provider's network.

Similarly, a cloud migration is an ideal time to review policy as to permitted protocols. A revised policy should banish, once and for all, insecure protocols such as FTP that allow cleartext transmission of passwords and other sensitive data. At the same time, policy also should specify which users can employ protocols that might leak data

Figure 4

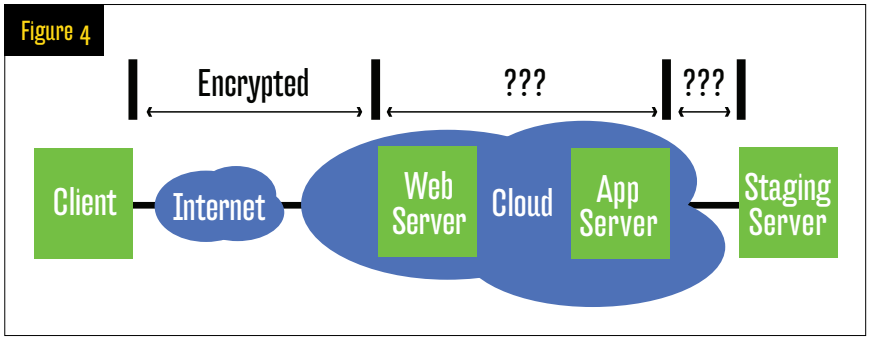

over encrypted protocols such as SSH and Secure Copy (SCP).

A redesigned DLP infrastructure can actually help solve some encryption problems by automating many processes. For example, DLP systems can automatically encrypt files attached to email and monitor traffic for files sent outside the enterprise using email or instant messaging. File-level encryption is also an option.

One final question to consider is whether the existing encryption and DLP infrastructure is adequate for cloud computing. Even if no upgrade to encryption and DLP is deemed necessary, network managers should consider how to implement these services within the cloud: As VM versions of existing appliances, as hardware devices between VMs and the network, or some combination of these.

A DIFFERENT ACCESS CONTROL MODEL

Cloud computing also changes long-standing concepts about access control. Historically, enterprises have used IP-centric access control models, where rules were based on criteria such as source and destination subnet addresses. That doesn't make much sense in a cloud context, where users can connect from anywhere, on any device, and where servers may be cloned or move around within the cloud.

Access control in the cloud should follow the NAC model of applying rules dynamically, in real time, as endpoints appear on the network.

Cloud computing changes access control from an IP-based to a user-based model. Essentially, cloud computing adopts the network access control (NAC) credo that who you are governs what resources you can reach. Because both clients and servers can be mobile in cloud computing, a dynamic approach to security policy is needed. Access control in the cloud should follow the NAC model of applying rules dynamically, in real time, as endpoints appear on the network. This approach is equally valid for clients and servers.

Of course, user-based access control supplements, but does not replace, the old IP-centric rules. Any sound migration strategy should include a review of existing access control lists (ACLs) on enterprise routers. It may make sense to rewrite and tighten ACLs so that inbound traffic for key applications comes from, and only from, the cloud provider. Similarly, new rules may be necessary to enable users to reach newly migrated applications in the cloud.

AUTHENTICATION REQUIREMENTS

Cloud computing stretches authentication requirements, both figuratively and literally. Anywhere, anytime client connectivity may require new, stronger forms of authentication. At the same time, the move to place services in the cloud extends the trust

domain enterprises need to protect. For both clients and services, strong control over password and key management is a must, as is better break-in detection.

With cloud computing, clients no longer cross a single, well-defined security perimeter before being granted access to enterprise resources. Clients also may connect to these resources from shared public networks such as Wi-Fi hotspots, increasing the risk of password interception. A move to two-factor authentication, for example, tokens plus some biometric mechanism, makes sense to ensure clients are properly authenticated. Some well-known public cloud services such as Google Apps also support passwords plus tokens for authentication.

Password synchronisation is also important. Maintaining separate sets of user accounts and passwords, one apiece for resources in the cloud and in the enterprise, is not a sound practice. Besides the added administrative overhead, two sets of accounts also inconveniences users and doubles the likelihood they will write down one or both passwords and save them in public view. A single sign-on system covering both enterprise and cloud-based user accounts can help here.

There's also an imperative to protect authentication mechanisms in the cloud, including both passwords and API keys. Many cloud services make use of representation state transfer (REST) Web services, which in turn use API secret keys for authentication.This raises a couple of potential risks. First, REST security can be poorly implemented. For example, a security researcher has demonstrated how a major hosting provider transmits the secret key in plaintext as part of an authentication request. Although the request must be made over SSL, any compromise of either side of the SSL tunnel would also result in loss of the secret key.

At least initially, cloud computing complicates the security auditor's job, since the systems and processes to be audited will be much more widely distributed.

Second, even in a well-designed system the API key represents an extremely valuable resource, with serious consequences if it's lost. For example, enterprises on Google Mail identify themselves to Google's servers using an API key associated with the entire enterprise, not individual users. If this secret key were stolen, an attacker could impersonate any email account or share any Google document associated with the enterprise. Sound practices to protect the API key include encryption and a software audit to review API usage.

A review of IDS/IPS and DLP configurations also is in order. If signatures to detect cleartext transmission of passwords aren't already in place—for example, in IMAP and POP email—they should be added.

COMPLIANCE COMPLICATIONS

At least initially, cloud computing complicates the security auditor's job, since the systems and processes to be audited will be much more widely distributed. And there are certain

to be regulatory considerations when it comes to moving sensitive data to and from the cloud.

Logging and monitoring is critical in the cloud, but also more complicated, with large cloud providers' networks spanning multiple continents. While this has the advantage of moving content closer to users, it complicates timestamp synchronisation between server logs. Without rigorous time synchronisation among servers, troubleshooting becomes very difficult. Setting all system clocks in a single time zone, such as coordinated universal time (UTC), also is essential for taking the guesswork out of distributed log analysis.

A move to the cloud may increase the number of servers involved, especially where virtualisation's cloning features are used, and this in turn increases the volume of logs to be analysed. Network managers may want to consider implementing a unified log analysis system to collect and synthesise data from all the new sources.

Various regulatory regimes require data sanitising as data moves to and from the cloud. This is similar to the encryption issues previously discussed, where cleartext transmission might be acceptable within a secure data centre, but is never permitted across a public network. The Payment Card Industry Data Security Standard (PCI DSS) specifications for credit card handling offer a well-known example of data sanitising. Among other things, these specifications require credit card data to be encrypted, obfuscated, or deleted before storage.

Cloud providers must be PCI-compliant to handle such data, and also must have auditing measures in place to maintain that compliance. To mitigate risk, enterprises also should require insurance coverage on the cloud provider's part in the event of a data breach in the cloud, and build such coverage into any service contract.

In some cases, enterprises may have more rigorous compliance requirements than a cloud provider can meet. This isn't necessarily a dealbreaker for a given cloud provider, but it may require the enterprise to implement its own compliance framework within the cloud.

Cloud computing's benefits are real: a lower IT profile, faster provisioning, and global availability of new services. At the same time, network managers need to think carefully before making the transition. Every challenge discussed here can be resolved, but each will require careful planning before and during the move to the cloud.›

David Newman is president of Network Test, an independent test lab and engineering services consultancy based in Westlake Village, CA. He is the author of IETF RFCs on firewall performance measurement and many articles on network device performance and security. Send comments on this article to feedback@infosecuritymag.com.

Extra Layer of DEFENCE

USE APPLICATION WHITELISTING AS ANOTHER WEAPON IN THE BATTLE AGAINST MALWARE.

BY ERIC OGREN

APPLICATION WHITELISTING is an idea that makes too much pragmatic sense for it not to have appeal as an anti-malware mechanism. Intuitively, a technology operating in the kernel that detects suspicious changes in an IT-controlled software configuration should be a simpler and easier-to-scale solution than a technology that looks at all files to identify and clean all attacks known to the world.

Application whitelisting (AWL) came into the security scene several years ago with an active approach to combat the relentless success of malware infiltrating endpoints. Signature matching antivirus has been overmatched in keeping pace with the volume of new attacks. While antivirus diligently scans to detect attacks against its blacklist of malware signatures, attacks continue to sneak through, undetected by security software. In contrast, AWL validates that the programme the user requests to run is on the IT-approved software list and analyses the integrity of the programme before making an allow or block

decision. The whitelist approach of approved applications and programmes has to be considered a valuable, manageable, and effective layer of defence that can complement the attack blacklist approach favored by antivirus vendors.

Unfortunately, application whitelisting followed the path of host intrusion prevention, with vendors positioning the technology as a replacement for antivirus. This has confused enterprise security organisations and created a competitive environment where security vendors are not cooperating effectively to solve a critical business problem for customers.

There are practical ways that companies can use AWL today to improve their endpoint security.

Fortunately, there has been traction within enterprises for a coordinated defence of application whitelisting and antivirus products in the fight against malware. There are practical ways that companies can use AWL today to improve their endpoint security. And with some improvements, the technology could serve as a significant layer of a larger endpoint management strategy in the future.

NO ANTIVIRUS KILLER

The surge in malware creates expensive problems for businesses by placing regulated data at risk and disrupting IT operations to clean infected devices. Application whitelisting tries to tackle the problem based on these premises:

- **Only malware changes programmes without IT knowledge.** Malware needs to modify executable programmes to launch attacks and survive reboot cycles on the endpoint. A pragmatic alternative to scanning for malware is to simply detect changes to programmes that are not associated with patches or software upgrades.
- **Identifying compliant configurations is easier than identifying malware.** Through the first three quarters of 2010, McAfee Labs reports identifying more than 14 million unique pieces of malware (.pdf), a rate of more than 60,000 new infections per day, continuing the trend of year-over-year growth in malware. Intuitively, checking a list of valid software configurations in real-time is a smaller problem to solve than checking files for traces of malware.
- **The concept of trusted sources, fueled by feeds from software vendors, simplifies management of compliant configurations.** Platform vendors, especially Microsoft, automatically supply application whitelisting vendors with detailed information on the files contained in released software products. This relieves IT of the burden of having to figure out what is legitimate system software, enabling IT to focus on defining approved custom applications.

However, the shared belief that there must be a better way to secure endpoints led to the unfortunate positioning of application whitelisting as an antivirus replacement.

Every application whitelisting vendor believed that AWL would put AV on the road to obsolescence. Ultimately, the technology has not been able to supplant the antivirus grip on endpoint security because it does not by itself fundamentally solve the malware problem. AWL has proven to be very effective in the hands of skilled IT, but there are flaws that impact usability and security that have yet to be overcome:

- **Most organisations cannot lock down user endpoints.** The concept of locking down IT policy-compliant endpoint configurations sounds good in theory, but in practice, users need the flexibility to install applications and personalise their PCs. Too tight a lockdown of the endpoint disrupts user productivity; too light a lockdown weakens the security benefits of application whitelisting.
- **Many threats are delivered as active code through the browser and do not modify whitelisted programmes.** Application whitelisting is good at making allow or block decisions when a programme is launched, but cannot easily make decisions on active code that is delivered to the browser. The problem will get worse as users become more dependent on browser-driven applications. For example, the number of social networking users actually surpassed email users last July, according to a report by Morgan Stanley (.pdf). The browser is now the target of choice for malware developers.
- **IT security teams are forced to decide which user applications should be allowed or blocked.** IT must not only deploy and administer an additional endpoint security product, but it must also make timely allow/block policy decisions on user application requests. Although automatically allowing applications from trusted sources saves time, security teams must be willing to commit extra time for application whitelisting support.

Application whitelisting vendors have been challenged to establish AWL as a vibrant segment of the endpoint security market.

Application whitelisting vendors have been challenged to establish AWL as a vibrant segment of the endpoint security market. Lumension, McAfee, and Microsoft have integrated application whitelisting into next-generation endpoint security and management solutions, while Bit9 and CoreTrace remain as major independent whitelisting suppliers. Thus far, enterprise security teams have spoken via product purchase decisions and the verdict is that application whitelisting is finding broader appeal as a key element of a comprehensive endpoint security strategy rather than an outright replacement for antivirus.

There are important business considerations that application whitelisting has not been able to overcome. The first is that the technology is an incremental product to purchase and administer. Enterprise security budgets for endpoints are committed to antivirus, and that is not going to change with compliance mandates and the absence of reasonable alternatives. In addition, application whitelisting has been unable to overcome resistance from the antivirus industry with its lucrative subscription revenue

streams to protect. While antivirus vendors are in the business of protecting endpoints, they must be careful in not de-valuing their solutions by being too quick to embrace innovative approaches. For instance, most AV vendors will tell sales prospects that they have whitelisting; they'll also say it's not application whitelisting that makes allow or block decisions on programme launch requests, but rather a performance-enhancing technique indicating that a file has been unchanged since the last scan (so only new signatures need to be checked). It's hard to imagine many AV vendors admitting that they need application whitelisting when their business depends upon scanning for attacks. This resistance has caused confusion among IT decision makers.

BEST PRACTICES FOR THE SHORT TERM

There is no question that application whitelisting works well to protect executables, providing a defence against zero-day attacks and custom attacks that evade antivirus detection. AWL backs up AV and will detect unauthorised modifications to programmes and enforce security policy, either allowing the programme to run or blocking execution of the programme. AWL's ability to look inward towards compliant software configurations for symptoms of an attack provides a complementary layer to AV's ability to mitigate damage from identified attacks. In the short term, organisations leveraging the combined strengths of both approaches will significantly enhance their resistance to malware outbreaks.

- **Use application whitelisting to secure system-level components and antivirus to vigorously scan other programmes.** Best practices call for locking down critical software against unapproved changes, blocking execution of explicitly unauthorised user-installed programmes, and closely monitoring the use of all other programmes. Programmes delivered from trusted sources that are unmodified copies from the distribution media do not need to be scanned for attacks. Security teams can focus the separation of security powers by coordinating application whitelists with antivirus exclusion lists to reduce functionality overlap and increase performance.
- **Evaluate integrated management of endpoint security technologies.** Vendors are integrating application whitelisting, antivirus, patch management, and application intelligence into single endpoint security management consoles. An integrated approach can save administration time and effort, and also ensure that there are no gaps in security coverage.
- **Prioritise computing assets requiring application whitelisting defences.** Mission critical command and control stations, IT operations and service desk computers, and sensitive servers are more appropriate for cooperative AWL and AV solutions than devices that require a higher level of user application customisation. Start deploying application whitelisting to bolster antivirus defences on devices that are needed to keep the technical infrastructure operational, even in the face of a new attack.

AWL's ROLE IN FUTURE SECURITY STRATEGIES

The concept of a balanced approach to endpoint security with application whitelisting is compelling, with the technology evolving to support emerging endpoint security strategies. There has to be a significant role for application whitelisting to play as organisations evolve their physical devices, deploy virtualisation services for desktops, and shift their infrastructure into the cloud and handheld devices. While it is not clear what direction application whitelisting will take, these are some areas that demand attention in order for whitelisting to remain viable in the future:

- **Extend the concept of trusted sources to include applications and active code from Web downloads.** While this may sound like a tall order, electronic storefronts such as Apple's already employ a form of application whitelisting; an iPad or iPhone will not allow an unauthorised programme or modified programme to run. AWL vendors can federate trusted sources, perhaps with reputation-based services, to provide more protection against browser-based attacks.
- **Automate reporting of application intelligence.** It will take years for organisations to evolve to application-centric firewalls. However, application whitelisting already produces intelligence on actual application usage on a user-by-user basis. Reporting application intelligence derived from whitelisting through systems such as a SIEM or protocols like

Taking a Different Tack

Customers of Cisco's defunct host intrusion prevention software are turning to application whitelisting.

MANY USERS OF the now retired Cisco Security Agent are replacing their CSA host intrusion prevention (HIPS) software with application whitelisting. The HIPS promise was to correlate file, network, and operating system activity to detect the presence of attacks that evade antivirus, and leverage the IT-defined policy rule set to block further execution of the attack. With AWL, the focus shifts to protecting executable software and file, network, and system resources by blocking the ability of zero-day attacks to execute. AWL is a simpler model based on the premise that only malware makes unauthorised changes to programmes.

The critical weakness limiting broader deployment of Cisco CSA and HIPS in general is the need for IT to define and maintain a complex rule set to enforce security policy. Since IT owned the rule set, any software upgrades or new software installations would generate trouble tickets to the security service desk for re-calibration. The Cisco CSA administration effort was difficult to scale to large distributed organisations. The AWL administration burden is significantly lighter than CSA since there is no longer a need for IT to define and maintain complex rules defining acceptable file, network, and system activity.

In many ways, the application whitelisting ability to thwart malicious code fulfills the goals for host intrusion prevention. Companies that added a HIPS layer to their endpoint security to complement antivirus now have an opportunity to evolve that strategy to application whitelisting.›

–ERIC OGREN

the Trusted Computing Group's IF-MAP would provide organisations the application information they need to streamline network processing without having to refresh their firewalls.

- **Add the ability to transparently replace infected software elements.** Virtualisation allows IT teams to automatically replace non-compliant software; as software becomes more disposable, the emphasis will shift from identifying and cleaning attacks to detecting change and replacing software. Whitelisting is a technology that is perfectly suited to provide attestation services to ensure the integrity of virtualised software. In addition to enforcing allow/block policy decisions, IT would be able to automate the recovery from attacks with an additional "replace and allow" decision. The ability to replace infected or obsolete elements would fundamentally change endpoint management strategies, and it would be enabled by whitelisting's ability to detect modifications.
- **Enrich antivirus subscription services.** The winning application whitelisting vendor will find resources that can be added to AV subscription services. AWL and AV vendors have the security of user endpoints as a common interest, even though they take opposite technical approaches. The motivation is there on both sides if application whitelisting vendors can show a plan that protects the antivirus business model. Perhaps AV vendors can stream reputation scoring for AWL to act on active code requests, or AWL can upload application configurations to streamline AV scanning. Enterprises need application whitelisting and antivirus to work together; the sooner that happens the better it will be for everyone.

The winning application whitelisting vendor will find resources that can be added to AV subscription services.

Application whitelisting vendors are researching ways to add most of these capabilities in their products. Right now, though, AWL solves a hard problem of detecting the presence of unauthorised software before it can execute to launch an attack. It is not—and will never be—a replacement for antivirus. However, application whitelisting approaches will be a critical element in the evolution of endpoint security strategies. With foresight and execution, application whitelisting is well positioned to reduce the impact of malware.›

Eric Ogren is founder and principal analyst of the Ogren Group, which provides industry analyst services for vendors focusing on virtualisation and security. He previously served as a security industry analyst for the Yankee Group and ESG, and has also served as vice president of marketing at security startups Okena, Sequation and Tizor. Send comments on this article to feedback@infosecuritymag.com.

TARGET: THE HUMAN

Cybercriminals are using social engineering fueled by social media to attack users and break into companies.

BY MARCIA SAVAGE

ON THE SURFACE, the email looked completely legitimate. It appeared to come from an employee within the U.S.-based Fortune 500 manufacturing firm and talked about a corporate initiative the CEO was pushing. Four high-level executives received the email; one clicked on a link embedded in the message. That was all the attackers needed. The unwitting click unleashed malware that infected the executive's computer and gave them a foothold into the company's network, where they sniffed for passwords and gained access to multiple systems.

Until the FBI notified it, the manufacturing firm—which was negotiating to acquire a Chinese company—had no idea the intruders were stealing data on a weekly basis. The stolen data was highly sensitive—critical emails with details of the negotiations. In the end, the company scuttled its acquisition plans, says Frank Nagle, senior consultant at MANDIANT, an Alexandria, Va.-based information security firm that investigated the case.

The attack is a stark example of the kind of social engineering techniques being used against companies today. Gone are days of the mass emails with misspelled messages. Criminals today are doing more reconnaissance than ever before—aided by social networks and all the personal information loaded onto them—to craft targeted emails or instant messages that trick people into opening malware-rigged attachments or divulging passwords and sensitive information. Social engineering is a common technique used in advanced persistent threat activity—like the intrusion into the manufacturing firm—raising the stakes as coordinated, state-sponsored groups infiltrate US companies, hunting for corporate secrets.

Social engineering is a common technique used in advanced persistent threat activity–like the intrusion into the manufacturing firm–raising the stakes as coordinated, state-sponsored groups infiltrate US companies, hunting for corporate secrets.

Defending against today's social engineering attacks is difficult but not impossible, security experts say. It requires focusing on the human element of the equation with better security awareness training that gets employees to think twice about clicking on certain emails. Let's look at some of the social engineering ploys used against enterprises, what's helping to fuel them, and strategies that can help a company fend off these attacks and protect its valuable data.

SOCIAL NETWORKS: SOCIAL ENGINEERING GOLD MINE

Social engineering is nothing new in the digital age, of course, but security experts say criminals are using it more as companies have gotten better at securing their networks.

"Before it would have been easier to take advantage of unpatched systems," says Mike Murr, a certified SANS instructor and author of the upcoming *Human Compromise: The Art of Social Engineering*. "Now it's often easier for the attacker to get code running on a remote system by persuading a user using social engineering to click on a link, execute code, or enter their password.

"We're getting better at locking down the digital assets. We're not perfect, but it's to the point now where the attacker is getting more ROI on the user vector than some of the digital vectors," he adds.

A common mistake enterprise security managers make is focusing on infrastructure and system defences instead of people, says Shawn Moyer, managing principal research consultant with Accuvant LABS R&D team. "A lot of defenders still think in terms of an attacker on the Internet externally trying to find a way in. …The reality is, if I'm the outside threat, I find an insider and that insider becomes your threat," he says.

Targeting the insider has never been easier, thanks to the rise of social media like Facebook, LinkedIn, and Twitter, security experts say. Outsiders researching a company can search the sites to find out who works there, who the top executives are, what they're talking about, and their contact information: all data that can be used to personalise an attack, making it more effective.

"Information is much easier to mine," says Moyer, who conducts penetration tests for clients. "I can find out who the IT security manager is much easier in 2011 than in 1991."

"Information is much easier to mine. I can find out who the IT security manager is much easier in 2011 than in 1991."

–SHAWN MOYER, managing principal research consultant, Accuvant Labs R&D

Chris Nickerson, founder and principal consultant at Denver-based Lares Consulting, which provides pen testing and other security services, says his tests for customers will use their Twitter or Facebook accounts to collect information and successfully social engineer employees. "These are corporate-sanctioned accounts. They're huge attack vectors," he says.

People generally have become aware of broad-based social engineering attacks like spam or 419 scams, so attackers have turned to social networks to create more targeted emails, says MANDIANT'S Nagle, who is also a PhD candidate at Harvard. They'll also send messages on the social networks themselves to exploit the trusted environment they provide. "Attackers started tailoring these messages so they're no longer poorly spelled and generic," he says.

In the case of APT, attackers are particularly adept at such reconnaissance to craft targeted attacks, Nagle says: "They do their homework better."

For example, someone wanting to break into a defence contractor could first identify five to 20 employees to target, then research those people, says Lance Spitzner, director of SANS Securing the Human Program. From publicly available information on the Internet, they could find out that those employees recently attended a conference and create a spear phishing email that pretends to be follow-up from the conference.

"By customising the email, two things happen: They're far more likely to click on it, and by having a small number [of targets] it's more likely to slip through. It goes under the radar of the antivirus companies because they don't have signatures [for it]," he says.

ATTACKS ON CORPORATE DATA

Social engineering played a role in this year's attack on RSA, the security division of EMC. The attack, which RSA described as an APT, began with two different spear phishing emails with the subject line "2011 Recruitment Plan" sent to two small groups of employees, according to Uri Rivner, head of new technologies, identity protection and verification at RSA. "The first thing actors like those behind the APT do is seek publicly available information about specific employees—social media sites are always a favourite," Rivner wrote in a blog post. One employee clicked on the spreadsheet attached to the email, which contained a zero-day exploit; attackers then were able to infiltrate RSA and steal information related to its SecurID products.

Social engineering also factored into two other high-profile breaches over the past 18 months: The attack by the "Anonymous" group on security firm HBGary Federal earlier this year, and the attack on Google, which the company disclosed in January 2010. In February, HBGary founder Greg Hoglund told investigative reporter Brian Krebs that the attackers tricked a network administrator into giving up access to Rootkit.org, a research website maintained by Hogland; from there, they gained access to systems containing sensitive emails and other data.

The attack on Google's intellectual property, which put APT into the popular lexicon of the security industry, started with reconnaissance that targeted specific Google employees. As described by Heather Adkins, information security manager at Google, during a presentation at the Forum of Incident Response and Security Teams (FIRST) Conference 2010 last year, the attackers gathered information posted by the employees on social networks like Facebook and LinkedIn, set up a Web server hosting a phony photo website, then sent emails containing links that appeared to come from people the employees trusted. Clicking on the links sent them to the website, which downloaded malware and ultimately gave the criminals an opening to infiltrate Google servers.

> **"Spear phishing in and of itself is not a particularly sophisticated form of attack, but its exploitation of a person's trusted relationships can make it quite effective."**
>
> –HEATHER ADKINS, information security manager, Google

"Spear phishing in and of itself is not a particularly sophisticated form of attack, but its exploitation of a person's trusted relationships can make it quite effective," Adkins says in an email this year. "Relatively simple spear phishing can also be a gatekeeper to more complex operations like APT and others."

APT attackers usually rely on targeted email and instant messages to trick employees into downloading a malicious link or attachment, Nagle says. Another common tactic in APT is to hack into one person's account and use that person's contacts to send an email or instant message and try to break into another company. That tactic "drastically increases the chances that the second victim will click on the infected PDF file or whatever it is they send," he says.

In one MANDIANT investigation, attackers who had infiltrated a company looked through its emails and intercepted an exchange with someone who worked at another company in the same industry. The attackers then added a malicious payload to a document that was included in that email exchange.

"The second company clicked on it because it was coming from someone they already talked to, it was related to a subject they already talked about, and it was a version of a document they already opened," Nagle says. "That's not uncommon… If the APT is interested in that industry, once they've compromised one company in that industry, they'll use that as leverage to go after other companies in that industry."

He's also seen cases in which attackers will intercept IMs to break into other companies or other parts of an enterprise. "They'll use contacts on MSN Messenger, then jump into a conversation that people are having" and add a malicious link, he says.

"When email is really targeted, it's tough to come up with technical means, so you need to rely on employees to be educated and on alert for those types of things."

–FRANK NAGLE, senior consultant, MANDIANT

SECURITY AWARENESS

Security experts acknowledge that targeted social engineering attacks can be tricky for companies to combat. Keepings systems patched and updated is critical, of course, but technology only goes so far, making effective employee security awareness training essential, they say.

"Employees in these social engineering attacks are really on the front lines," Nagle says. "When email is really targeted, it's tough to come up with technical means so you need to rely on employees to be educated and on alert for those types of things."

It's important to train users to trust their instincts on any email that seems at all suspicious, experts say. "Organisations would do well to caution employees to be wary of unexpected messages or unsolicited links, even if they appear to come from friends or co-workers," Google's Adkins says. "A quick phone confirmation in suspicious cases is a much better option than becoming a victim of spear phishing."

Lares Consulting's Nickerson says there are usually always signs that expose an email as a phishing attack. For example, phishers don't understand tonality; it's usually easy to tell if someone you know wrote an email from the tone. Users can also mouse over a

link to identify it; if it looks suspicious, don't click on it, he says.

"Pay attention to the details, like the email address it's coming from, the links you're being sent to, and the tone of the email," Nickerson says.

Part of the problem is that companies have fallen short when it comes to security awareness training, which often ends up being a cursory exercise for compliance purposes, says Spitzner. "We've done tremendous work to secure computers but nothing to secure the human operating system. That's why these social engineering techniques are so prevalent," he says. "To change human behaviour, you need to educate and train employees, not just once a year but continuously. Like you continually patch computers and applications, you're continually training and patching human operating systems."

Employees who are trained to be security aware are more likely to realise if they're victimised by a spear phishing email and quickly call the security team, Spitzner says. That speeds response, which is particularly critical with APT activity, he adds.

Chris Hadnagy, aka loganWHC and operations manager for Offensive Security, which provides security and pen testing services, says companies need to create a programme that makes security awareness personal for employees with hands-on training that demonstrates how easy it is to profile them online or crack their password.

"I've heard employees say, 'What do I care, it's not my data.' Now, security awareness has become personal for them. It's not just about protecting their employer's data but their life," says Hadnagy, who also is lead developer of Social-Engineer.org and author of *Social Engineering: The Art of Human Hacking*.

Dave Marcus, director of security research and communications at McAfee Labs,

Rules of Engagement

Policies for acceptable social media use can help build a company's defences.

WITH SOCIAL MEDIA helping to fuel social engineering attacks, companies can protect themselves by backing up their security educational efforts with policies for social media use, advises Frank Nagle, senior consultant at information security firm MANDIANT.

More and more companies are realising it's important to have rules for acceptable use of social networks, not only as protection from social engineering scams but also to protect their corporate reputation and image, he says.

"They can't stop their employees from using social networks, but they can encourage them to use social networks in a responsible way that reflects on the company," he says.

He cites the IBM Social Computing Guidelines as an example of a corporate acceptable use policy extended to social media. They include protecting personal privacy, not disclosing IBM confidential data, and respecting copyright laws.

–MARCIA SAVAGE

agrees: "If you show them how easy it is to mine their own data, they'll get it."

Pen testing is another step companies can take to help protect their employees and their data against social engineering attacks, experts say.

"Red team penetration testing is important. Not just getting a vulnerability scan to pass an audit, but engaging in open scope attack simulation," Moyer says, acknowledging his bias on this front. Generic network pen tests have specific methodologies and are designed as one-size-fits all, he says. "Real hacking is a creative exercise. Whenever you have a rigid set of rules of engagement, it doesn't leave the door open for creativity. Not that a red team will find every single path into an environment, but it will find paths that your standard methodology did not."

PEN TESTING TECHNIQUES

One social engineering tactic that pen testers are using with a lot of success to break into

Social Engineering Demo

Capture the Flag event reveals corporate security weaknesses.

A SOCIAL ENGINEERING contest at last year's DefCon18 targeted 17 companies, including Walmart, McAfee, Cisco, and Apple, and all failed. Contestants were each assigned to one target and managed to get a piece of information using social engineering.

That shows that enterprise security awareness programmes aren't working, says Chris Hadnagy, lead developer of Social-Engineer.org, which hosted the Social Engineering Capture the Flag event. Companies should pay heed to the report produced after the CTF–and the report from this year's upcoming contest–to improve their programmes, he says.

Last year's final report (.pdf) was downloaded more than 300,000 times, and prompted requests from companies on how to use it. "We worked with companies to improve their security awareness programmes," Hadnagy says. "We accomplished our goal, which was to raise awareness of the threat social engineering poses to corporate America and to provide something that companies can use as a tool for improving their awareness programmes."

In the CTF event, contestants have two weeks to collect information and build a profile of the target company. At DefCon, they're given 25 minutes to call their target and collect as many "flags" as possible: information such as VPN software, type of browser, employee schedules, and food supplier. No directly sensitive information such as passwords, IP address, Social Security numbers or credit card numbers is targeted.

This year's contest at DefCon19 in August in Las Vegas featured15 targets, including two "premier targets," which Hadnagy says are companies that have agreed to work with SocialEngineer.org and be used as social engineering targets. Contestants this year were given an actual professional audit report they could mimic in their information gathering. Also, a new target ranking system was introduced; organisers didn't list the data extracted from the targeted companies but did rank how they fared overall, Hadnagy says.

–MARCIA SAVAGE

companies is what Nickerson calls "polo shirt attacks." Testers will use intelligence gathered about the client and impersonate a representative from a cleaning crew, auditing firm or other service organisation by wearing a polo shirt with a logo. The shirts are easy to have made, Hadnagy says.

"I can show up at your building, say I'm from such-and-such waste management. I heard one of your dumpsters is damaged and I need to go on site," he says. "Who is going to stop the dude with the clipboard and shirt with the logo? You find the dumpster, pull out papers and discs and load the car."

Nickerson says he's seen a number of companies experience asset losses, theft and even corporate espionage when criminals have used this kind of impersonation scheme. "Impersonation happens at all levels, and in my opinion, is responsible for a great deal of loss in many businesses, he says. Clients are beginning to strengthen their defences by checking for identification of service providers, having processes for calling the normal service representative to see who was sent for the job, and not accepting unscheduled visits, he says.

PHOTOGRAPH BY APNEET JOLLY

"Impersonation happens at all levels, and in my opinion, is responsible for a great deal of loss in many businesses."

–CHRIS NICKERSON, founder and principal consultant, Lares Consulting

When he's on a red team case, Moyer starts out with the domain name and domain name registration of the company to collect information about a company. Other tools include the American Registry for Internet Numbers (ARIN), the company's website (especially the "About Us" page), and Web forums. The most valuable piece of information he digs up is the company's email naming convention. From there, he constructs a scenario specific to the company.

For example, in a case involving a retail company, he found employees chatting on a Web forum. One worker mentioned that the company didn't offer an employee discount. Moyer created a fake employee discount programme by registering a separate domain name. Then he sent emails to about 20 employees telling them they were enrolled in private early testing of the discount programme and asked them to forward the message to five co-workers they'd like to participate. "Once I have that first click, I can pivot into the environment from that victim's machine," he says.

"I got an understanding of the employees and tailored a scenario specific to them," Moyer adds. "Most companies aren't thinking of an attacker specifically targeting them."

Hadnagy describes another case in which his team gathered data from social media sites and Internet forums to create a successful spear phishing email. The client company

had just upgraded its firewall and IDS systems; three IT workers discussed a problem they were having with the firewall configuration on sites like LinkedIn and Twitter. "Now, I'm the secretary to the CIO, here's a PDF that's a solution to that. Personal phishing attacks are very successful with the use of social media," Hadnagy says. "It makes your whole story line credible because you know something personal about them."

LESSONS LEARNED

After cases involving APT and a social engineering attack, MANDIANT often sees companies step up their awareness efforts, Nagle says. They won't divulge the details of the attack to employees, but they'll caution them to be on alert for suspicious emails. In the case of the manufacturing firm, remediation included boosting the company's ability to monitor its internal network, something Nagle says many businesses neglect.

He's seen heightened awareness really pay off, with employees spotting suspicious emails and forwarding them to IT, which sends them to MANDIANT for analysis. "We'll confirm it's malicious. That to me is great because it's education working," he says.

Marcia Savage is editor of Information Security. *Send comments on this article to feedback@infosecuritymag.com.*

INFORMATION SECURITY EUROPE is published quarterly by TechTarget Member Services, Marble Arch Tower, 55 Bryanston Street, London W1H 7AA; Toll-Free 888-274-4111; Phone 617-431-9200; Fax 617-431-9201.

See ad page 2

- Blue Coat Mid-Year Security Report 2011
- Corporate Web Security - Market Quadrant 2011
- Magic Quadrant for Secure Web Gateway

- Choosing a Cloud Provider with Confidence
- E-commerce 101: A Guide to Successful Selling on the Web
- Securing Multiple Domains with SSL