Realtime publishers

The Shortcut Guide™ To

Prioritizing Security Spending

sponsored by

IBM®

Dan Sullivan

Copyright Statement

Chapter 3: Security and the Dynamic Infrastructure

Information security practices have evolved along with advances in hardware, software, and systems architectures. Security in mainframe and mini-computer environments was dominated with issues of physical access and user account management. Distributed architectures introduced the need for authentication and authorization mechanisms that accommodated multiple clients and servers. Early widespread use of the Internet led to the adoption of perimeter security techniques, such as network DMZs. Today, those measures must be augmented with multiple layers of defenses on individual devices and network security controls, such as intrusion prevention systems (IPSs) and content-filtering applications. This pattern of changing security practices shows no sign of abating.

A number of technologies and advances in system designs are introducing opportunities to improve security as well creating the need for additional security measures to ensure the continued protection of the confidentiality, integrity, and availability of information resources. This chapter will examine four sources of change in information security:

- Virtualization
- Cloud computing
- Distributed information flows
- Web 2.0 technologies

Together, these technologies are enabling a more dynamic infrastructure than we have previously had to manage. Virtual servers can be provisioned and moved into production faster than physical servers. On-demand computing and storage resources in the form of cloud computing is changing the way we handle peak demand periods and reducing barriers to expansion, especially with regard to the need for capital expenditures. Businesses' relentless drive for efficiency coupled with improvements in workflows and information interchange are leading to distributed information flows that cross organizational boundaries. Rich Internet applications are bringing desktop application functionality to the Web but at the expense of an increasingly complex set of software components. All these technologies provide new degrees of flexibility in how we configure and manage IT infrastructure. With a focus on keeping a dynamic infrastructure secure, we will examine each of these technologies with an eye toward discerning the security implications of each.

Security and Virtualization

Advances in hardware design have provided a steady stream of increasingly faster CPUs, larger storage devices, and more efficient bus architectures for communicating between components. In theory, these advances alone should be enough to generate more productive IT operations, but that is not always the case.

Server virtualization is the process of simultaneously running multiple instances of operating systems (OSs) on a single physical computer. With virtualization, applications can run in their own virtual machine without concern for conflicts with other applications. Prior to virtualization, it was not uncommon to find conflicts when multiple applications were run on the same devices. One application might need a particular version of a dynamic link library (DLL) while a second application needed another. A security patch to fix a flaw might not be installed because the patch broke one of the applications running on that server, thus leaving all the applications vulnerable. Virtualization can eliminate such problems by isolating each application in a separate OS.

The cost benefits that drive virtualization adoption are based on several factors:

- Better utilization of CPU resources, which results in the need for fewer physical servers
- Reduced operating costs for data centers because of less demand for power, air conditioning, and physical space to house servers
- Reduced maintenance costs because virtual machines can be standardized and instances of virtual machines can be created from a single virtual machine image.

Another cost-saving factor that may be less apparent is the improved security that can come with virtualization. We can see this in terms of virtualized network protection, server virtualization, and desktop virtualization.

Virtualized Network Protection

Networks require several security and monitoring systems that can place significant demands on CPUs and network interfaces. Commonly used network protection systems include:

- Intrusion prevention systems and intrusion detection systems
- Content-filtering systems for blocking malware, spam, and spyware
- Log management and reporting
- Firewalls, including stateful and application-layer firewalls

Each of these network protection systems performs distinct functions and each has a role in defense-in-depth strategies. However, in networks of any significant size, there will be need for multiple instances of these applications. For example, a network segment for databases and application servers, another for client devices, and a third segment for Web servers may each require instances of an IPS, content-filtering system, log management system, and firewall. Unless the proliferation of these security applications is well managed, network administrators can find themselves battling appliance sprawl. Another option is to deploy these security applications in a virtual machine.

Combining best-of-breed applications in a virtualized environment allows organizations to configure a standard set of controls and deploy over multiple network segments. When there is a high volume of network traffic, running multiple network monitoring systems can require high-end processors, such as clusters of blade servers. The same virtualized suite of applications could run in lower traffic segments, such as a remote office, with correspondingly less hardware, thus providing the benefits of a standardized set of security applications managed through the deployment of a single image rather than configuring multiple applications on non-virtualized hardware The same virtualized network security systems can be used in different environments, each configured with the appropriate hardware.

Server Virtualization

Server virtualization can significantly contribute to maintaining a secure network environment by easing standardization. Consider how virtualization can help in a typical development environment.

Software engineering practices dictate the use of separate development and test environments. In large development efforts, it is not uncommon to divide development resources among distinct parts of the project. Doing so allows the team developing the user interface (UI) to run a stable version of the application middleware on one development application server while the middleware developers could work their code on a different application server without having to worry about disrupting other developers. This also allows developers to test in realistic environments without disrupting or damaging the production systems and databases.

Constant change is the hallmark of the developers' environment, but testers need a stable platform. Application testers will typically run a suite of tests on each new version of code; thus, it is essential to keep other elements of the test environment fixed. If new code is being tested at the same time developers are updating a database, it could be difficult to determine whether changes in test results are due to changes in the code or in the data.

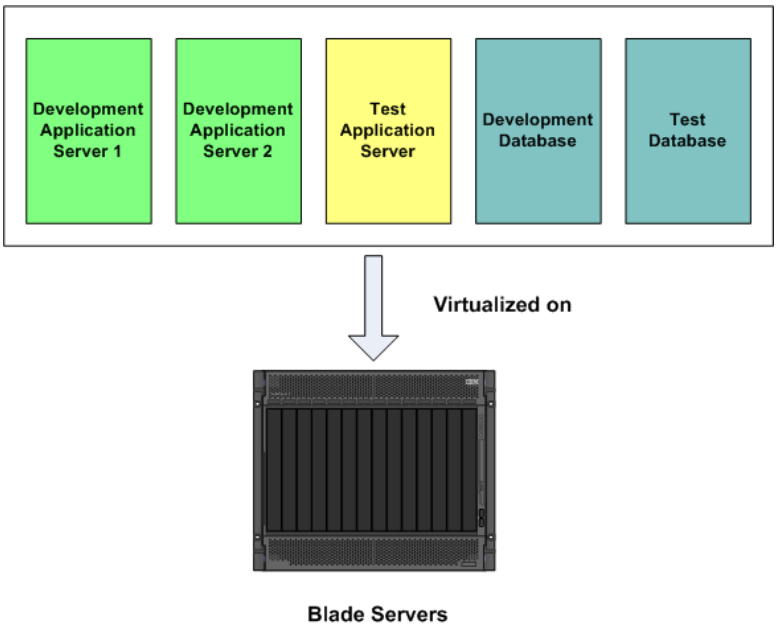

Figure 3.1: Common IT operations, such as maintaining development and test environments, require multiple servers for running nearly identical functions.

In this simple scenario, we can readily see the need for at least two development application servers, a test application server, and two database servers (see Figure 3.1). If these were five physical servers each running conventional OS configurations, systems administrators would have to configure and harden five OSs. With virtual servers, those systems administrators would have to start only five copies of the same virtual server, or more likely, three copies of an application server virtual machine and two copies of a database server virtual machine.

Systems administrators still need to contend with patching multiple instances of virtual machines. One option is to rebuild a base image with patches and redeploy; another option is to patch individual instances. The first option may be preferable when a patch is required for all instances; the second option may be preferable when a patch is only required for a subset of the virtual machine instances. The security benefits of virtualization are even clearer when we consider desktop virtualization.

Desktop Virtualization

To maintain a defensible desktop, we need to, at a minimum

- Install anti-malware software and keep it up to date
- Properly configure a local firewall
- Apply security patches for the OS and applications as they become available
- Limit access to administrator accounts
- Protect information with encryption

This list is easily manageable when dealing with a small number of desktops, but once we start to multiply these tasks by the hundreds or thousands of desktops in an enterprise, the volume of work and costs grows significantly. For example, with an asset management system, one could, in theory, push patches to devices across the network and remotely manage all these devices. In reality, some devices will be powered down and not patched, some will have insufficient local disk space and the patching process will fail, and other older devices will not support remote management protocols required for more complex remote management operations. As a result, systems administrators are left to track down failed patches, diagnose reasons for failures, and in some cases, make costly desk-side visits to correct the problems. Virtual desktops hosted in a centralized hardware environment can eliminate some of these problems.

Virtual desktops provide advantages even when asset management systems are in place to push patches to client devices. If a patch is pushed to hundreds or thousands of desktops in the middle of the night (to minimize the impact on the clients and the network during the work day), there is a good chance some of the machines will be powered down. Many of these will likely not have remote management functionality that allows, for example, remote power up. As a result, these devices will not be patched. Even in cases in which desktops are powered up, other problems, such as insufficient disk space, can disrupt the patch operation. Service support personnel are left to review the patch logs, determine which devices failed to patch correctly, and either attempt another remote patch or make a costly desk-side visit. None of this is necessary with virtual desktops.

The rapid adoption of virtualization technology is driven primarily by cost considerations, but there are added security benefits as well. A properly configured virtual machine can be deployed repeatedly with low marginal cost. Patching can be done with fewer potential glitches and at less cost. Both servers and desktops can be virtualized, enabling organizations to realize savings across IT assets.

Up to now, we have considered the benefits of virtualization as it is deployed within a single organization. Virtualization is also one of the enabling technologies of cloud computing.

Cloud Computing and the Security Issues It Presents

Cloud computing is a service model by which computing and storage resources are delivered on demand and maintained only as long as needed. This elastic model of delivering computing services is in the process of radically reshaping the economics of delivering IT services because of two key benefits: scalability without internal infrastructure and global access. For the purposes of this chapter, we will limit our discussion to public clouds and not address security issues with private clouds.

> **Public vs. Private Clouds**
>
> The term "cloud computing" can refer to either public or private clouds. Public clouds are computing and storage services offered by third parties, such as Amazon's EC2 and S3 offerings. Private clouds offer similar functionality, such as on-demand provisioning, but are provided within an enterprise. Both models have advantages and disadvantages. Public clouds require no capital expenditures or infrastructure management on the part of customers. Private clouds, which are built and managed by an organization, allow for greater control and none of the security concerns over placing data and applications on third-party devices.

Types of Cloud Computing

There are many ways to organize cloud computing offerings into various taxonomies; for our purposes, we will focus on two distinct categories: desktop software replacement services and back-office infrastructure.

Desktop Software Replacement

The time and cost of maintaining desktop software may be reduced with the advent of cloud-based desktop software replacements. With a cloud-based service, organizations may lower software licensing costs and reduce maintenance overhead.

Google Docs, Zoho, and ThinkFree Office are a few examples of cloud-based alternatives to traditional desktop software. These services provide the core functionality one would expect from a desktop office suite, including word processing, spreadsheets, presentation software, and in some cases, databases. Zoho, for example, shows how far cloud-based services can be pushed with additional support for online document management, project management, customer relationship management (CRM), and human resources applications. Virtualized desktops running office suites on servers within the enterprise have some similarities to cloud services but are a distinct model that is different from cloud computing.

Back-Office Infrastructure

Back-office infrastructure includes servers and storage arrays as well as higher-level, application-specific functionality. Collectively, the higher-level functions are described as "X as a Service" where X could include middle-tier services, such as databases and applications services or broader services such as CRM, HR, and security management.

One of the distinguishing characteristics of back-office offerings is the level of control over management and design the customer retains. Consider three example scenarios:

- Scenario 1: A customer purchases access to a servers and storage services as needed. The customer determines which OS is run on the servers, when the servers are started, how long they are run, and what level of access controls will be applied to the server. Storage is allocated as needed, and the customer retains responsibility for backup and disaster recovery. (In which case, the customer may assume that the cloud providers redundancy in the storage service is sufficient, but that is a risk-management decision that may not be appropriate for all customers).
- Scenario 2: The customer purchases a database and application server service. The customer determines which database and application server it will run and the number of instances of each. The cloud service provider manages the physical aspects of the database, ensuring space is allocated on underlying file systems and the database is sufficiently patched and properly configured. The customer designs the overall database architecture and monitors performance, but the cloud provider attends to implementation details.
- Scenario 3: A customer purchases a cloud-based enterprise application service, such as an HR management system. The customer manages data in the system and determines user access and privileges, but relies on the service provider to ensure availability of the system, appropriate backup and recovery operations, architecture and application design of the system, patching, and performance monitoring.

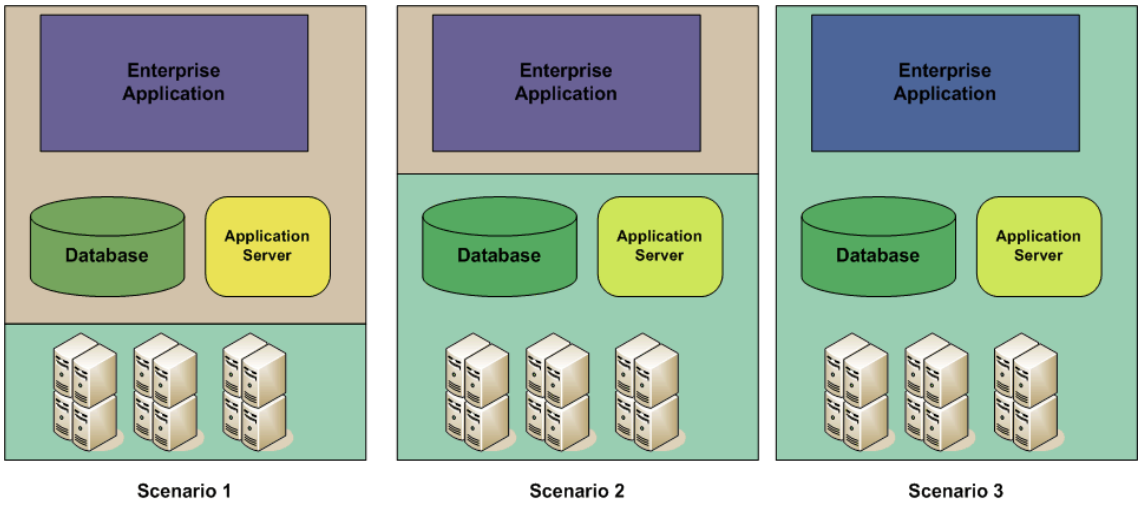

Figure 3.2: Cloud providers may control multiple layers within an application stack (shown with blue-green background) while the customer maintains control of other layers of the stack (shown with brown background). The appropriate choice is a function of several factors.

As these three scenarios demonstrate, cloud computing is a range of services defined by customer and service providers finding an appropriate distribution of labor between the two. In Scenario 1, the customer retains the most control but also has the most responsibility for developing and maintaining applications. In Scenario 3, the customer has the least responsibility for management details and, presumably, the least control over implementation details. Choosing the right combination of services is largely dictated by the customer's core competencies, ability to design and maintain IT applications, and the fit of service provider offerings to the customer's business strategy. Security considerations, in many cases, will factor heavily in cloud computing decisions.

Security Considerations with Cloud Computing

Regardless of whether a customer uses desktop software services or some combination of back-office applications and services, there are security issues to take into consideration:

- Encryption and other data security measures
- Availability and service level agreements (SLAs)
- Compliance with government and industry regulations
- Ensure cloud based applications are secure

As with the types of services offered, cloud service provider offerings can vary in their security characteristics.

Encryption and Other Data Security Measures

One of the first questions to consider about cloud security is: What could happen with your data? Confidentiality is an obvious concern and encryption is usually part of the solution when confidentiality is required. There are different ways to approach the use of encryption.

The cloud provider could encrypt data stored in its data centers. This is an approach taken by Amazon's S3 storage service. Customers generate a key that is associated with an account, and data is encrypted with that key when it is stored in the cloud. The advantage of this approach is that all data stored to the cloud is encrypted. The disadvantage, at least to some, is that the cloud provider controls the encryption process.

An alternative approach is to encrypt data locally before sending it to the cloud. This setup might appeal to those who need to maintain finer-grained controls over the encryption process, but there is the possibility that someone could upload confidential data that has not been encrypted.

Another advantage of encryption is that cloud providers would have less need to sanitize data blocks after they have been deallocated by a customer. The data is encrypted, so the next customer to use that data block, if they could read it before writing to it, would not be able to make sense of the data (assuming strong encryption and private keys of the previous user are not known to the current user of the data block). If data is not encrypted, there is more need for sanitizing storage before allocating to another user.

One way to evaluate cloud storage is to consider how the cloud provider's practices compare with the ones used with physical devices in your own company or organization. For example, when a server is removed from service, the hard drives are probably overwritten using some kind of hard drive overwrite software so that data cannot be recovered after you lose physical control of the device. Is the functional equivalent of disk overwriting available from the cloud provider?

Availability and SLAs

Highly distributed architectures, such as used in cloud environments, can take advantage of redundancy to ensure availability. If a data center on the East Coast of the US is inaccessible, customers could still access their applications and information using a data center in the Midwest. A bad controller in one disk array would not result in lost data because the same data is written to multiple other storage devices. This is the theory, at least, when it comes to availability.

In practice, well-defined SLAs trump theory. Availability and SLA issues with cloud computing include:

- The total amount of contracted downtime over some period of time (for example, per month or per year)
- The longest acceptable continuous period of downtime; downtime in excess of that presumably results in compensation to the customer
- Backup services, if any

Regarding the last bullet point, with highly redundant systems, there is less concern from losing data due to a hardware failure because the latest data can be recovered from other data blocks. There are cases in which rolling back to earlier versions of data becomes necessary. For example, if an application bug corrupts data but is not discovered for days, would it be possible to restore the data back to the last-known good version of the database?

One should also consider cases in which cloud services are not available. If a cloud-based application or data storage service is unavailable for hours or days, how would that affect operations?

An essential but much more difficult question to assess, is how likely is an occurrence of unavailability? From a risk analysis perspective, one could use past performance as a basis for estimating the likelihood of an outage; however, past conditions may not be the same as current or future conditions. Cloud providers may have many more customers in the future and have to accommodate larger volumes of data. Will their architectures continue to scale? Are there potential bottlenecks outside of their control, such as an ISP that cannot scale up bandwidth as fast as a data center needs for peak demand? Of course, serious cloud providers build redundancy and sufficient capacity into their infrastructure, but these are still questions to consider when outsourcing computing and storage services.

Compliance

Compliance issues will also require careful consideration. A CTO asked to sign off on a Sarbanes-Oxley Act compliance report will want to know their cloud provider's procedures and practices are sufficient to maintain compliance. There are a range of topics that could fall under compliance:

- Access controls to data to ensure that only users authorized by the customer have access to data
- The cloud provider offers protections to prevent potential abuse by administrators and other privileged users operating the cloud infrastructure
- When data is deleted, it becomes irrecoverable in all redundant copies and backups, if any
- Sufficient logging and monitoring is in place to meet compliance requirements

Shifting responsibilities to cloud providers to meet some of the compliance requirements on a company should be done only after ensuring the cloud provider can actually meet audit and compliance requirements.

Infrastructure Security in the Cloud

When we put money in a bank, we usually assume it is safe. Banks have developed a security infrastructure and risk management procedures that have, at least until recently, presumed to be sufficient to protect depositors' assets. Even in cases in which individual banks fail, federal government guarantees virtually eliminate the risk of a loss. Some day, we may have the same level of trust and guarantees in the cloud computing industry, but they are not in place yet. Customers conducting due diligence on cloud providers will want to understand the providers' policies and procedures with regard to physical security in data centers, access controls, identity provisioning and de-provisioning, protection for data during transmission, disaster recovery procedures and guarantees, and employee background checks, to name a few.

Cloud computing is changing the economic equation of IT services, but along with the benefits come variations on long-understood security concerns. As consumers of cloud-computing services, we need to adapt our security strategy to accommodate these new concerns.

Securing Distributed Information Flows

Another significant way in which IT service delivery has changed is the demise of traditional organization boundaries with respect to information sharing. The benefits of specialization and the ability to move information quickly and inexpensively around the globe is one of the enabling technologies of globalization. Distributed information flows are so prevalent now that we can, in the words of Thomas Freidman, view the world as flat. A business with headquarters in Chicago could have a manufacturing partner based in Shanghai, receive accounting and finance services from a company in Mumbai, look to a firm in Brussels for legal advice, and collaborate with a distributor in Buenos Aires. Once again, we have an example of a compelling economic argument for an innovative way of doing business with significant security implications. We will consider three:

- Protecting data in transit and the demise of network boundaries
- Sharing data with trusted business partners
- Employees and personal information devices

As we will see, distributed information flows must be protected at a macro level (business to business) and at a micro level (business to employee).

Protecting Data in Transit and the Demise of Network Boundaries

Data moving between organizations can give the impression that network boundaries no longer exist. This is an exaggeration, but an illustrative one. Of course, business and organizations continue to use firewalls, network segments, and other means to isolate resources. At a physical and architectural level, boundaries still exist, but at the logical level of data flows, these boundaries are more porous than a network architecture diagram might indicate. Orders can flow from a sales management system to a manufacturing partner who then transmits data to the accounts receivable system which then issues an invoice to a distributor halfway around the world.

Protecting data in a highly distributed, multi-organization system such as this requires attention to

- Data classification—Businesses need to know what data to protect. Not all data is created equal; some requires more protection than others, either for regulatory or business strategy reasons. Personally identifying information (PII), credit and financial information, and trade secret information should be governed by appropriate controls.
- Data in transit—Businesses need to know where protected data flows. Manufacturing partners may need some insight to a trade secret related to a product design but do not need customer accounting information. Information flows are dynamic, but they should not be free form.
- Confidentiality—Businesses, government agencies, and other organizations maintain substantial amounts of private information on individuals and businesses. State, provincial, national, and trans-national regulations dictate protections of such information in many parts of the world. A data breach in a Mumbai data center can have multiple implications when lost data includes information on customers from California to the European Union (EU).

Encrypting communications is one control, but knowing appropriate data classifications and implementing controls on where data flows is also required to protect data in transit.

Sharing Data with Trusted Business Partners

Sharing data with trusted business partners has similar security implications to those found when utilizing cloud computing. First, you need some way to establish who you want to share the data with. Federated identity management systems allow for this by providing the means to determine who is a trusted business partner. After you have identified your trusted business partners, there are issues associated with compliance implications and data loss prevention.

With regards to compliance, a business must understand how the data shared with business partners relates to compliance requirements. A well-formed and well-managed data classification system can help organizations understand how data flowing out of the organization should be protected. Agreements between business partners can be used to bind parties to particular responsibilities regarding data protections, including measures to protect against data loss.

Employees and Personal Information Devices

Sharing data with other businesses or organizations is just one way protected data can leave the controlled infrastructure of a business. Employees using personally owned information devices are another.

The increasing use of personal devices for work-related tasks has created something of a grey area for IT security. On the one hand, these devices are not owned by the business or government agencies, so they are not generally at liberty to dictate what device the employee should purchase, what OS to run, or the applications that the employee should use. On the other hand, individuals downloading corporate data have a responsibility to protect that data. The meeting ground seems to be that businesses should establish policies and practices that define minimum security requirements for devices that will house company data. These can include:

- Establishing polices on the use of encryption, limits on the amount or types of data that can be downloaded, restrictions on backing up corporate data from a personal device, and requirements for the use of passwords or other means of authentication on the device.
- Network security professionals can also use network access controls to prevent devices from connecting to the network that do not meet minimal security standards. This can include proper OS patch levels and up-to-date antivirus software.
- Organizations can also provide security awareness training with an emphasis on data loss prevention and social engineering attacks.

Corporate and government information is flowing more easily to devices controlled by other companies, agencies, and in some cases employees. The drive for efficiency and the willingness to adapt innovative processes will likely perpetuate and perhaps accelerate this process. Attending to the security implications is best done sooner rather than later in the adoption process.

Application Security and Web 2.0 Technologies

Innovation is introducing new security challenges through end user applications as well as in the back-office, infrastructure operations. Consider the case of rich Internet applications (RIA). The goal is to provide desktop application-like functionality in Web-based systems. Web developers are succeeding, and we all benefit with more interactive, faster responding applications. The cost (there is always a tradeoff) is an increasingly complex application stack. Data may be drawn from multiple databases, which means database connections, each authenticated to some user with a specific set of rights that must be managed and maintained. Multiple application servers may be involved, such as .NET applications and J2EE application servers providing an array of services to the UI layer. The simplistic HTML-generated UI is being replaced with interactive AJAX functionality and specialized components such as Adobe Flex. With complexity comes the risk of vulnerabilities. We shall consider four distinct types of risk areas that one should be aware of when deploying Web 2.0 technologies.

- Browser-based attacks
- Application-level attacks
- Social engineering attacks
- Need for trusted identities

The first set of areas is primarily technical, the second shows how a single individual can compromise security, and the third is an example of a broader, more fundamental need within highly distributed, heterogeneous systems managed and used by many parties.

Browser-Based Attacks

As more and more of us use browsers, it is not surprising to see an increase in browser-based attacks. These attacks can come in the form of malware and exploit vulnerabilities in browsers or browser-based applications.

Browser-based malware can infect a device when a user simply visits a site that has been compromised. Researchers have identified four ways in which Web sites can become infected with malware: through insufficient Web server security, user-generated content, advertising content, and third-party add-ons. (Source: Niels Provos, et. al. “The Ghost in the Browser: Analysis of Web Based Malware” http://www.usenix.org/events/hotbots07/tech/full_papers/provos/provos.pdf. Also see IBM X-Force Threat Reports at http://www-935.ibm.com/services/us/iss/xforce/trendreports/). Web servers can be compromised through vulnerable scripting languages. The popular Web application framework PHP has been vulnerable to attacks that can result in remote code execution and rootkit installation.

Malicious user-provided content can include scripts that are injected through anonymous comment postings and executed by unsuspecting victims when they display pages with such comments. Web advertising often depend on snippets of JavaScript code to display ads; attackers can use this to tunnel malicious content to victims. Third-party add-ons take the form of scripts that add functionality, such as advertising, page hit counters, and similar services. If a trusted third-party service is compromised, Web sites using those services can be compromised as well. What we find is that when we visit a site, we are not only trusting the providers of that site but also all the users that post to unedited forums, the advertisers on that site, and the providers of functional components for that site.

Other kinds of browser-based attacks include cross-site scripting and cross-site request forgery attacks that exploit the ways users trust sites and sites trust users. In both cases, content from the trusted party, in the form of HTML and script from a server or input from a user, can be used to perform malicious functions. A related kind of attack, a SQL injection attack, uses trusted input from a user to alter the expected response by the database management system to a query. Technically, SQL injection attacks are not limited to browsers, but these attacks have become more well known with the increased use of Web-based database applications.

Too Much Trust?

The advent of browser-based malware highlights a fundamental design assumption of the Internet: services trust other services that communicate with it. When we type a URL into a browser, we trust the Domain Name System (DNS) will return the legitimate IP address of the server. We trust a Web site to return only content we want and not the drive-by downloads we sometimes get. We make at least implicit assumptions about what we can trust and not trust every time we use the Internet. The assumption of trust certainly made sense in the early days of the Internet, and its predecessor the Arpanet, when only a small number of government agencies, universities, and research institutions used the network. Those early days were analogous to living in a small town where everyone knows each other and no one feels the need to lock their doors; now we are all living in the big city and things have changed. Today, we have spam, DNS poisoning, and Web-based malware to contend with and the assumption of trust no longer seems justified. Revising Internet protocols to require proof of identity is difficult but pieces of the solution exist in the form of more secure email and DNS protocols; much more remains to be done. See the Clean Slate Project at Stanford University (http://cleanslate.stanford.edu/) for more information about what is required for a more secure Internet.

Application-Level Attacks

As organizations improve their network and endpoint security, attackers are shifting their focus to easier targets: applications. Web applications can be vulnerable to a variety of attacks that can compromise application integrity. For example, SQL injection attacks are used by attackers to essentially "walk through the front door" and capture data from underlying databases. The fundamental problem is that application developers do not often build security into development processes (properly code applications to prevent tampering with the intended use of database queries).

As businesses move applications to the Web and adopt Web 2.0 frameworks, it is imperative that they use application security testing solutions to identify weaknesses in applications and build security testing into development. Policies and procedures must also be in place to ensure that discovered vulnerabilities are corrected and that the threat of application-level attacks is accounted for in incident response plans.

Social Engineering Attacks

The blending of private and public spheres is having some unanticipated consequences for businesses. Social networking allows individuals to share their interests and aspects of themselves with friends and colleagues. Services such as Facebook, Linked In, and My Space are well known and widely used. An unintended consequence of using social networking is that information about individuals can spill out beyond the group of trusted social network friends. Phishers, hackers, and other scammers can use this type of information to conduct targeted phishing (also known as "spear phishing") attacks.

The problem, from the business perspective, is that personal information about an employee can be used by a phisher to build trust that is then exploited to have the employee disclose business-related information. There are so many sources of information, an attacker could lure multiple individuals into each, disclosing distinct but related pieces of information. The attacker could then put the proverbial puzzle pieces together to understand their target. Social engineering attacks, by their nature, are designed to circumvent technical controls. Training employees about security risks and how to mitigate those risks should be part of any enterprise security strategy.

Occasionally, we hear a skeptic claim that if security training were going to work it would have worked by now so don't bother. They are wrong. Do not dismiss security awareness training because (1) some do it poorly and live with the results and (2) it is not a panacea. There is no silver bullet in security, and that is why we have practices such as defense in depth. Training that is engaging and relevant will do more to mitigate the threat of phishing and other social engineering scams than sticking our heads in the sand or giving up on training because we don't get it right the first time.

One of the reasons phishing attacks can succeed is that we use a variety of cues to verify the trustworthiness of an individual. Those cues probably evolved over time in social environments that do not translate well to digital communications.

Increasing Need to Trust Identities

Web 2.0 technologies are highlighting our need to trust identities and authentication procedures. Today, a social networking site will trust that I am the account owner I claim to be if I know the proper username and password. Today, banking customers trust they are communicating with their bank if their personally selected image is displayed during the logon process. This is not enough. Today, we use services that combine services from multiple sites using mashups of multiple services. We need to be able to trust not only the Web site we navigate to but also that the service provider has sufficiently authenticated the service providers it depends on.

The move to more distributed services, the use of cloud computing resources, broader flows of information across organizational boundaries, and increasingly sophisticated attacks will promote an increased need to trust identities. Financial transactions, government services, healthcare, and other applications have moved to multi-factor authentication with biometrics, smartcards, and other security capabilities to provide greater assurance when establishing identities.

Summary

The evolving patterns of IT infrastructure showcase two characteristics: less fixed infrastructure and greater use of on-demand resources. Ultimately, these characteristics are a product of the search for more efficient means to operate businesses, governments, and other organizations. Virtualization allows for more efficient use of computing and storage resources. Cloud computing allows businesses to satisfy the need for spikes in computing and storage without maintaining infrastructure sufficient for peak demand. Dynamic information flows across organizational boundaries allow for specialization and a focus on core operations. Web 2.0 technologies are essential to bringing rich application interfaces to the Web, which in turn will allow more efficient delivery of services to end users. All these benefits have associated security implications. How well we address those security issues will ultimately influence the benefit we realize from the new, more dynamic IT infrastructure.