Realtime
publishers
"Leading the Conversation"

# The Shortcut Guide™ To

# Securing Automated File Transfers

*sponsored by*

ssh® tectia™

Ed Tittel

## Copyright Statement

# Chapter 2: Planning Secure File Transfer Deployment

The first chapter set the background and laid the foundation for concepts explained and explored in this chapter. As you now know, today's typical enterprise file transfer environment is a result of long-term development and deployment of multiple forms of data transfer, often involving different hardware and software platforms. Typically, ad hoc solutions focus around the lowest-common-denominator or principle of least effort, which means leveraging native file transfer utilities to facilitate data transfers. Many times these tasks incorporate scripted interaction and scheduled automation to create periodic transfers for the most common and repetitive forms of shared or archived data.

FTP's popularity stems from its seamless platform integration, lack of interoperability issues, and nearly ubiquitous presence on everything from Windows workstations to IBM mainframes. Although this sort of "anytime, anywhere" availability seemingly makes FTP well suited to the task of file transfer, the FTP protocol and applications leave much to be desired in terms of centralized management, user accountability, data security, and real-time event auditing. Occasionally, this sort of convenience pays off in the near-term only to produce long-term security consequences by which confidential data can be revealed, intercepted, and possibly even tampered with by intermediate or external parties. There is also the possibility that an organization's FTP servers can be subverted for use as third-party storage repositories for illicit content or controlled as proxies for external network reconnaissance techniques. Eavesdropping, snooping, and hijacking attacks are three primary vulnerabilities that face improperly secured FTP installations.

**Security Breaches**

Security breaches happen all the time. Some instances are quietly dealt with as they occur, whereas others are too high profile or involve high-visibility targets (such as major corporations, organizations, or individuals) to avoid publicity. Some such situations now occur as a result of mandated reporting when security or customer confidentiality is breached, as the following examples reveal.

Mistaken postings of data occasionally occur, resulting in the disclosure of what should be confidential information. In January 2007, for example, an announcement from the Ohio Board of Nursing indicated that names and Social Security numbers for more than 3000 newly licensed nurses had been posted twice on the agency's Web site because of errors in sanitizing a posted report. Although other Ohio state agencies have elected to pay for a year of credit monitoring for those with compromised information, the Board of Nursing declined to do so.

The TJX Companies disclosed on January 17, 2007, that it had discovered the theft of credit and debit card information from its systems, which could affect customers of stores such as TJ Maxx, Marshalls, and others for whom the company processes card-based purchases. The intrusion resulted from a compromise of the companies' networks that handle credit card, debit card, check, and merchandise return transactions.

The University of Texas at Dallas discovered a computer attack on January 19, 2007, that may have resulted in exposure of Social Security numbers and other personal information for as many as 35,000 current and former students as well as faculty and staff members at the institution. This included 29,000 card holders at the University library, and a mixed group of 6000 other students, faculty, and staff. Though they have declined to discuss the details of the penetration, the university has indicated that it has taken steps to strengthen system and network security to prevent further attacks in the future.

These breaches do not account for other sources of exposure in previous years, such as malicious system cracking or poorly configured systems. For an excellent source of information about security breaches in general, please consult the archives of the Privacy Rights Clearinghouse at http://www.privacyrights.org/ar/ChronDataBreaches.htm.

Proper planning, implementation, and execution of secure file transfer methods, especially where compliance considerations are concerned, requires a multi-layered approach and thorough attention to detail. This chapter discusses several crucial points to help IT administrators address the issues involved in implementing secure file transfers within the network infrastructure.

## Implementing FTP Inventory

It's not difficult to inventory plain-vanilla FTP transfers simply and easily by using a network protocol analyzer, but this activity may be better facilitated by way of a more robust traffic monitoring application such as Cisco Systems' NetFlow or the Multi-Router Traffic Grapher (MRTG). NetFlow is an open, but still proprietary, network protocol developed by Cisco Systems to operate Cisco IOS-enabled equipment for collecting IP traffic information.

✎ IP Flow Information Export (IPFIX) is described in Request for Comment (RFC) 3917 and other documents under development under the aegis of the Internet Engineering Task Force (IETF) IPFIX Working Group, as documented at http://tools.ietf.org/wg/ipfix/. IPFIX promises to replace NetFlow as the premier source for network traffic monitoring, analysis, and reporting, and is attracting interest from vendors other than Cisco.

MRTG is a freeware/sponsorware application that uses the Simple Network Management Protocol (SNMP) to collect statistical IP traffic information for accounting purposes. Both approaches have their unique strengths and merits and can report on the same type of information, which is useful when it comes to characterizing traffic and identifying traffic trends on any given network segment.

Several FTP-related or specific tools are also available. These include the AWStats log file analysis tool that is a member of the SourceForge family of projects (http://awstats.sourceforge.net); likewise, other free log analysis tools with strong FTP capabilities include Analog (http://www.analog.cx) and Webalizer (http://www.mrunix.net/webalizer). All these tools can report on who has been accessing an FTP server, for how long, and which files have been accessed, uploaded, or downloaded. Software Assist also offers free FTP analysis tools that include FTP Audit, which discovers all FTP servers in an enterprise or organization to help identify potential exposures, as well as FTP Sampler, a tool that may be used to dig into file transfers, identify sources, and document activity, including failed transfers and potential performance issues (http://www.softwareassit.net).

Taking one or more of these approaches will help administrators identify and evaluate file transfer activity on their networks, and provide a starting point for investigating sources of such activity. NetFlow can record a wealth of information about the network activities it tracks and report highly descriptive details about them, including statistical analyses, source and destination ports and addresses, protocols and flags, service values, duration periods, and more. FTP-specific tools, however, will provide the details that administrators need to identify and handle obvious, non-covert use of file transfer tools and utilities. Together, these tools combine to provide a complete picture of the situation.

As Chapter 1 stated, FTP may be characterized by a few predictable parameters through which all traffic can be detected, identified, and analyzed. The ease or difficulty in recognizing and identifying this traffic depends on placement of monitoring equipment and application behavior on the network under observation. Not all file transactions are FTP-based, however. Some transactions do not even employ a client-server model, opting instead for peer-to-peer arrangements whereby endpoints operate in an ad hoc fashion, thereby merging client and server roles and switching back and forth between them. In fact, some peer-to-peer services are structured to switch randomly among multiple sources for data during file transfers to purposefully mask that very activity!

It is worth recalling that TCP ports can be either permanent or short-lived. Servers listening specifically for FTP traffic may be permanent in the sense that a port is always kept open, awaiting inbound connections from calling client applications. That said, active and passive transfer modes directly influence the nature of any particular FTP exchange.

> 📖 Refer back to Chapter 1 for an illustrated explanation of these modes of operation.

Client-side ports remain open only for the duration of a transfer and are not bound to a specific "listener port" as are server-side applications. However, server-side ports may also be temporarily allocated to some higher-numbered port, following the long-standing UNIX tradition of using low-numbered ports (below 1023) for privileged processes only. Depending on the transfer method employed, an FTP server may also make a temporary port allocation for the duration of an exchange, which serves only to complicate spot-check analysis. For this and other reasons, real-time traffic monitoring solutions are described and discussed later in this chapter.

On-site network appliances are ideal for observing localized traffic usage patterns where file transfers occur, but this is a naive perspective at best. Large-scale site-to-site file transfers may encompass many endpoints across multiple disjoint regions, which may be controlled by an off-site management application that negotiates cross-site server-to-server exchanges. Several combinations of software and hardware must often be employed to identify network-driven file transactions.

For example, Javvin's Packet Analyzer is a Java-based application that can track and report on four specific types of traffic, two primarily interesting types being FTP and HTTP. The Packet Analyzer permits an administrator to view reconstructed content and offers a simple tabbed interface and free trial-use period to evaluate its capabilities. Wireshark (previously known as Ethereal) is a free full-fledged Ethernet network protocol analyzer with a variety of filtering, sorting, and dissection capabilities as well as advanced analysis features. The former is a commercial product written in a cross-platform language, the latter is a free but fully capable cross-platform solution. Both provide the ability to observe and monitor network traffic. There are many similar solutions available on the open market, each with its own options, benefits, and drawbacks. What remains consistent across all of them is the training necessary to make best use of whatever products and technologies are acquired to capture and categorize network traffic by site or location.

### *Capture and Characterize Network Traffic by Location*

Knowledge about where organizational units and their employees are housed is essential when determining the role and purpose of any given network segment. In other words, it's important to know where network segments are located; which of them have desktops, servers, and other devices attached to them; what they use them for; how those segments are attached into the network core; how they obtain Internet or wide-area network (WAN) access; and so forth.

For example, a financial institution may organize its accounting department assets into logically and physically separate office spaces and network segments, perhaps tied together through a local switch, itself attached to a corporate backbone. Web-based consulting services may also be separated internally and externally either as standalone internal servers and services, or perhaps operating in a demilitarized zone situated between the internal firewalls that guard the corporate network, and the external firewalls or routers that attach to one or more wide-area links. Every department also typically has its own separate physical and logical access to different kinds of confidential client data. Therefore, each department must operate under the same canopy of compliance but with slightly different procedures.

Traffic captures should be characterized according to their points of origin. You must give special attention to how and where confidential information passes through each department. Web-based transactions may seem transient and sporadic, while internal database accesses often appear to be more deterministic and predictable. Not only should ongoing traffic monitoring be your guide to finding file transfer activities and users, local system accounting and log files also leave an electronic trail that you must examine to establish usage types and trends for each given server.

On BSD, Linux, UNIX, and Windows hosts, an application such as Sawmill can process log files and generate dynamic statistics reports. Sawmill can parse UNIX FTP logs, import them into Symbolic Query Language (SQL) format in bulk, and report on a variety of details including date and time of a given transaction, file type, host name, domain name, geographic location, authenticated user, transfer mode, and the direction for a given data exchange (or file transfer). This kind of information arms network administrators with exactly what is needed to evaluate file transfers as they occur locally on any machine under their purview. Other FTP-specific log file analysis tools such as AWStats, Webalizer, and FTP Sampler might also be brought to bear, as described earlier in this chapter. The Institute of Internal Auditors (IIA) has also published a very informative article entitled “11 Steps to an Effective FTP Audit” (http://www.theiia.org/ITAudit/index.cfm?catid=21&iid=513) that provides an overview of this entire process and includes a list of default FTP server log files that administrators could (and should) search for along the way. This kind of information can be quite illuminating, particularly when used in tandem with a suitable log analysis tool.

### Perform a Port and Protocol Inventory

A port and protocol inventory is best handled by network traffic graphing and monitoring utilities placed at major junctions of the network, such as enabling NetFlow on WAN-attached routers or at key infrastructure cross points, or by using MRTG to interrogate SNMP-capable network devices. Host-based protocol analyzers such as the robust open source Wireshark application can also perform spot checks, providing additional traffic flow detail and analysis reports for file transfer–related protocols and ports in use on any given segment. Wireshark also offers numerous relevant protocol dissector plug-ins that can assist network enumeration procedures and traffic identification, characterization, and security checks.

The objective is to determine what ports are in use on the network segment you're characterizing, then to zero in on protocols or ports that can indicate file transfers may be occurring or involved in observed activity. The idea is to identify all potential file transfers, to determine their sources and destinations, and ultimately to identify the applications used to perform such transfers. A solid, comprehensive port and protocol tool is invaluable when it comes to tracking down this information.

### Trace Items of Potential Interest

Large-scale one-to-one file transfers, such as explicit FTP transfers, are the most obvious manifestation that you're seeking throughout this exercise. But one-to-many mappings such as peer-to-peer transfers, which may involve a single peer client on one end but any number of peers on the other end, tend to follow non-linear and almost non-deterministic patterns. Such is the case with Torrent-based peer-to-peer transfers that can select portions of a single file from numerous active BitTorrent participants and effect transfers from all of them in no particular order. This explains why identifying certain types of file transfer can be more difficult than others. It also explains why you should note anything that stands out against a normal network traffic baseline, such as the sudden and unexpected appearance of BitTorrent metadata files (which typically end with a .torrent extension) and for which related software typically listens for download requests on port numbers in the range from 6881 through 6889 then work through each such annotated transaction manually. This also helps to affirm the contention that identifying applications in use, or associated with specific ports and socket addresses, is also quite important.

### Track Unknown Items Back to Their Original Locations

Be aware of any unidentified sources of traffic. Not every protocol is well known, well behaved, or sufficiently well established on the network to make enumeration easy. There may be some proprietary protocol active on your network that lacks a proper protocol decode module for your network analyzer of choice, or that uses non-standard ports for communications. Where a packet trace yields little evidence or information to describe the kinds of network activity that are underway, protocol analysis based upon specific circumstances takes on a more personal and perhaps more intrusive and hands-on feel. You must trace difficult transactions back to their points of origin and isolate one or more involved endpoints, then zero in on all active services in each system's process space. When all attempts at identification and characterization fail—and they sometimes will, though not often—you can always block the protocols or hosts involved and wait to see what kinds of complaints may be forthcoming. Obscurity of this nature is seldom associated with mission-critical applications, but it's always wise to check first with in-house developers just to make sure what you're observing isn't simply evidence of a poorly documented internal application before shutting down such traffic, even temporarily.

### Identify All Types of File Transfer

For the most part, inventorying most file transfer protocols and applications is pretty straightforward. (To help make this process easier, this guide provides an appendix that itemizes all well-known file transfer and peer-to-peer protocols and related ports.) Many forms of file transfer contain obvious traffic signatures that may be readily identified by protocol analyzers and protocol analysts, intrusion detection systems (IDSs), firewalls and traffic monitoring appliances, and so forth. Ultimately, ports, protocols, and payloads are the three key elements by which file transactions may best be recognized and identified.

That said, not all forms of file transfer are neatly parameterized and clearly defined on the network. One easy way to evade simple detection by a network protocol analyzer is for client and server applications to choose non-standard ports and then to encrypt end-to-end communications. Services may run only for short periods of time, awaiting the arrival of some predetermined sequence of packets (called port knocking) or some other signifying event before negotiating any active connections. But such behavior is often indicative of malign forces at work and should trigger a full-fledged security investigation once they're detected. Most routine, normal file transfers will be entirely overt, and much easier to recognize and identify.

Many regulations require that data shuttled across public networks be analyzed for risk of exposure to PII and that any at-risk data be encrypted to meet security standards to avoid its exposure to unauthorized third parties. In some cases, this may involve encryption of data while it's at rest inside a file system somewhere and again whenever it is in transit from an authorized sender to an authorized recipient.

## Perform a Platform Inventory

Proper documentation is part of the compliance process required by many regulations. This task includes describing and detailing all the security controls, policies, procedures, and facilities used to ensure compliance. Documentation for organizational hardware and software assets is crucial to understanding your security landscape. It's essential that you consult such documentation as part of your research process, and that you be prepared to amend and update such documentation as you conduct your own research (or at least to submit such information to the proper authorities).

### *Security Documentation*

Documenting security information is crucial to standard IT management and should have been performed as a standard of due care already. Establishing a comprehensive, up-to-date inventory of hardware and software components related to security also assists in the detection, discovery, and development of your IT infrastructure and the activities that take place upon it. This means an inventory of devices and applications down to a highly granular level of detail including the directories, files, permissions, registries, databases, and computers involved, along with all relevant configuration parameters. Know what hardware and software platforms comprise the corporate environment and pay particular attention to software versions. Each operating system (OS) has its own methods of implementing security controls and may be subject to weaknesses or exposures depending on patch, hotfix, or update levels in place. Although the concept of security and compliance is common to all platforms, the components used to implement or verify compliance will typically vary on a per-platform basis (if not from situation to situation, as is very often the case).

### *Compliance Management*

A compliance management solution such as Cendura's Cohesion software utilizes the Control Objectives for Information and related Technology (COBIT) framework and assists IT staff and auditors in the design and review of IT process controls supporting IT governance and regulatory auditing initiatives. Cendura provides continuous support for planning and organizing IT environments, acquiring and implementing program development, delivering and supporting application operations and data access, and monitoring and evaluating these environments. Their software solution, Cohesion, can inventory and audit application infrastructures, compliance management automation, policy compliance enforcement, and many other relevant aspects of regulatory compliance controls.

### Identify Platforms and OSs in Use

UNIX and Windows are among the more common platforms deployed in anywhere from small business to large-scale enterprise network environments. In addition, mainframe platforms, typically the IBM z/OS systems, play an important role in the enterprise data transfers and storage. The following list highlights several names for the many OSs found in modern network environments. This list is not meant to be comprehensive and encompasses only more typical OSs found on most networks.

> ✎ What is interesting to note is that most of these platforms include built-in FTP and other file transfer clients, and nearly all of them support similar server-based applications as well. Thus, the means and the opportunity for file transfers are invariably present, which also means they can sometimes show up in interesting and unexpected contexts.

There are many variations of each UNIX and Windows-based platform, many of which are still in use today. Some of these operating systems include:

- IBM z/OS hosts and environments
- UNIX and Linux—FreeBSD, NetBSD, and OpenBSD; Debian, Redhat, SuSE Linux (among others); Mac OS X; AIX, HP-UX, SCO, Solaris, System V; OS/2 Warp
- Windows—9x, NT, 2000/2003, XP, and (more recently) Vista hosts
- MacOS—There is the original, proprietary MacOS in any of a wide number of versions (1.x through 9.x)

The least intrusive method for inventorying these platforms involves consulting administrative documentation that enumerates and describes their installations (where applicable and available, as will be the case in environments in which such documentation is both present and current). In the absence of complete, current, and accurate documentation, you must turn to either passive or active means for fingerprinting endpoints and performing platform discovery that require observation of network traffic (passive discovery) or some sort of stimulus-response interrogation (active discovery).

Two open source tools used for large-scale fingerprinting of network endpoints include the passive OS fingerprint application (p0f) and network mapper (nmap), both very popular in their respective camps.

> 📖 For more information about the passive OS fingerprint method, p0f, visit http://lcamtuf.coredump.cx/p0f.shtml. For information about nmap operation and usage, see http://insecure.org/nmap/download.html.

With p0f, a host is discovered during initial packet synchronization (SYN) and identified by default start-up values contained in these packets. p0f can determine the distance to a remote host and the structure of a local or foreign network.

In contrast, nmap seeks to interrogate a remote host actively and performs network discovery differently. That is, nmap uses a battery of tests to provoke the remote network stack into revealing specific properties that are more-or-less unique to each platform. Thus, nmap can typically determine a multitude of things right down to a particular version and patch level for a given OS. The nmap application is fully capable of fine-tuned probes that can distinguish among OS versions and discriminate between workstations and servers as well as routers and printers, among various other network appliances and intermediate devices.

### *Identify Secure File Transfer Solutions*

There are many types and variations of secure file transfer solutions available on the market, each with its own advantages and disadvantages that must be weighed and considered carefully on a case-by-case basis. What is applicable for one organization may be deemed unsuitable for another. In fact, there is no one-size-fits-all solution.

A variety of secure transport implementations, for example, can provide consolidated and centrally managed systems for the management of secure site-to-site and application-to-application file transfer activity. Such exchanges occur over FTP, SFTP, HTTP, HTTPS with Transport Layer Security (TLS) or Secure Socket Layer (SSL), secure shell (SSH) via SFTP and Secure Copy (SCP), and the Applicability Statement 2 (AS2) specification. They can operate across a variety of platforms, are compatible with multi-tiered networks, comply with both SSL and Lightweight Directory Access Protocol (LDAP), operate with Active Directory (AD) file routing and handling, and integrate with Java 2 Enterprise Edition (J2EE) applications. Perhaps most important, such secure transport implementations also typically conform to industry, corporate, and government regulations including the Sarbanes-Oxley Act (SOX), the Gramm-Leach-Bliley Act (GLBA), and the Health Insurance Portability and Accountability Act (HIPAA).

AS2 is a specification that describes the safe, secure, and reliable transport of data via the Internet as defined in RFC 4130. This specification describes how to transport data safely and reliably across the Internet, where data consists of electronic data interchange (EDI) messages or any other message type. AS2 envelopes a message using digital certificates and encryption methods for transport between client and server application software to safely push confidential data across the network or Internet. AS3 is a current working standard that combines AS2 capabilities with FTP and Secure Multi-purpose Internet Mail Exchanges (S/MIME) to establish secure real-time connections between business partners using SSL.

📖 AS2 RFC 4130 may be perused by visiting http://www.ietf.org/rfc/rfc4130.txt.

Disadvantages to an AS2 approach are that endpoints rely on static IP addresses and fixed locations, imply specific firewall considerations, and require equally specific administrative expertise. These drawbacks automatically rule out the roving nature of mobile computing platforms and infrastructures. Endpoints cannot pull data or continue failed or truncated file transfers and require certificate management to establish secure connections. The costs of AS2 software can be high but may offer low return on investment (ROI) for high transaction volumes. Also, AS2 only works across native TCP/IP networks; thus, older X.25 and ISDN-based technologies can't use its services. Nevertheless, plenty of high-profile companies such as Wal-Mart have opted for AS2-based technologies to replace their former reliance on dial-up technologies for purchase orders that occur at periodic intervals to avoid the otherwise costly implementation of a value-added network (VAN).

The Odette FTP (OFTP) specification offers an alternative to AS2. OFTP 2.0 describes secure transfer of business documents via the Internet and includes X.25 and ISDN networks. The details of OFTP v1.3 can be found in RFC 2204, while v2.0 is currently available as an Internet draft. OFTP works either point to point or indirectly via a VAN, where a single OFTP entity can place and receive data, exchange files in either direction, and operate in push or pull mode, unlike AS2. Furthermore, OFTP can encrypt and digitally sign message data, request signed receipts, offer high levels of file compression, and, when run over TCP/IP networks, make delivery using TLS. OFTP's advantage for business applications stems from its origin by and for the automotive industry for e-Business communications and engineering data as a protocol for automated EDI. Although it was initially designed primarily for the automotive industry, OFTP can also be found in retail, petrochemical, and tax submissions and banking sectors, among others. OFTP 2.0 provides session and file security and secure authentication and offers more robust capabilities than its AS2 counterpart.

> 🕮 OFTP RFC 2204 resides at http://www.ietf.org/rfc/rfc2204.txt. To read the latest Internet draft, dated September 2006, visit http://www.ietf.org/internet-drafts/draft-friend-oftp2-03.txt.

Here again, interoperability is a key ingredient. Interoperability within the network infrastructure is as crucial to business operations as site-to-site interoperability across multiple network infrastructures. The protocols, specifications, standards, and hardware must share one or several common denominators to pull everything together. Such a process is much more involved with spoke-hub scenarios such as health care payer-provider derived systems versus more collaborative peer-to-peer platforms.

### *Look for Intersection Across Platforms*

Intersection across platforms refers to common components or capabilities that are not unique to one platform but ordinary to many. One such major intersection is in the TCP/IP network stack utilized by many networked computer systems from BSD to Windows on notebooks and laptops, workstations and servers, as well as network appliances and other widely used devices. Basic connectivity between any two machines can be created where TCP/IP protocols and equipment provide a common denominator. That said, interoperability issues are likely to arise in the following key areas:

- Authentication methods
- Communications protocols
- Software applications
- Data and information

## Communications Protocols

Systems should be able to pass and receive information in some generic manner without any pre-programming of sender or recipient that is party to such transactions. The parameters for both originator and destination endpoints need not be explicitly known by both parties to create ordinary transactions, just as a Windows workstation client should not particularly care that a UNIX server acts as its host platform. Both platforms are fluent in TCP/IP as a primary bi-directional communication scheme and quite capable of sharing files in a number of ways, most notably FTP.

Where EDI transactions are concerned, such as in the health care industry, this may require an advance arrangement that employs one or more designated communications protocols to permit out-of-the-box interoperability between source and destination. This sort of interoperability is not covered in HIPAA and indeed there is a significant gap in the health care industry regarding an all-purpose protocol framework within which to achieve such standard communication. A direct side effect of this set of circumstances is that each trading partner (a covered entity—or CE, in HIPAA terms, meaning a pharmacy, any health care provider, and any insurer that handles medical claims) necessitates an inefficient validation process for each and every implementation as new partners come online.

## Software Interoperability

Software interoperability ensures that two or more products are tested for compatibility and provide an interoperable user interface and interoperable commands; likewise for scripts, macros, and data files. In our imperfect world, we cannot always have access to development procedures, standards, and specifications, nor will the most extensive testing always make all interoperability issues apparent.

Here again, FTP's success owes to its seamless interoperability across different hardware and software platforms. This popularity is heavily battle tested and has evolved from years of development and deployment in large-scale university and enterprise network environments, though as we've already observed repeatedly, FTP is also not without its drawbacks and limitations. But this same desirable all-purpose characteristic is essential for any secure solution destined to fill the role of an enterprise-grade replacement file transfer solution.

The secure transport approach described earlier in this chapter works across disparate systems using platform-independent open standard protocol specifications, as do SSH-based approaches. These replacements simultaneously eliminate vendor-specific tie-ins to any architecture, infrastructure or framework and extend security and regulatory compliance coverage to virtually all current platforms in a business or enterprise environment. Financial transactions, critical business files, large documents, and EDI transactions can then take place securely over the Internet and across private IP networks. This setup permits organizations to supplant expensive legacy methods of data transfer and to avoid the potentially costly establishment of virtual private networking (VPN) technologies, leased-line subscriptions, or VAN deployment—and particularly avoid unsecured FTP transactions.

## Data and Information

Last but not least, you must also consider the importance of data and information interoperability. Particularly where different platforms, interfaces, and applications are used, generic or portably defined documents and information specifications provide a basic framework for interoperability. At the data level, where data produced by one product is readily understood by a product designed for an entirely different platform, both products can share and process common information. Information passed within and about this data can also be produced and presented in platform-independent ways.

Email is a common point for information exchange, where interoperable clients exist for virtually every kind of computing platform. An email composed on a Windows-based Outlook client can be formatted, delivered by a UNIX-based SMTP server, then read on a Linux-based Thunderbird recipient. Defining appropriate data sets, assessing the correct level of data representation, and evaluating the right solution are all part of the same interoperability examination process when it comes to file transfer.

## Making a Technology Selection: Performance Issues

Several high-profile data theft cases in the recent past sent an industry-wide rude awakening that underscores the need for heightened awareness and better security procedures for storing and transmitting confidential information. Perimeter security is not enough; when data passes beyond the scope of perimeter-based security mechanisms, it can be exposed to tampering, interception, and possibly even modification or manipulation. Tampering can happen at any time—an unattended notebook or external drive stolen off-site, an unsecured connection on a foreign network, or even a momentary physical access by an unscrupulous third party all provide opportunities for tampering. Not to mention bad things happening, by accident or maliciously, by authorized network users inside the perimeter.

These examples may seem naive at first blush, but consider a traveling IT professional who will pass through numerous airports and cabs or drive multiple rental cars. At any point during this sequence of activities, his or her notebook could be left unattended, only to be stolen by some opportunistic criminal. A publicly accessible wireless hotspot or Internet cafe may make a perfect temporary arrangement to access the Internet and upload or download information, but may also provide eavesdropping opportunities or involve connections through rogue or attacker-modified access points. And if his or her notebook should break on the road, there is always the possibility that the notebook will end up in a third-party technical service department, where unauthorized outsiders could dig into its contents.

It may be unrealistic to audit where and when all pieces of confidential information are left exposed, even with the strictest auditing regimes and controls in place, so it may never be fully known what confidential data remains unsecured on a notebook's drives at any given time. The one real certainty is that only data stored or transmitted in a cryptographically secure format can provide adequate protection against leakage to or access by unauthorized parties.

When it comes to implementing cryptographic solutions within a business environment, one of the first queries is how they address business needs. First and foremost, an encryption solution will resolve most rigorous compliance issues with regard to data storage and collection. Ensuring privacy for credit card transactions is exemplified in Requirement 3.4 of the PCI Data Security Standards (DSS), which requires that sensitive cardholder data be rendered unreadable anywhere it is stored. That said, many regulations also provide a "safe harbor" clause for companies that implement encryption. When a breach of security occurs, companies acting in good faith will have limited or no liability for such incidents.

Performance arises as another issue for encryption deployment, especially where large-scale enterprises or multiple-client organizations are concerned. This issue is discussed further in the following section.

### Costs of Encryption and Compression

The performance of a given cryptosystem or encryption solution will be directly related to the implementation approach. There are host and application-based encryption solutions that burden local general-purpose processing resources during normal operation. A CPU already under load from other tasks will be greatly impacted by localized software-driven encryption methods that rely on its processing capability, unlike an appliance or storage-based solution that exports this burden to specialized standalone hardware resources. The first solution will affect notebook, desktop, or workstation computers at best; at worst, software-based encryption often requires significant additions or major overhauls to existing server hardware.

In general, the following types of encryption solutions are available to plug existing gaps in data security and to help meet compliance requirements for data privacy and confidentiality and customer security:

- Host and application-based solutions are implemented on each at-risk computer that requires cryptographic components for secure operation. This method introduces the least amount of change to an existing hardware infrastructure at the lowest possible performance level. For underpowered systems, encryption and compression routines may be too resource intensive and require substantial hardware upgrades to bring such systems up to compliance standards.
- Storage-based encryption solutions utilize native processing resources to apply encryption and compression routines. As such, this solution leverages much of the operational overhead away from hosts and host applications to achieve the same or similar goals. This method introduces a moderate amount of change to existing hardware infrastructures and may require additions to or replacements of client and server software.
- Appliance-based encryption solutions are implemented in specialized standalone units utilizing hardware-accelerated encryption and sometimes compression routines. This method also introduces marginal change to an existing network infrastructure and minimal disruption to services. Some units are capable of interfacing different backend storage technologies including NAS, FC, and IP SAN. As you would expect from the best-performing solution, it is also the most expensive to implement.

Budget and performance costs are largely determined by the implementation approach taken for each solution. Host or application-based solutions may include per-seat or per-user licensing costs and additional periodic upgrade fees, where turnkey storage and appliance solutions involve upfront product costs and may also entail subsequent support and service subscription fees. Hardware-accelerated encryption solutions are considerably more expensive than software applications that use general purpose computing resources. Furthermore, each such solution brings with it a unique set of challenges and potential changes to the existing hardware or network infrastructures, all of which must be factored into the selection process.

Storage and appliance solutions continue to grow in popularity as an alternative to resource-consuming host or application implementations. Some offerings are built around gigabit-capable encryption appliances that integrate transparently into NAS, SAN, DAS, and tape backup environments and provide cryptographically secured storage compartmentalization. Other similar solutions use enterprise-class network appliances that non-intrusively safeguard critical database information and that deliver enterprise-class transparent cryptographic transmission and storage. Each of these devices can pass data compressed and encrypted at line speeds for superior performance with minimal impact to other networking resources and tasks. There are also specialized parts such as encrypted tape storage solutions and encrypted notebook drives that provide coverage for data at rest that may venture off-site beyond perimeter security controls.

### *Identifying and Minimizing Potential Bottlenecks*

Encryption merely complements or expands an existing IT security posture and should be applied anywhere sensitive data is stored, collected, and maintained. This protection comes at a considerable cost: good software solutions can impose additional overhead in terms of resource consumption and levy performance penalties that range from mild (barely noticeable) to 15 to 20 percent reductions in processing and throughput, and effective hardware solutions to accelerate good software solutions are costly to purchase. That said, not all SSH implementations impose noticeable performance reductions, and most enterprises and organizations are happy to accept minor performance degradation in exchange for improved security and the ability to hit mandated compliance targets.

Performance suffers wherever compression and encryption tasks double-up alongside other process-intensive tasks. Workstations with general purpose CPUs that process application data with these other two resource-consuming tasks can become overburdened. Network transactions also lean on the CPU unless some kind of onboard hardware acceleration is built into the network interface, which is more the exception than the rule. But indeed, both of these chores can be offloaded to add-in PCI cards for those willing to purchase and deploy them. Such specially designed cards can deliver moderate-cost hardware compression and encryption solutions with throughput suitable for use on T3, E3, and Fast Ethernet technologies.

When you examine your situation and start working toward a solution, it’s important to assess what data realistically needs cryptographic coverage. You should perform a risk analysis and restrict the scope of cryptographic coverage to only the most at-risk business assets. Not only is this a good practice, it is also required by numerous regulations. Not every file and network transaction must be cryptographically secure; neither must every network endpoint. Establishing an appropriate level of precaution against illegal tampering or interception of safely stored data reduces liability, remains within compliance requirements, and reduces the waste and potential overuse of resource-intensive encryption and compression products.

### *Benefits of Hardware Acceleration*

Storage security appliances can be placed inline between storage volumes and client endpoints to reduce processing power required for local host-based resources. Hardware acceleration and integrated encryption components built-in to these appliances facilitate the compression and encryption routines transparently for network-driven transactions. They can also provide excellent performance returns for the investment involved. In many cases, such appliances introduce little or no changes to existing network infrastructures, provide file access monitoring, and might possibly even restrict unauthorized use. All of these factors add up to sizable potential gains. Storage appliances, network appliances, and add-in peripherals all tend to incorporate some form of hardware acceleration to drive compression and encryption routines while also delivering reasonable levels of performance to enterprise users.

## Making a Technology Selection: Security

Trust is an important and considerable factor in any cryptographic exchange. We trust that the cryptographic algorithms used in an exchange are reasonably strong, that we have chosen appropriately strong passwords, that the implementation is relatively invulnerable to most forms of attack, and that both parties involved can be identified, authenticated, and accounted for. Sometimes a trust relationship is formed in a small window of time between the originator and recipient of a message, as happens with online retail outlets where a customer renders electronic payments. It's next to impossible to verify anyone's identity one-on-one through an Internet connection alone. Therefore third-party entities are entrusted to bind and manage identities mapped to credentials so that authorized parties can establish confidentiality with one another and begin transmissions in a cryptographically secure manner.

### *Key Management Schemes*

Standards and protocols such as IPSec and TLS/SSL use server-level Public Key Infrastructure (PKI), passwords, and shared secrets to provide a secure foundation for network communications. By design and by default, these solutions do not always match up to the comprehensive security, scalability, or manageability required by compliance regimes and provisions offered by true end-to-end PKI solutions.

To ensure that these point-to-point transactions across diverse platforms and applications are secure, users can use digitally signed transactions to provide the paper trail necessary to meet compliance standards. Provisions for public key encryption and digital signature services and manageability are made possible through the PKI. An organization establishes and maintains trustworthy environment network partnerships through the PKI in a transparent manner—that is, an end user is not required to understand how PKI works to take advantage of its services.

An alternative approach to public keyed information exchange is a web of trust scheme, a method of using self-signed certificates and third-party verifications for such certificates. Pretty Good Privacy (PGP), GNU Privacy Guard (GnuPG), and OpenPGP are increasingly popular as a result of their inclusion in secure email and Web transactions, among other practical uses. Simple PKI (SPKI) is born of three independent efforts to reduce the complexity of X.509 certificates and avoid the anarchy of PGP's web of trust. SPKI binds people and systems to their respective keys using a local trust model like the web of trust with integrated authorization capabilities but has no definition for the role of Certificate Authority (CA).

## Public Key vs. Private Key Infrastructures

PKI encompasses a set of developed and formalized standards, products, solutions, policies, and practices for secure data transmissions over unsecured lines. As such, PKI is an arrangement that provides for trusted third-party examination, evaluation, and verification of user identities. This method of secure transaction binds public keys to user identities and exchanges certificates on behalf of properly authenticated and presumably legitimate parties. PKI is a stable cryptographic framework that sees widespread usage and deployment across the often unruly Internet and appears in several types of products from email and HTTP clients to VPN concentrators and wireless infrastructures.

To examine basic PKI features reveals many of the reasons that private key technologies do not work in the context of mutual public exchanges. For starters, how does one safely exchange private keys, ensuring that no party has intercepted or modified such keys? With a single all-purpose private key, anyone can use it to encrypt and decrypt messages keyed with its content. Therefore, how can you trust that all subsequent communications are made only by the intended recipient of the private key? These and many other questions are resolved by using some form of PKI, where each implementation has its own solutions for each problem associated with public key exchanges.

There are many ways to safely and securely move data; exchanging private keys is not one of them. Where a cryptosystem uses the Digital Encryption Standard (DES) algorithm to create all-purpose private keys, RSA uses an asymmetric algorithm to achieve the public/private key pairs necessary for PKI. Digital signatures and certificate systems bind the original party with the public encryption key together with some key distribution system with management facilities such as a CA or web of trust.

## Benefits of Centralized Key Management

Key backup and key escrow are two features provided through centralized key management facilities. In key escrow systems, a third party can obtain decryption keys for encrypted data. Such need arises when a federal investigation requires access to information secured by encryption keys kept by an escrow service, or when an authorized user accidentally loses access to or destroys a key that must be replaced so that data encrypted with that key remains accessible.

Key backup solutions are ideal in less legally binding cases, such as when employees change departments or employers and for long-term data storage. Passwords and logon credentials change over time and may easily be lost or forgotten, unless little variation is introduced into your routine. Your password or passphrase of 10 years ago surely isn’t the same one you use today. But your public and private keys are another matter altogether. Thus, key backup serves an important recovery purpose—to keep secured data accessible to authorized parties when the original party or original key is no longer available.

## PKI and Interoperability Issues

Interoperability has proven to be a crucial issue for PKI implementations. It’s simply not safe to assume that PKI components from vendor A will always work with those from vendor B, or that components from one or both A and B will also work with vendor C, and so forth. Unfortunately, this puts IT professionals in the interesting situation of having to seek definite proof of interoperability as part of the purchase process, or to stick to products from a single vendor. The only other sensible multi-vendor alternative is to purchase small numbers of products for test or pilot implementations, to weed out products or solutions that don’t prove to interoperate properly, and to construct workable systems through repeated testing that may not exactly qualify as “trial and error” but often looks like something close to that by the time all the necessary work is done.

Another approach is to look for products built around common underlying PKI technology, which may not require a single-vendor architecture but ends up employing a single-supplier PKI nevertheless. This explains why some security companies make a business out of licensing certificate toolsets specifically to support the Internet X.509 PKI certificate and certificate revocation list as specified in RFC 3280, aka PKIX for Public Key Infrastructure (X.509). Although it may not always be possible to purchase PKI software or components from the same vendor, to meet client needs for email, VPN, file transfer, and so forth as well as to support PKI in an enterprise (CAs, revocation list management, Online Certificate Status Protocol services, and so forth), make sure that at least some of the different components employ common underlying PKI technology. Doing so simplifies the job of vetting interoperability and may even make it possible to demonstrate interoperability without extensive testing. The most responsible vendors maintain extensive libraries of compatibility or interoperability test documents that are well worth researching as you work your way into a practical and manageable security implementation, especially where PKI is involved.

## Avoid Security Through Obscurity

Simply put, security through obscurity is wishful thinking and doesn’t work. Here, obscurity means the act of hiding implementation-specific details governing the operation of a presumably secure protocol, application, or infrastructure. This might be a software product, a communications protocol, or a topological arrangement of network devices. By leaving crucial details invisible to observers, the idea is that little can be derived, discovered, or determined about whatever security measures are in place. The assumption is that nothing shows to reveal its secrets or to exploit its weaknesses. In actuality, this couldn’t be further from the truth.

Secrecy as a method of security is flawed for many reasons. Though a system relying on such logic may have theoretical or actual vulnerabilities, its designer makes the assumption that its weak points are unlikely to be identified by attackers because no properties about such conditions have been revealed. However, this does not mean vulnerabilities cannot be found, and certainly does not mean such a system can safely be entrusted with protected information or allowed to operate independently within a secure setting.

There is nothing wrong with public or private scrutiny of a security or cryptographic system's source code. In fact, there are benefits to having professional developers, security experts, and systems analysts review a given application, protocol, or framework for weak points that could lead to vulnerability or exploitation. A statement to this effect, known as Kerckhoff's principle, was made in 1883 by Auguste Kerckhoff, "A cryptosystem should be secure even if everything about the system, except the key, is public knowledge."

Indeed, a cryptosystem is much like any other security system and should always be reviewed by many specialists and experts in various security fields to ensure proper fitness during operation. World-renowned security expert and cryptography specialist Bruce Schneier extends this reasoning to all security systems when he states the following, "Kerckhoff's principle applies beyond codes and ciphers to security systems in general: every secret creates a potential failure point. Secrecy, in other words, is a prime cause of brittleness—and therefore something likely to make a system prone to catastrophic collapse. Conversely, openness provides ductility." This notion of ductility, or structural rigidity and resistance to deformation, is vital to every security framework and infrastructure.

Therefore, whenever you select technology for enterprise deployment, insist on security-evaluated or certified software products. When researching a potential security product for use within a production environment, inquire about its internal security polices and practices and whether the system has been audited by any external agency or testing laboratory. In lieu of such assurances, you might have an independent auditing firm inspect or test the product for potential problems. Accept no vendor assurances to the contrary: such systems must be verified and validated. Any sign of resistance, short of non-disclosure agreements to third-party sources, should be viewed with suspicion. After all, the safety and livelihood of your confidential business practices are at stake and there is no sense in leaving such critical items of interest unchecked or untested.

## Compliance and Security Audits

In terms of networking, compliance refers to the state or act of conforming with or agreeing to provide adequate security and privacy according to mandated and regulated specifications. For IT departments, the following three areas are most affected when processing confidential data:

- Data authorization and access
- Data retention
- Data retrieval

Network and host-based storage infrastructures are susceptible to damage or loss of data, potentially vulnerable to security breaches, and challenged by the dynamic nature of an ever-changing security landscape. To further complicate matters, organizations that deal in PII are perpetually regulated and routinely examined for compliance. IT's major role and responsibility in this process is to maintain a strong security posture that satisfies both sides of the at-risk equation. Execution requires a methodical approach involving risk assessment, implementation evaluation, and periodic monitoring practices. It also strongly urges the need for enhanced security awareness, training, and expertise, and demands the following IT issues:

- Uncovering existing problems
- Mapping potential solutions to business values
- Justifying ROI
- Targeting high-return countermeasures
- Rigorously using management-based measurement and monitoring methods

With SOX, GLBA, and HIPAA now part of the common business vernacular, a proliferation of many more mandatory regulations and regulatory regimes give rise to innumerable IT security requirements, some of which are vague and difficult to translate into technical controls. Regulatory compliance requirements tend to be broadly worded and use generalized concepts, while asset compliance policies are precise and accurately descriptive. Where a regulation is worded too broadly, there is the potential for misinterpretation of the definition of compliance depending on the subject making the interpretation. The evaluation of a solution that correctly addresses all such broadly defined mandates may not even improve the existing security stance. This is what makes internal and external audits so important to ensuring compliance—both give ample opportunities to identify ambiguities, resolve misinterpretations, and to ensure that implementations adhere to both the letter and the spirit of governing regulations. This also explains why compliance experts stress the need to treat compliance needs as an opportunity to improve overall IT operations, not just to bring security up to minimally acceptable levels.

Securing IT assets requires a specialist, someone who possesses a specialized knowledge on a variety of security subjects and is willing and able to research and learn about new technologies (and presumably, related regulatory requirements that may trail in their wakes). Navigating such a labyrinthine maze of options is a daunting task even for a well-schooled and seasoned Chief Security Officer (CSO), Chief Technology Officer (CTO), or IT administrator. IT security involves complex subject matter, including numerous specific network protocols and convoluted security regimes, plus access rights and permissions for systems across multiple domains and geographically dispersed organizational assets.

HIPAA imposes national standards for securing and maintaining privacy and security for protected health information (which includes medical data and other personal information such as name, address, and so on). The Privacy Rule applies to protected health information in all forms; the Security Rule applies to electronic data. This makes it imperative that all types of information, electronic or not, are safeguarded appropriately. The HIPAA standard encompasses three defined types of organizations: health care providers (for example, clinics, hospitals, and pharmacies), health care insurers, and health care clearinghouses.

GLBA, or the Financial Services Modernization Act of 1999, applies to financial institutions and organizations that deal in third-party monies such as banks, brokerage firms, and consumer credit reporting agencies. Credit counseling services, debt collection agencies, real estate services, and income tax preparers are also regulated by GLBA. GLBA comprises three portions, only one of which applies to IT professionals—the Safeguards Rule. Section 501 of GLBA addresses "Protection of Nonpublic Personal Information," which requires that financial institutions adhere to administrative, technical, and physical safeguards for handling customer information. The Safeguards Rule section governs the collection and disclosure of sensitive client financial information and requires that regulated companies specify a person or group responsible for GLBA compliance, identification of security risks, safeguard assessment, safeguard implementation, security monitoring, compliance assurances, and completion of necessary upgrades to maintain compliance.

SOX sections 302 and 404 are the portions that affect IT departments for regulated companies. All companies registered under the Security Act of 1993 must comply. Under these sections, companies are required to establish an infrastructure secure from unauthorized data access or alteration, electronic damage and loss. This also means documenting all security measures taken.

Other relevant legislation includes California Senate Bill (SB) 1386, which went into effect on July 1, 2003. This bill requires all state agencies, persons, or businesses that conduct business in the state of California that also own or license computerized data (including personal information) to disclose any breach of security of that data as mandated in various ways by the bill. What is interesting about this legislation is that it requires reporting of any and all security breaches involving access to "unencrypted personal information" by unauthorized parties. Thus, in effect, this legislation not only mandates reporting of breaches in keeping with other similar regulatory requirements at all levels of government but also explicitly recognizes that encrypted data is protected and need not be reported should it be accessed by outsiders. This provides yet another strong argument for the value of encryption for data, either when stored at rest or in transit during file transfer, transaction processing, email, and so forth.

✎ For more information about SB 1386, please visit http://info.sen.ca.gov/pub/01-02/bill/sen/sb_1351-1400/sb_1386_bill_20020926_chaptered.html.

Third-party services and vendors can be recruited where in-house technical knowledge or resources are scarce to none. Various storage vendors provide assessment services from auditing a storage environment for security weaknesses to drafting open solutions to meet business-specific needs and standards. Their data storage network experts can assist in the evaluation of services and products tailored to fit a given organizational structure and data storage requirements. Such services take a holistic life cycle approach with coverage through a four-cycle process: risk assessment, prioritization, implementation, and monitoring.

### *Stick to Security Standards*

If there is a moral that emerges from the preceding section, it provides the title for this one—namely, that working with well-documented and well-understood industry standards remains an enterprise's best hope for establishing and maintaining security. Though it's tempting to view Andrew Tannenbaum's often-cited observation, "The nice thing about standards is that there are so many to choose from," with a certain amount of irony, a strict focus on established industry security standards nevertheless offers the best (and some would argue, only) chance of implementing security tools and controls that implement security policy correctly and conform to applicable laws, regulations, and best practices. To that end, a thorough understanding of TCP/IP security protocols and services comes highly recommended and should include working knowledge of IETF and W3C standards and protocols for secure transport, secure sockets, secure encryption and authentication, PKI, distributed applications, and secure remote access. This represents a huge body of knowledge and is probably best approached by seeking out experienced, knowledgeable individuals who have worked with and around at least some of these technologies for some time. It may also be useful to think about finding people who are certified with credentials such as Certified Information Systems Security Professionals (CISSP), Certified Information Security Managers (CISM), Cisco Certified Security Professionals, or even Cisco Certified Internet Experts (CCIEs) from the security track.

## Making a Technology Selection: File Transfer Topology

The type of connections among sites is what determines network topology in an enterprise; this in turn dictates how file transfers occur. Although many organizations employ a mixture of types, common network topology models include hub-and-spoke and partial mesh organizations. A hub-and-spoke model arranges satellite sites around a hub (like a wheel, where spokes link the central hub to various locations on the "rim"), where traffic between satellites must pass through the hub to get from one to another (more commonly, large enterprises might have multiple hubs, each with its own spokes, and high-speed links between hubs, or perhaps even a hierarchy, with a central hub at headquarters connected to distribution hubs, each with its own regional spokes). A complete mesh is a topology whereby every node in a graph is connected to every other node, and reflects a situation where every site is connected to every other site in an enterprise network. This is prohibitively expensive once the number of sites goes much above half a dozen, so the reality is that sites that need links (or that are geographically closest) to one another get connected, and sites that communicate infrequently use Internet links to access otherwise inaccessible sites.

### Hub-and-Spoke Transfer Environment: Pros and Cons

Many large organizations employ a traditional hub-and-spoke model to network computer systems. In fact, basic data transfer protocols such as HTTP and FTP are inherently hub-and-spoke by nature, as a spoke needs to initiate contact and possibly handle specific rules for login credentials, file names, and so forth.

As a networking model, the hub-and-spoke methodology is characterized by hub locations at corporate and regional headquarters or a single data center, with many point-to-point spoke connections extending out to branch offices. Using telephony as an illustration, each phone line represents a spoke where each phone company point of presence represents a hub. Ease of implementation, efficient cost allocation, and simple management are primary motivations for using hub-and-spoke configurations.

With health care systems, hub-and-spoke arrangements were the only initial way to create connections, where a subscriber establishes listening technology with a provider serving as the spoke to automate dial-up technology. A large subscriber company could potentially establish thousands of connections with no real unified management for these connections, as no all-purpose standard protocol exists in the health care industry. Small subscriber spokes are forced to tailor their systems to larger companies of importance, and it’s not possible for a provider to tailor its communications scripts for all subscribers owing to a lack of uniformity. This drawback also impedes hub-to-hub and spoke-to-spoke communications, so that specially configured spokes are difficult to impossible to support using automation software.

As application needs and usage patterns transition into more collaborative knowledge-sharing trends, it becomes necessary for organizations to rethink this model. Yielding higher end-user productivity—especially through collaborative efforts and global pooling of organizational resources—will exercise the hub-and-spoke paradigm in ways it may not adequately handle. Multi-Protocol Label Switching (MPLS) and VPN technologies are two proven alternatives to the hub-and-spoke design that vastly improve upon the performance and capabilities of traditional point-to-point topologies without requiring a plethora of WAN links to interconnect multiple sites.

## Finding the Best Fit for Your Environment

Each environment introduces its own challenges to achieving regulatory compliance, but the fundamental principles remain largely the same. The following checklist briefly describes these principles:

- Authorization and access should remain consistent across the infrastructure irrespective of the implementations and frameworks used to exercise these controls.
- Data retention and retrieval procedures should also remain consistent, in that information should be restricted only to authorized parties where it gets stored and obtained.
- Accountability can be maintained any number of ways, from platform-specific software monitoring components to general-purpose traffic analysis, intermediate network appliances, or distributed measurement facilities.

### *Reworking Automation*

Mission-critical software, like mission-critical hardware, should be made to operate in a cyclic and continuous manner. There are far too many types and classifications of network-driven events that occur on a regular basis for any one person or group of persons to maintain and manage. Endpoints change or may come and go at a whim for mobile users, topologies may change due to impaired routing conditions, and threat levels are dynamic and therefore constantly change. Assessing, evaluating, and documenting these changes for compliance purposes can be exhausting as is, so automated batch processing of the most mundane and easily reproducible tasks should be a natural extension of any existing IT administration process.

### *Scripting, Scheduling, and Batch Processing*

Task queues are essential for provisioning and managing modern computer systems. Even workstation and notebook systems are task oriented nowadays and are capable of scheduling services to operate at specific times or at regular intervals. Multiple similar tasks may occur concurrently and in parallel as schedules dictate, and run scripted or unscripted, tailored to meet site-specific needs and situations. At least some of these tasks are likely to involve file transfers of some kind, for everything from new commands, software updates, antivirus or anti-spyware signature files, directory updates or replication, backups and archives, and so forth.

Just about every platform from workstation to server has integrated scheduler capabilities and facilities for batch processing tasks. These automated chores take away most of the mundane administrative tasks that would otherwise be the responsibility of IT staff and end users, which is especially ideal for large-scale networking environments. UNIX and Windows hosts in particular ship with native scripting and scheduling facilities for a variety of automation tasks and purposes. Take advantage of these properties as they apply to creating a well-maintained business computing environment.

## Summary

This chapter covers a lot of ground from scoping the network environment to enumerating organizational assets. Security is a process, not a product. There is no one-shot, cure-all, set-it-and-forget-it solution. The security and compliance process requires astute attention to detail, due diligence, and an unflinching persistence. Security is a deep and complex issue that is difficult to get right, even the second and third time around.

The following list summarizes the content covered in this chapter:

- Start by inventorying organization assets. Take stock of all server hardware, software, services, and protocols within the network environment.
- Identify types of file transfer protocols and classify data exchange-based services. Create a baseline of known-good services and protocols and take note of any network anomalies.
- Consult with compliance management officers, teams, and third-party professionals to obtain an accurate view of your business-specific needs. Ensure that the rules of compliance are clearly defined and strictly followed in practice.
- Analyze the IT infrastructure for risk and vulnerability and document every security implementation used to reduce risk factors for each case.
- Seek insecure file transfer solutions where PII is at risk of exposure to unauthorized parties. Legacy applications with no suitable drop-in replacement should be secured as well, perhaps by creating a cryptographically secure network connection between endpoints.
- Interoperability between platforms, protocols, and services is instrumental to maintaining consistency and reliability. One secure solution may be implemented differently on a Windows host versus on a UNIX host, but they should interact in an identical fashion.
- Verify and validate everything. Though a software protocol, package, or platform may advertise itself as undefeated or unbreakable, it must be validated and verified by auditors, penetration testers, or other security professionals time and time again.
- Automate everything possible. In doing so, you relieve strain from IT staff having to perform repetitive or routine tasks that are better handled by integrated system scheduling services.

## Download Additional eBooks from Realtime Nexus!

Realtime Nexus—The Digital Library provides world-class expert resources that IT professionals depend on to learn about the newest technologies. If you found this eBook to be informative, we encourage you to download more of our industry-leading technology eBooks and video guides at Realtime Nexus. Please visit http://nexus.realtimepublishers.com.