netwrix

Data Security Best Practices

Table of Contents

1. Understand data technologies and databases

Databases have become increasingly sophisticated over the last decades. The relational database is the most common. This technology allows data to be viewed in dynamic ways based on the user's or administrator's needs. The most common language used to communicate with databases is Structured Query Language (SQL). SQL allows users to pass queries to database servers in real time. This flexibility causes a major vulnerability when it isn't securely implemented. Don't confuse the language SQL with Microsoft's database product SQL Server, though, like most databases, SQL Server uses Structured Query Language.

Early database systems connected the end user directly to the data through applications. In a private network, physical security was usually all that was needed to protect the data. Now, to improve the performance and security of databases, companies chose one of these models:

- **One-tier model**. In a one-tier, or single-tier, model, the database and the application exist on a single system. This is common on desktop systems running a standalone database. Early Unix implementations also worked in this manner; each user would sign on to a terminal and run a dedicated application that accessed the data.

- **Two-tier model**. In a two-tier model, the client workstation or system runs an application that communicates with a database that is running on a different server. This is a common implementation, and it works well for many applications.

- **Three-tier model**. The three-tier model effectively isolates the end user from the database by introducing a middle-tier server. This server accepts requests from clients, evaluates them and sends them on to the database server for processing. The database server sends the data back to the middle-tier server, which then sends the data to the client system. This approach is becoming common today. The middle server can also control access to the database and provide additional security.

NoSQL

NoSQL is a relatively new concept. As noted above, most commercial relational database management systems (Oracle, Microsoft SQL Server, MySQL, PostGres, and so forth) use SQL. A NoSQL database is not a relational database and does not use SQL. These databases are less common than relational databases but often used where scaling is important. Here are some key differences:

Feature	NoSQL Database	SQL Database
Database type	Non-relational/distributed	Relational
Schema type	Dynamic	Pre-defined
Data storage	Stores everything in a single nested document, often in XMLformat (document-based)	Individual records are stored asrows in tables (table-based)
Benefits	Can handle large volumes ofstructured, semi-structured, and unstructured data	Widely supported and easy to configure for structured data
Typical scaling model	Horizontal (add more servers)	Vertical (upgrade the server)
Popular vendors/ implementations	MongoDB, CouchDB	Oracle, Microsoft, MySQL
Susceptible to SQL injection attacks?	No, but susceptible to similar injection-type attacks	Yes

Big data

Some organizations store extremely large amounts of data, often many terabytes. This “big data” normally cannot fit on a single server, so instead it is stored on a storage area network (SAN). A SAN is a separate network that is set up to appear as a server to the main organizational network. For example, multiple servers and network storage devices might be configured as a mini-network that has one purpose: to store several terabytes of data. It is connected to the main network so users can access the data in the SAN without being concerned about the complexities involved there. SANs usually have redundant servers, and they are connected via high-speed fiber optic connections or iSCSI running on copper.

One of the issues with big data is that it can reach a size where it becomes difficult to search, store, share, back up and manage it.

File systems

File systems are another way to store unstructured data and control how it is retrieved. Without a file system, information placed in a storage medium would be one large body of data with no way to tell where one piece of information stops and the next begins. Separating the data into pieces and giving each piece a name enables the information to be easily isolated and identified.

File systems can be used on many different kinds of media, such as SSDs, magnetic tapes and optical discs. File system types depend on the operating system used; for example, Linux uses file systems such as ext family, xfs and jfs; Windows OS uses fat, fat32 and ntfs; and MacOS uses apfs and hfs+.

2. Identify and classify sensitive data

To protect data effectively, you need to know exactly what types of data you have. Data discovery technology will scan your data repositories and report on the findings. Then you can organize the data into categories using a data classification process. Data discovery engine usually uses regular expressions for its searches, which is a very flexible thing but quite complicate in creation.

Using data discovery and classification technology helps you control user access to critical data and avoid storing it in unsecure locations, thus reducing the risk of improper data exposure and data loss. All critical or sensitive data should be clearly labeled with a digital signature that denotes its classification, so you can protect it in accordance with its value to the organization. Third-party tools, such as Netwrix Data Classification, can make data discovery and classification easier and more accurate.

As data is created, modified, stored or transmitted, the classification can be updated. However, controls should be in place to prevent users from falsifying the classification level. For example, only privileged users should be able to downgrade the classification of data.

Follow these guidelines to create a strong data classification policy. And don't forget to perform data discovery and classification as part of your IT risk assessment process.

3. Create a data usage policy

Of course, data classification alone is not sufficient; you need to create a policy that specifies access types, conditions for data access based on classification, who has access to data, what constitutes correct usage of data, and so on. Don't forget that all policy violations should have clear consequences.

4. Control access to sensitive data

You also need to apply appropriate access controls to your data. Access controls should restrict access to information based on the principle of least privilege: Users have to get only those privileges that are essential to perform their intended function. This helps to ensure that only appropriate personnel can access data. Access controls can be **physical**, **technical** or **administrative**, as explained below.

Administrative controls

Administrative access controls are procedures and policies that all employees must follow. A security policy can list actions that are deemed acceptable, the level of risk the company is willing to undertake, the penalties in case of a violation, etc. The policy is normally compiled by an expert who understands the business's objectives and applicable compliance regulations.

Supervisory structure is an important part of administrative controls. Almost all organizations make managers responsible for the activities of their staff; if an employee violates an administrative control, the supervisor will get held accountable as well.

Personnel education and awareness

Training should be provided to make users aware of the company's data usage policies and emphasize that the company takes security seriously and will actively enforce the policy. In addition, users should be periodically reeducated and tested to reinforce and validate their comprehension. Security measures are in place to limit what users can do but those tools aren't perfect. If users open every attachment in every e-mail, chances are high that some zero-day attack or other exploit not listed in your antivirus database will compromise a machine. Therefore, users need to be educated about their responsibilities and best practices for proper computer usage.

Employee termination procedure

Ensuring that each departing employee retains no access to your IT infrastructure is critical to protecting your systems and data. You need to work with HR to develop an effective user termination procedure that protects your organization legally and technologically from former employees. Follow these user termination best practices in order to achieve that goal.

Technical controls

In most cases, users should not be allowed to copy or store sensitive data locally. Instead, they should be forced to manipulate the data remotely. The cache of both systems, the client and server, should be thoroughly cleaned after a user logs off or a session times out, or else encrypted RAM drives should be used. Sensitive data should ideally never be stored on a portable system of any kind. All systems should require a login of some kind, and should have conditions set to lock the system if questionable usage occurs.

Permissions

User permissions should be granted in strict accordance with the principle of least privileges. Here are the basic file permissions in Microsoft operating systems:

- **Full Control** — The user can read, execute, modify and delete files; assign permissions; and take ownership.
- **Modify** — The user can read, write and delete the file.
- **Read and Execute** — The user can run the executable file.
- **Read** — The user can read but not modify the file.
- **Write** — The user can read and modify the file but not delete it.

Folders have the same permissions, plus one more — list folder contents, which enables the user to see what is in the folder but not to read the files.

Access control lists

An access control list (ACL) is a list of who can access what resource and at what level. It can be an internal part of an operating system or application. For example, a custom application might have an ACL that lists which users have what permissions in that system.

ACLs can be based on whitelists or blacklists. A whitelist is a list of items that are allowed, such as a list of web-sites that users are allowed to visit using company computers, or a list of third-party software that is authorized to be installed on company computers. Blacklists are lists of things that are prohibited, such as specific websites that employees are not permitted to visit or software that is forbidden to be installed on client computers.

In the file management process, whitelist ACLs are used more commonly, and they are configured at the file system level. For example in Microsoft Windows, you can configure NTFS permissions and create NTFS access

control lists from them. You can find more information about how to properly configure NTFS permissions in this list of NTFS permissions management best practices. Remember that access controls should be implemented in every application that has role base access control (RBAC); examples include Active Directory groups and delegation.

Security devices and methods

Certain devices and systems help you further restrict access to data. Here is the list of the most commonly implemented ones:

- **Data loss prevention (DLP)** — These systems monitor the workstations, servers and networks to make sure that sensitive data is not deleted, removed, moved or copied. They also monitor who is using and transmitting the data to spot unauthorized use.

- **Firewall** — A firewall is one of the first lines of defense in a network because it isolates one network from another. Firewalls can be standalone systems or they can be included in other infrastructure devices, such as routers or servers. You can find both hardware and software firewall solutions; some firewalls are available as appliances that serve as the primary device separating two networks. Firewalls exclude undesirable traffic from entering the organization's network, which helps prevent data leakage to third-party rogue servers by malware hackers activity. Depending on the organization's firewall policy, the firewall might completely disallow some traffic or all traffic, or it might perform a verification on some or all of the traffic.

- **Network access control (NAC)** — This involves restricting the availability of network resources to endpoint devices that comply with your security policy. Some NAC solutions can automatically fix a non-compliant node to ensure it is secure before access is allowed. NAC is most useful when the user environment is fairly static and can be rigidly controlled, such as enterprises and government agencies. It can be less practical in settings with a diverse set of users and devices that are frequently changing. NAC can restrict unauthorized devices to access your data directly from your network.

- **Proxy server** — These devices act as negotiators for requests from client software seeking resources from other servers. A client connects to the proxy server, requesting some service (for example, a website); the proxy server evaluates the request and then allows or denies it. In organizations, proxy servers are usually used for traffic filtering and performance improvement. Proxy devices can restrict access to your sensitive data from Internet.

Physical controls

Physical security is often overlooked in discussions about data security. Having a poor policy on physical security could lead to a full compromise of your data or even network. Each workstation should be locked down so that it cannot be removed from the area. Also, a lock should be placed so that the case cannot be opened up, exposing the internals of the system; otherwise, hard drives or other sensitive components that store data could be removed and compromised. It's also good practice to implement a BIOS password to prevent attackers from booting into other operating systems using removable media. Mobile devices, such as smartphones, tablets, laptops, USB flash drives, iPods, and Bluetooth devices require special attention, as explore below.

Laptop security

With laptops, the biggest concerns are loss and theft, either of which can enable malicious parties to access the data on the hard drive. Full-disk encryption should be used on every laptop within an organization. Also, using public wi-fi hotspots is never a good idea unless a secure communication channel such as a VPN or SSH is used. Account credentials can be easily hijacked through wireless attacks and can lead to compromise of an organization's network.

Mobile device security

Mobile devices can carry viruses or other malware into an organization's network and extract sensitive data from your servers. Because of these threats, mobile devices need to be controlled very strictly. Devices that are allowed to connect should be scanned for viruses, and removable devices should be encrypted. It is great to use NAC for that purpose.

It is important to focus in on the data, not the form factor of the device it resides on. Smartphones often contain sensitive information, yet less security protection is applied to them than to laptops that contain the same information. All mobile devices that can access sensitive data should require the same-length passwords and have the same access controls and protection software.

Another big data leakage instrument is a smartphone with a camera that can take high-resolution photos and videos and record good-quality sound. It is very hard to protect your documents from insiders with these mobile devices or detect a person taking a photo of a monitor or whiteboard with sensitive data, but you should have a policy that disallows camera use in the building.

Network segregation

Network segmentation involves segregating the network into logical or functional units called zones. Each zone can be assigned different data classification rules, set to an appropriate level of security and monitored accordingly. Segmentation limits the potential damage of a compromise to a single zone. Essentially, it divides one target into many, leaving attackers with two choices: treat each segment as a separate network, or compromise one and attempt to jump the divide. Neither choice is appealing. Treating each segment as a separate network creates a great deal of additional work, since the attacker must compromise each segment individually; this approach also dramatically increases the attacker's exposure to being discovered. Attempting to jump from a compromised zone to other zones is difficult because, if the segments are designed well, then the network traffic between them can be restricted. There are always exceptions that must be allowed through, such as communication with domain servers for centralized account management, but this limited traffic is easier to characterize.

Video surveillance

Monitoring all critical facilities in your company by video cameras with motion sensors and night vision is essential for spotting unauthorized people trying to steal your data via direct access to your file servers, archives or backups, as well as spotting people taking photos of sensitive data in restricted areas.

Locking and recycling

Your workspace area and any equipment should be secure before being left unattended. For example, check doors, desk drawers and windows, and don't leave papers on your desk. All hard copies of sensitive data should be locked up, and then be completely destroyed when they are no longer needed. Also, never share or duplicate access keys, ID cards, lock codes, and so on.

Before discarding or recycling a disk drive, completely erase all information from it and ensure the data is no longer recoverable. Old hard disks and other IT devices that contained critical information should be physically destroyed; assign a specific IT engineer to personally control this process.

5. Implement change management and database auditing

Another security measure is to log all database and file server activities. Login activity has to be maintained for at least one year for security audits. Any account that exceeds the maximum number of failed login attempts should automatically be reported to the information security administrator for investigation. Being able to spot changes to sensitive information and associated permissions is critical. Using historical information to understand what data is sensitive, how it is being used, who is using it, and where it is going gives you the ability to build effective and accurate policies the first time and anticipate how changes in your environment might impact security. This process can also help you identify previously unknown risks. There are third-party tools that simplify change management and auditing of user activity, such as Netwrix Auditor.

6. Use data encryption

All critical business data should be encrypted while at rest or in transit, whether via portable devices or over the network. Portable systems should use encrypted disk solutions if they will hold important data of any kind.

For desktop systems that store critical or proprietary information, encrypting the hard drives will help avoid the loss of critical information even if there is a breach and computers or hard drives are missing. For example, the most basic way to encrypt data on your Windows systems is Encrypting File System (EFS) technology. If you use EFS to protect data, unauthorized users cannot view a file's content even if they have full access to the device. When an authorized user opens an encrypted file, EFS decrypts the file in the background and provides an unencrypted copy to the application. Authorized users can view or modify the file, and EFS saves changes transparently as encrypted data. If unauthorized users try to do the same, they receive an "Access denied" error.

Another encryption technology from Microsoft is BitLocker. BitLocker complements EFS by providing an additional layer of protection for data stored on Windows devices. BitLocker protects devices that are lost or stolen against data theft or exposure, and it offers secure data disposal when you decommission a device.

Hardware-based encryption

In addition to software-based encryption, hardware-based encryption can be applied. Within the advanced configuration settings on some BIOS configuration menus, you can choose to enable or disable a Trusted Platform Module (TPM). A TPM is a chip that can store cryptographic keys, passwords or certificates. A TPM can be used to assist with hash key generation and to protect smartphones and devices other than PCs as well. It can also be used to generate values used with whole disk encryption, such as BitLocker described above. A TPM chip may be installed on the motherboard.

7. Back up your data

Critical business assets should be duplicated to provide redundancy and serve as backups. At the most basic level, fault tolerance for a server means a data backup. Backups are simply the periodic archiving of the data so that if there is a server failure you can retrieve the data. From a security point of view, there are three primary backup types with which we are concerned:

- **Full** — All data is archived.
- **Differential** — All changes since the last full backup are archived.
- **Incremental** — All changes since the last backup of any type are archived.

Full backup

Full backup is the best backup strategy but it has drawbacks. Consider a scenario where you do a full backup at 1 a.m. each morning. You are concerned about the possibility of a server crash before the next full backup, so you want to do a backup every two hours as well. Which type of backup should you choose? Let's consider each option and what would happen if the system crashes at 5:10 a.m. If you do a full backup every two hours beginning at 1 a.m., then when the system crashes at 5:10 a.m., you simply need to restore the 5:00 a.m. full backup. However, running a full backup every two hours is very time consuming and resource intensive, and it will have a significant negative impact on server performance.

Differential backup

In this scenario, you do a full backup at 1 a.m. and then perform a differential every two hours thereafter. When the system crashes at 5:10 a.m., you have to restore the full backup from 1 a.m. and the differential backup from 5 a.m. This takes just one more step than restoring the full backup. Keep in mind, however, that the differential backups are going to get larger each time you do them and thus more time consuming and resource intensive. Although they won't have as much impact as the full backups, they will still slow down your network.

Incremental backup

In this scenario, you do a full backup at 1 a.m. and then an incremental backup every two hours. When the system crashes at 5:10 a.m., you need to restore the last full backup done at 1 a.m. and then each incremental backup since then — and they must be restored in order. This is much more complex task, but each incremental backup is small and does not take much time or resources to create.

There is no single correct choice of what backup to use. The proper choice depends on your organization's needs. Whatever backup strategy you choose, you must periodically test it. The only effective way to test your backup strategy is to restore the backup data to a test machine. One of the top best practices is to store your backups in geographically different places to prevent disasters such as acts of nature or accidents (e.g., hurricanes, fires or hard-disk failures) from destroying the business's IT core. Backups should be performed incrementally across multiple disks and servers, and on different time schedules (daily, weekly and monthly). Preferably, these incremental backups should save a base copy and each modification should reflect only the changes to the base copy, or a closely matching previous version. This allows for proper versioning and can help to serve as a form of data control.

8. Use RAID on your servers

RAID is a fundamental tool for fault tolerance that helps protect against data destruction and system downtime. RAID is a redundant array of independent disks. RAID allows your servers to have more than one hard drive so that if the main hard drive fails, the system keeps functioning. The primary RAID levels are described here:

- **RAID 0 (striped disks)** — Data is distributed across multiple disks in a way that provides improved speed (read/write performance) at any given instant but does not offer any fault tolerance. A minimum of two disks are needed.

- **RAID 1** — This RAID level introduces fault tolerance because it mirrors the contents of the disks; it is also called mirroring. For every disk you need for operations, there is an identical disk in the system. A minimum of two disks are needed and 50 percent of your total capacity is used for data and the other 50 percent for the mirror. For instance, a server with two hard drives will be able to store data equal to the size of one of the disks. With RAID 1, if the primary drive fails, the system keeps running on the backup drive. If you add another controller to the system, it is still RAID 1, but it is now called duplexing.

- **RAID 3 or 4 (striped disks with dedicated parity)** — This RAID level involves three or more disks with the data distributed across the disks. One dedicated disk is used to store parity information, so the storage capacity of the array is reduced by one disk. If a disk fails, there is only a partial loss of data. The data remaining on the other disks, along with the parity information, allows the data to be recovered.

- **RAID 5 (striped disks with distributed parity)** — This RAID level combines three or more disks in a way that protects data against the loss of any one disk. It is similar to RAID 3, but the parity is distributed across the drive array. This way, you don't allocate an entire disk for storing parity bits.

- **RAID 6 (striped disks with dual parity)** — This RAID level combines four or more disks in a way that protects data against the loss of any two disks. It accomplishes this by adding an additional parity block to RAID 5. Each of the parity blocks is distributed across the drive array so parity is not dedicated to any specific drive.

- **RAID 1+0 (or 10)** — This RAID level is a mirrored data set (RAID 1), which is then striped (RAID 0), which is the reason for the "1+0" name. Think of it as a "stripe of mirrors." A RAID 1+0 array requires a minimum of four drives: two mirrored drives to hold half of the striped data, plus another two mirrored drives for the other half of the data.

- **RAID 0+1** — This RAID level is the opposite or RAID 1+0. Here, the stripes are mirrored. A RAID 0+1 array requires a minimum of four drives: two mirrored drives to replicate the data on the RAID 0 array.

9. Use clustering and load balancing

RAID does a fantastic job of protecting data on systems (which you then protect further with regular backups), but sometimes you need to grow beyond single systems. Connecting multiple computers to work together as a single server is known as clustering. Clustered systems utilize parallel processing, which improves performance and availability, and add redundancy (but also add costs).

High availability can also be obtained through load balancing. This allows you to split the workload across multiple computers. Those computers are often servers answering HTTP requests (often called a server farm), which may or may not be in the same geographic location. If you split locations, this is called a mirror site, and the mirrored copy can add geographic redundancy (allowing requests to be answered quicker) and help prevent downtime.

10. Harden your systems

Any place where sensitive data could reside, even temporarily, should be adequately secured based on the type of information that system could potentially have access to. This would include all external systems that could get internal network access via remote connection with significant privileges, since a network is only as secure as the weakest link. However, usability must still be a consideration, and a suitable balance between functionality and security must be determined.

OS baseline

The first step to securing your systems is making sure the operating system's configuration is as secure as possible. Out of the box, most operating systems come with unneeded services running that serve only to give an attacker additional avenues of compromise. The only programs and listening services that should be enabled are those that are essential for your employees to do their jobs. If something doesn't have a business purpose, it should be disabled. It may also be beneficial to create a secure baseline image OS that is used for the typical employee. If anyone needs additional functionality, those services or programs will be enabled on a case-by-case basis. Windows and Linux operating systems will each have their unique hardening configurations.

Windows

Windows is by far the most popular operating system used by consumers and businesses alike. But because of this, it is also the most targeted operating system, with new vulnerabilities announced almost weekly. There are a number of different Windows versions used throughout different organizations, so some of the configurations mentioned here may not translate to all of them. Here are some things that should be done to enhance security:

- Disable LanMan authentication.
- Ensure that all accounts have passwords, whether the account is enabled or disabled.
- Disable or restrict permissions on network shares.
- Remove all services that are not required, especially telnet and ftp, which are clear-text protocols.
- Enable logging for important system events.

You can find more Windows hardening best practices in this Windows Server hardening checklist.

Linux

The Linux operating system has become more popular in recent years. Even though some claim that it's more secure than Windows, some things still must be done to harden it correctly:

- Disable unnecessary services and ports.
- Disable trust authentication used by the "r commands."
- Disable unnecessary setuid and setgid programs.
- Reconfigure user accounts for only the necessary users.

Web servers

Web servers are one of the favorite areas for attackers to exploit because of the reach they have. If an attacker can gain access to a popular web server and take advantage of a weakness there, they have the opportunity to reach thousands, if not hundreds of thousands, of users who access the site and their data. By targeting a web server, an attacker can affect all the connections from users' web browsers and inflict harm far beyond the one machine they compromised.

Web servers were originally simple in design and used primarily to provide HTML text and graphics content. Modern web servers allow database access, chat functionality, streaming media and many other services; this diversity enables websites to provide rich and complex capabilities to visitors. Every service and capability supported on a website is potentially a target for exploitation. Make sure that they're kept up to the most current software standards. You must also make certain that you give users to have only the permissions necessary to accomplish their tasks. If users are accessing your server via an anonymous account, then common sense dictates that you must make certain that the anonymous account has the permissions needed to view web pages and nothing more.

Two particular areas of interest with web servers are filters and controlling access to executable scripts. Filters allow you to limit the traffic that is allowed through. Limiting traffic to only that which is required for your business can help ward off attacks. A good set of filters can also be applied to your network to prevent users from accessing sites other than those that are business related. Not only does this increase productivity, but it also reduces the likelihood of users obtaining a virus from a questionable site.

Executable scripts, such as those written in PHP, Python, various flavors of Java and Common Gateway Interface (CGI) scripts, often run at elevated permission levels. Under most circumstances, this isn't a problem because the user is returned to their regular permission level at the conclusion of the execution. Problems arise, however, if the user can break out of the script while at the elevated level. From an administrator's standpoint, the best course of action is to verify that all scripts on your server have been thoroughly tested, debugged and approved for use.

Email servers

Email servers provide the communications backbone for many businesses. They typically run either as an additional service on a server or as dedicated systems. Putting an active virus scanner on email servers can reduce the number of viruses introduced into your network and prevent viruses from being spread by your email server. It is worth noting, though, that most scanners can't read Microsoft's open files; to scan Exchange mail stores, you need a specific email AV scanner, some of them even try to detect phishing, such technology is based on machine learning engine and has good perspectives combating social engineering attacks.

Email servers are being inundated by automated systems that attempt to use them to send spam. Most email servers have implemented measures to prevent this. The threats, however, are becoming increasingly more sophisticated. You may be able to reduce these attempts to access your system by entering the TCP/IP addresses in your router's ACL Deny list. Doing so will cause your router to ignore connection requests from these IP addresses, effectively improving your security. Such measures can be also done with spam filters.

FTP servers

File Transfer Protocol (FTP) servers aren't intended for high-security applications because of their inherent weaknesses. Most FTP servers allow you to create file areas on any drive on the system. You should create a separate drive or subdirectory on the system to allow file transfers. If possible, use virtual private network (VPN) or Secure Shell (SSH) connections for FTP-type activities. FTP isn't notable for security, and many FTP systems send account and password information across the network unencrypted. FTP is one of the tools frequently used to exploit systems.

From an operational security perspective, you should use separate logon accounts and passwords for FTP access. Doing so will prevent system accounts from being disclosed to unauthorized individuals. Also, make sure that all files stored on an FTP server are scanned for viruses.

You should always disable the anonymous user account. To make FTP use easier, most servers default to allowing anonymous access. However, from a security perspective, the last thing you want is to allow anonymous users to copy files to and from your servers. Disabling anonymous access requires the user to be a known, authenticated user in order to access the FTP server.

The best way to secure FTP is to replace it altogether. The same functionality can be found in more secure services such as Secure File Transfer Protocol (SFTP).

11. Apply a proper patch management strategy

Ensuring that all versions of the applications that reside on your IT environment are up to date is not an easy task but it's essential for data protection. One of the best ways to ensure security is to make the signatures for antiviruses and patch updates for systems automatic. For critical infrastructure, patches need to be thoroughly tested to ensure that no functionality is affected and no vulnerabilities are introduced into the system. You need to have patching strategy for both your operating systems and your applications.

Operating system patch management

There are three types of operating system patches, each with a different level of urgency.

- **Hotfix** — A hotfix is an immediate and urgent patch. In general, these represent serious security issues and are not optional; they must be applied to the system.
- **Patch** — A patch provides some additional functionality or a non-urgent fix. These are sometimes optional.
- **Service pack** — A service pack is the set of hotfixes and patches to date. These should always be applied, but test them first to be sure that no problems are caused by the update.

Application patch management

Just as you need to keep operating system patches current because they often fix security problems discovered within the OS, you need to do the same with application patches. Once an exploit in an application becomes known, an attacker can take advantage of it to enter or harm a system. Most vendors post patches on a regular basis, and you should routinely scan for any available ones. A large number of attacks today are targeted at client systems for the simple reason that clients do not always manage application patching well. Establish maintenance days where you will be testing and installing patches to all your critical applications.

12. Protect your data from insider threats

Organizations continue to spend an exceptional amount of time and money to secure the network at the perimeter from external attacks; however, insider threats are becoming a key cause of data exposure. Many surveys say insider incidents account for more than 60 percent of all attacks; however, many organizations don't report insider attacks out of fear of business loss and damage to their reputation.

Insider threats come in two forms. An authorized insider threat is someone who misuse their rights and privileges, either accidentally, deliberately or his credentials were stolen. An unauthorized insider is someone who has connected to the network behind the perimeter defenses. This could be someone who plugged into a jack in the lobby or a conference room, or someone who is using an unprotected wireless network connected to the internal network. Insider attacks can lead to data loss or downtime, so it's as important to monitor activity in your network as activity at the perimeter.

Insiders using remote access

Remote access to corporate networks is also becoming commonplace. Users are working from home at an increasing rate, so it's critical to secure the connections used for remote access. Strong authentication is essential when connecting remotely. It is also important that the machines users are employing for remote access to the network are also secured properly. In addition, remote sessions should be properly logged or even video recorded.

13. Use endpoint security systems to protect your data

The endpoints of your network are under attack constantly, so having the endpoint security infrastructure in place to deal with them is crucial to preventing data breaches. Unauthorized programs and advanced malware (such as rootkits) are some of the things to consider in your endpoint security strategy. With the increased usage of mobile devices, the endpoints of the network are expanding and becoming more and more undefined. Automated tools that reside on the endpoint system are essential to mitigating the effectiveness of malware. At a minimum, you should use the following technologies:

Antivirus software

Antivirus software should be installed and kept current on all servers and workstations. In addition to active monitoring of incoming files, scans should be conducted regularly to catch any infections that have slipped through, such as ransomware.

Antispyware

Anti-spyware and anti-adware tools are designed to remove or block spyware. Spyware is computer software installed without the user's knowledge. Usually its goal is to find out more information about the user's behavior and to collect personal information. Anti-spyware tools work much like antivirus tools; many of their functions overlap. Some antispyware software is combined with antivirus packages, whereas other programs are available as standalones. Regardless of the type you use, you must regularly look for spyware, often identified by the presence of tracking cookies on hosts, and remove any that gets installed.

Pop-up blockers

Pop-ups are not just irritating; they are a security threat. Pop-ups (including pop-unders) represent unwanted programs running on the system, so they can jeopardize the system's well-being.

Host-based firewalls

Personal firewalls are software-based firewalls installed on each computer in the network. They work in much the same way as larger border firewalls — they filter out certain packets to prevent them from leaving or reaching your system. The need for personal firewalls is often questioned, especially in corporate networks that have large dedicated firewalls that keep potentially harmful traffic from reaching internal computers. However, that firewall can't do anything to prevent internal attacks, which are quite common and often very different from the ones from the internet; attacks that originate within a private network are usually carried out by viruses. So, instead of disabling personal firewalls, simply configure a standard personal firewall according to your organization's needs and export those settings to the other personal firewalls.

Host-based IDSs

Intrusion detection systems (IDSs) are also available for individual hosts. Host IDSs will monitor only the internals of a computing system. A host-based IDS will look at the system state and check whether its contents are as expected. Most host-based IDSs use integrity verification, which works on the principle that most malware will try to modify host programs or files as it spreads. Integrity verification tries to determine what system files have been unexpectedly modified. It does this with computing fingerprints, in the form of cryptographic hashes, of files that need to be monitored when the system is in a known clean state. It then scans and will issue an alert when the fingerprint of a monitored file changes. The main problem with integrity verification is that it detects the malware infection after the fact and will not prevent it.

14. Perform vulnerability assessments and cybersecurity penetration tests

Vulnerability assessments usually consist of port scanners and vulnerability scanning tools such as nmap, OpenVas and Nessus. These tools scan the environment from an external machine, looking for open ports and the version numbers of those services. The results from the test can be cross-referenced with known services and patch levels that are supposed to be on the endpoint systems, allowing the administrator to make sure that the systems are adhering to the endpoint security policies.

Penetration testing is the practice of testing a computer system, network or web application to find security vulnerabilities that an attacker could exploit. Penetration testing can be automated with software applications or performed manually. The main objective of penetration testing is to identify security weaknesses. Penetration testing can also be used to test an organization's security policy, its adherence to compliance requirements, its employees' security awareness, and the organization's ability to provide security incident response and identification. Organizations should perform pen testing regularly — ideally, once a year — to ensure more consistent network security and IT management. Here are several of the main pen test strategies used by security professionals:

- **Targeted testing** is performed by the organization's IT team and the penetration testing team working together. It's sometimes referred to as a "lights turned on" approach because everyone can see the test being carried out.

- **External testing** targets a company's externally visible servers or devices, including domain servers, email servers, web servers or firewalls. The objective is to find out if an outside attacker can get in and how far they can go once they've gained access.

- **Internal testing** performs an inside attack behind the firewall by an authorized user with standard access privileges. This kind of test is useful for estimating how much damage a regular employee could cause.

- **Blind testing** simulates the actions and procedures of a real attacker by severely limiting the information given to the person or team performing the test. Typically, the pen testers are given only the name of the company.

- **Double-blind** testing takes the blind test and carries it a step further — only one or two people within the organization might be aware a test is being conducted.

- **Black box** testing is basically the same as blind testing, but the tester receives no information before the test takes place. Rather, the pen testers must find their own way into the system.

- **White box** (Crystal box) testing provides the penetration testers with information about the target network before they start their work. This information can include IP addresses, network infrastructure schematics, the protocols being used and so on.

Conclusion

As you have seen, data protection encompasses a lot of topics and areas. It's critical for good network administrators and security professionals to keep all their security tools up to date and to use good policy management. With so many policies to enforce and applications to keep up to date, this can seem like a daunting challenge for any security team.

Another challenge with data protection is minimizing the impact on the end user. Unfortunately, programs like antivirus software, personal firewalls and threat detection systems tend to sap bandwidth and processing power from important end-user functionality. For this reason, when deciding what programs to use to protect end users, look carefully at how big a footprint the program uses and its memory utilization.

About Netwrix

Netwrix is a software company that enables information security and governance professionals to reclaim control over sensitive, regulated and business-critical data, regardless of where it resides. Over 10,000 organizations worldwide rely on Netwrix solutions to secure sensitive data, realize the full business value of enterprise content, pass compliance audits with less effort and expense, and increase the productivity of IT teams and knowledge workers.

Founded in 2006, Netwrix has earned more than 150 industry awards and been named to both the Inc. 5000 and Deloitte Technology Fast 500 lists of the fastest growing companies in the U.S.
For more information, visit www.netwrix.com.

Corporate Headquarters:
300 Spectrum Center Drive, Suite 200, Irvine, CA 92618
Phone: 1-949-407-5125 **Toll-free:** 888-638-9749 **EMEA:** +44 (0) 203-588-3023
netwrix.com/social

Powerful Data Security Made Easy

- Understand which data needs protection and how exposed it is.
- Minimize the risk of a data breach.
- Promptly detect data security threats.
- Make more informed incident response decisions.
- Facilitate the recovery of key data and learn from past incidents.
- Achieve and prove regulatory compliance.

Download Free 20-Day Trial