Detecting and Mitigating Security Risks

Executive Summary

With the number of security threats increasing daily, enterprises are looking for ways to identify and eradicate vulnerabilities. Because many applications are now dependent on other networked applications, downtime can ripple across the organization and impact a greater number of services. What can enterprises do to detect and mitigate these risks?

at&t

Introduction

IT security threats are multiplying quickly. One reason is that sophisticated hacking tools have grown abundant, greatly simplifying the effort required for intruders to compromise a network. Second, while hacking once largely constituted an intellectual exercise for security enthusiasts, hacking for profit has now become a business. Entire underground communities are springing up that focus on making money by invading corporate and government network resources.

Finally, there's the changing nature of the software that is being installed throughout organizations. Applications now often share components and data with one another, using the enterprise network as their primary interface. There is appreciable business value in having networked applications share data enterprise-wide. When its application foundation operates holistically and correlates events, a company can become responsive to changing business conditions, such as when inventory runs low, when credit card fraud is in progress or when a machine is about to overheat.

With applications now depending to some degree on other networked applications, a level of complexity enters the picture that is accompanied by new security vulnerabilities. Bringing down one application, a costly event in itself, no longer has an isolated effect on just that application. Rather, given that application's dependencies on other software, service downtime could ripple across the enterprise and impact any number of services. This situation has heightened the potential for compromised assets, company reputation and revenues.

Detecting Vulnerabilities

What can enterprises do to detect and mitigate these vulnerabilities? There are multiple types of threats facing a given organization. Hackers might attempt to access private information resources from inside or outside the organization, for example. Meanwhile, the public Internet carries viruses, spyware and other types of malware that get introduced to unsuspecting users during everyday communications activities, such as opening an email attachment or downloading a file. Left untreated, malware usually causes disruption or complete denial of service (DoS) to one or more networked application services.

To launch such attacks, intruders can tamper with packets in various ways. Figure 1 shows several types of attacks that overburden network and computing resources to cause situations where users are denied access to servers, applications or networks.

There is virtually no limit to the resources that can be attacked using these common methods. Targeting an enterprise's Domain Name System (DNS) servers, which translate domain names to the corresponding underlying IP addresses that Internet routers understand, could topple an entire enterprise's operations. Hackers could flood a Web site's access network with traffic or a Web server with messages and transactions so that legitimate users would not be able to access the site. Exploiting a hole in an application that is a component of an organization's total operational and access system would have a similar effect.

So how does the enterprise get its arms around these potential exploits? Generally speaking, any noticeable activity change somewhere in the network reflects an anomaly that should be investigated.

Spotting these fluctuations in activity requires scanning traffic at all key network junctures. These are the places where one network or network segment meets another, such as between enterprise access and distribution switches, between data center switches and WAN access routers and between data center switches and servers. Scanning should take place between network segments in both the WAN and LAN and might use a combination of third-party security-monitoring services and self-managed security products, which will be explored further in the section, "Mitigating Threats."

The network conditions and threats described below are among some of the red flags indicating that a security violation might be pending:

- Unusually high levels of activity on a given application port. These indicate a potential distributed DoS (DDoS) attack aimed at a particular application, which could quickly turn into an Internet-wide issue. Viruses and worms use certain application ports to propagate themselves. An appreciable spike in activity on, say, TCP port 25 might indicate an attack on email applications that use the Simple Mail Transfer Protocol (SMTP), which is universally assigned to port 25. Such an attack could potentially impact anyone connected to the Internet if the virus or worm should propagate.
- Elevated traffic volumes on ports and protocols aimed at a particular IP destination address. This condition usually indicates that the attack is intended for a specific company, perhaps targeted by a competitor or disgruntled ex-employee, or for a specific individual.
- Elevated traffic volumes coming from a particular IP source address. Here, it is likely that an individual person or company is attempting to flood a resource with bogus messages or transactions.
- The use of "bot armies," sometimes called "botnets." These large groups of remotely controlled software are sometimes malicious in nature. Often conforming to industry-standard chat protocols and controlled by a chat server, they spread spam, launch DDoS attacks against Web sites, aid in conducting fraudulent activities and prevent authorized traffic from traversing the network. Because the perpetrators tend to continually change the IP address of the device(s) generating the bot attacks, it can be challenging to block traffic from the source server to curb the attacks.

Mitigating Threats

There are several layers of threat identification and management functions that are required to deal with the broad spectrum of attack types. Most organizations use a combination of outsourced security services and self-managed security products to help discover and mitigate threats.

Network security starts with authenticating users using firewall and related access-control capabilities that are designed to enforce customer's policies about which users are allowed to access which services. Firewalls, however, don't address "bad" traffic, such as worms and the other malware, which someone might wish to inject into the Internet at large or into a specific intranet to bring down server resources, applications and networks. Infected traffic can also find its way into an enterprise network by mistake. A corporate user who picks up bad code while surfing the public Internet could innocently introduce it onto the enterprise network.

The latest generation of detection services provides organizations with the ability to give permission to security detection service providers to inspect traffic for threats and attacks directed specifically at their enterprise networks. The provider can then supply actionable information to organizations based on their internal network's traffic flow.

There are managed security services that, operating at an Internet-wide macro level, help detect such anomalies and help determine whether they are precursors to worms or viruses. These services benefit the Internet community in general as they are designed to detect potential, pending attacks before they cause serious downtime. Providers of these services usually post alerts and direct their customers to software patches that clean the malware as soon as the patches become available.

Basic DDoS Attack Types

Attack Type	Description
Spoofed	Packet in which the source address has been forged
Malformed	IPackets sent with abnormal bits or flags
Floods	High rates of legitimately formed packets
Null	Null is a packet that has no payload. The protocol in the packet header could be TCP, UDP, etc.
Protocol	Packets sent with illegitimate protocol, a value that is not legal in the protocol field of the packet header

More importantly, the latest generation of detection services provides organizations with the ability to give permission to security detection service providers to inspect traffic for threats and attacks directed specifically at their enterprise networks. The provider can then supply actionable information to organizations based on their internal network's traffic flow. Enterprises using these services often can access a portal service where enterprise IT staff can view their own network traffic patterns and event information.

These portal-based services might spot anomalies in the Internet backbone destined for a given enterprise's network and deal with them before they even reach the premises. Once it discovers an aberration, the provider may divert traffic to a specialized device, either to quarantine the traffic stream or to "scrub" it. Scrubbing involves applying a series of algorithms to remove the packets causing DDoS, whether it is coming from a single IP address or botnet, allowing the device to then inject the valid traffic back into the routing path toward its destination on the enterprise network.

The other primary components of the security detection foundation include the following:

Firewall Filtering

To monitor and control access outside of the enterprise network, firewall filtering can be deployed in the form of a network-based service or using self-managed products installed at the edge of the enterprise WAN. As a first line of defense, firewalls make a "permit/deny" decision about whether to grant a particular user access to the enterprise network based on access-control lists (ACLs) created by the IT department.

For large installations, network-based firewall services scale much better for protecting against external attacks, because a "perimeter" firewall is not required behind the WAN access router in each and every enterprise location. In CPE-centric environments, capital and operational expenses inflate quickly as enterprises add increasing numbers of distributed sites. Scanning for attacks that originate inside the organization does require firewall-filtering capabilities directly on the enterprise premises between access and distribution switches and between core switches and servers. The internal firewalls can be managed either by the security services provider or by the corporate IT staff.

Firewalls have been enhanced in recent years with extra preventive features called application inspection, or deep packet inspection (DPI). These capabilities allow firewalls to examine, help identify and verify application types and treat traffic according to detailed policies that go beyond just network information at Layer 3. DPI allows networks to block, for example, traffic and users that unlawfully try to gain admittance using an open TCP application port.

Intrusion Detection/Prevention Systems (IDS/IPSs)

Together, these functions continually check traffic flows at the higher layers (4 – 7) in search of known malicious signatures or unusual conditions, such as those described in Figure 1. The "detection" component compares signatures to a known database of worms and other malware, and the "prevention" component filters any traffic that matches it off the network. Many of today's IDS/IPS systems and services can also address so-called Day Zero attacks by aiming to identify generally anomalous traffic and protocol behavior and then rate-limiting or blocking that traffic from the network. Day Zero attacks are suspicious signatures that have not yet been identified and stored for comparison.

Note that in the wireless LAN environment, there are wireless IDS/IPSs that scan at the lower, radio-frequency (RF) layers to ferret out identified unauthorized devices operating in an organization's air space. Wireless IDS/IPSs do not deeply inspect application-level traffic; they work in conjunction with an enterprise's wired IDS/IPS systems and services for that function. Rather, their purpose is to scan all Wi-Fi channels to help detect, alert and possibly act on the presence of unauthorized devices that are not known by the enterprise, but are connected to its network.

Endpoint Security and Antivirus Control

This security component involves running special client software on mobile and remote user devices. Network access control/protection software or appliances in the managed service provider's security operations center (SOC) or in the enterprise's data center check the client devices for viruses and helps ensure that client software is in compliance with the organization's current software versions and standards. The scans compare the operating system, application and antivirus software versions residing on the devices with the corporate policy. If there is a match, the connection is allowed. If not, the system takes the action as dictated by policy to block the connection, update the software or quarantine the connection for later remediation.

Security Information and Event Monitoring (SIEM)

SIEM tools and services enable enterprise-wide event logging, correlation to other events, incident management and reporting. Based on the resulting comprehensive network security picture, SIEM tools provide snapshots, trends and related incidents that help identify the number of security events that should be viewed and addressed by IT staffs. This helps IT staffs, which are facing increased workloads, prioritize security events, helping to decrease the number of events that need to be manually addressed. Automated capabilities that review and correlate hundreds or thousands of daily events leave staff able to handle a manageable number of events in a given day. According to the SANS Institute, a worldwide provider of information system security training and certification, about 1,000 events per day is a practical maximum of events to handle.[1] But some businesses without SIEM systems are seeing 100,000 or more events per day. SIEM tools can also be used in the context of a centralized security service from an ISP.

In the future, protection of data and information at the individual level will become of paramount importance. Security policies and monitors will be able to identify, down to the file level, those individuals who have read a document, and they will control who has access to a given document using versions of identity management that utilize biometrics, smart cards and other technologies.

As a service or on-site tool, SIEM plays a large role in detecting, alerting and remediating vulnerabilities. SIEM generally involves using automated security tools and services that integrate with other security devices such as firewalls, IDS/IPSs and user authentication systems to correlate events enterprise-wide. SIEM tools generally

consolidate logs, gathered from various monitors, into a centralized server, for example, where AI-based software quickly sifts through the logs to identify attacks and correlate them in an enterprise-wide context. Such a system might pick up on a repeated event over a period of time, such as multiple unsuccessful log-in attempts to crack a password (called a "brute force" method). If the system identifies a number of unsuccessful log-in tries, followed by a successful log-in and a network configuration change, this is the type of activity IT security staff would likely wish to be alerted about right away with an automated page or an email notification.

Once a threat is identified, SIEM systems and services also enable the automation of managing an incident, whether that entails an email alert, an automated remediation action or the creation of a trouble ticket. Security event reporting is also part of this discipline.

Today's business environment requires audits, logging and the tracking of network resource access. Tracking is performed through the use of Internet cookies, messages from a Web server that a Web browser stores in a text file. The message is sent back to the server each time the browser requests a page from the server to identify users and potentially prepare customized Web pages for them, based on information they have input in the past.

In the future, however, protection of data and information at the individual level will become of paramount importance. Security policies and monitors will be able to identify, down to the file level, those individuals who have read a document, and they will control who has access to a given document using versions of identity management that utilize biometrics, smart cards and other technologies. What will be required for this model is a universally accepted and trusted source for identity management, similar in concept to the public key infrastructure that binds public keys with respective user identities by means of a trusted certificate authority.

Centralizing network policies and security also helps overcome software-patching issues, which have the potential to cause significant vulnerabilities if a foolproof patching process is not in place. By pushing software updates out to predetermined network devices all at once from a central location, organizations keep patches updated and synchronized throughout the organization.

The Importance of Policies

Centralizing network security policies is a recommended industry best practice. To follow this model, enterprises create one central place for setting, maintaining and enforcing a common set of security policies across all network sites. The functions, services and systems described in the section above function as the "policy enforcers."

Policies are at the core of any security foundation. They can cover a lot of ground, including what action is to be taken if certain conditions are discovered. These conditions can ran ge from the discovery of malware to the type of connection being used to access the network. Enterprises may wish to bar certain legitimate traffic from the network based on circumstances surrounding the connection, such as device type, access method or even time of day.

One common enterprise policy is to prohibit the use of the older Telnet protocol to export traffic, because Telnet sends traffic in the clear rather than in an encrypted state. This leaves organizations vulnerable to so-called man-in-the-middle attacks; whereby, if a hacker grabs a piece of information about a client from an open session, that hacker could pose as a valid user and take control of the session.

Centralized policies can be enforced internally or through the use of a managed carrier service. When created and enforced internally, the policy resides in either an appliance or router in the enterprise data center, where scanning occurs. In the case of managed services, the policies can be uploaded from the data center to the service provider's SOC, where incoming requests are scanned on behalf of the business. This is designed to keep harmful traffic farther away from internal enterprise network components, lowering the risk of a breach.

Centralizing network policies and security also helps overcome software-patching issues, which have the potential to cause significant vulnerabilities if a foolproof patching process is not in place. By pushing software updates out to predetermined network devices all at once from a central location, organizations keep patches updated and synchronized throughout the organization. Without such a system, IT staffs usually follow a priority-based procedure, where certain locations take precedence over others to accommodate their time and load constraints. Depending on patching load and frequency, this system can result in some sites being left unpatched for a period of time, opening a pinhole to a potential enterprise attack.

Conclusion

The abundance of sophisticated hacking tools, network-security attacks and the interdependent nature of networked applications have created a complex environment that introduces new vulnerabilities. Mitigating risks now requires continuous scanning of internal and external traffic streams for unauthorized users and malware. Malware attacks could target a single application; however, by that application's association with other software, possibly bring down a good portion of the networked computing environment. The security foundation must comprise automated and multi-level detection and prevention functions that can alert IT staff and take appropriate policy-based action when such anomalies are discovered.

Detection services continually sift through traffic looking for nonstandard protocol use and comparing bit sequences to known threats in order to prevent malicious signatures from infecting the enterprise network. Some of these services function at a macro level to help identify general Internet threats that could affect nearly any Internet user if not tempered quickly. Others target Internet-based attacks aimed directly at a given enterprise's network. Still others might involve the managed service provider specifically monitoring a given organization's own virtual private backbone service by sampling traffic from the enterprise's WAN access routers to find infections that might have been generated from inside the corporation.

In addition to WAN-centric detection services, enterprises must institute firewall access control and detection scanning on their internal LANs to help protect against intrusions that might originate in-house, innocently or otherwise. Taking a centralized approach to enforcing policies, correlating events and automating actions helps the IT staff monitor and manage the entire security landscape so they can sleep soundly at night.

Reference

1. "A Practical Application of SIM/SEM/SIEM: Automating Threat Management," pp. 5-6. Published by the SANS Institute © May 2007; author, David Swift. Reference: http://www.sans.org/reading_room/whitepapers/logging/1781.php

For more information contact an AT&T Representative or visit www.att.com/business.