PowerTech
your security expert

WHITE PAPER

Secure Inside and Out: Maximizing Intrusion Detection on IBM i

Modern threats to information security have evolved to become more than mere annoyances. Hackers, no longer only interested in disrupting operations, are increasingly turning their attention to intellectual property and other sensitive information housed in corporate databases. And in the rapidly evolving threat ecosystem, many security incidents stem from inside an organization rather than from hackers. Organizations therefore have more incentive than ever to use intrusion detection systems (IDS) to prevent data leaks and operational disruptions.

Broadly defined, an intrusion detection system is the set of technologies or features used to identify malicious activity. Intrusion detection is a common security challenge facing organizations, as demonstrated by a recent high-profile incident in which the August 2012 attacks on South Carolina's Department of Revenue were not discovered until October. As a result of these intrusions, 3.6 million Social Security numbers were compromised, the latest in a recent trend of attacks targeting government websites. The attack also marked a broader trend: Detecting such attacks is becoming increasingly difficult.

There are two primary types of intrusion: external and internal. External threats include denial-of-service (DoS) attacks and other common hacker activity. The good news is that IBM i offers a robust set of features for detecting and mitigating external intrusions. However, a comprehensive intrusion detection strategy must also incorporate effective, comprehensive protections against internal intrusions.

Before delving too much into the specifics of external and internal intrusion types—and detection and remediation methods—it may be beneficial to analyze the areas in which IBM i excels the most.

The IDS of the i

IBM i ships with robust functionality for intrusion detection that provides excellent protection against a large number of threats. In addition to capabilities for highly secure default monitoring and prevention, the IBM i IDS can be expanded by implementing user-defined policies. This way, when an event that surpasses user-defined thresholds occurs, the monitoring system sends a notification to the built-in message queue. Optionally, these messages can be sent to the user's email. The IBM i IDS also supports a built-in audit reporting system that gives administrators visibility to network activity and breach attempts.

The PowerTech Group, Inc.
www.powertech.com • info@powertech.com
TEL USA: 253.872.7788
TOLL FREE: 800.915.7700

Users can create a set of default intrusion detection policies for the entire system, which include attack, scan, and traffic management policies. In addition to monitoring functionality, the built-in IDS interface also offers robust support for intrusion prevention with mechanisms like variant dynamic throttling, packet filtering, and Quality of Service (QoS) controls.

How the i Protects Against External Intrusions

IBM i's prepackaged features allow for a high level of layered security, which is particularly important given the evolving nature of cyber threats. As IBM's *X-Force 2012 Trend and Risk Report* indicates, traditional cybercriminal activities, such as distributed-denial-of-service (DDoS) attacks, have become more advanced. Strategies like "spoofing" allow attackers to disguise the origin of their traffic, which limits the effectiveness of packet filters and reliance on single security solutions in general.

The built-in IDS mitigates the risk of several common methods of attack:

- **Spoofing attacks:** IBM i monitors for several types of spoofing attacks. Address poisoning is a form of spoofing in which the hacker redirects traffic to a different system to steal data; or to nonexistent addresses simply to take up bandwidth and reduce overall network performance. The IDS monitors for changes in Address Resolution Protocol and the Neighbor Discovery cache to identify the type of activity normally associated with address poisoning.
- **Denial-of-service attacks:** Denial-of-service attacks have become a staple of many cyber-criminal operations. In some cases, they serve as a distraction to disguise data theft attempts. Other attackers may use this type of attack to disrupt network operations. To orchestrate a DoS attack, attackers flood a server with traffic in an attempt to slow or shut down the host. Common methods include the Fraggle attack and "the perpetual echo." In both cases, the attacker targets echo port 7 using User Datagram Protocol (UDP). Fortunately, IBM i IDS receives a notification whenever there is a UDP echo request, and if the destination is a broadcast or multicast address, IDS sends an attack notification.
- **Restricted IP protocol:** (This is an area in which configuration errors can lead to unnoticed service disruptions.) The attacker attempts to use an unrecognized IP protocol to bypass TCP/IP programming. However, if there is no policy for existing protocols, the IDS notification goes unrecorded.
- **Scan events:** Scans can be used legitimately by system administrators to monitor network security. However, an attacker may also scan the network for unused ports in an effort to identify weaknesses. IDS policies can be set to alert on both slow and fast scans. (A fast scan is generally used to gather vulnerability information or to plan a DoS attack. Slow scans indicate the attacker is looking for more detailed information, such as which operating system is running.) IDS can detect intrusions even if no policies exist, though an IDS scan policy will allow IDS to start an audit record once a scan event occurs.
- **Traffic regulation events:** Traffic regulation policies are used to detect abnormal rates of traffic like those of a DoS attack. These features can also be used to detect UDP errors, including socket errors and buffer overflow. Traffic policies may take some tweaking, as a high level of traffic sometimes indicates that a large number of legitimate users are all accessing the system at once.

Object-Level vs. Menu-Level Security

The popular misconception is that IBM is so secure that it doesn't need any help. It is true that IBM i provides excellent network security and monitoring capabilities, which enables system administrators to fine-tune their policies. However, there are still intrusion detection gaps that must be filled. According to IBM, "You may need tools to supplement the built-in IBM i security, such as restricting access during certain

time periods, or allowing users to read a particular file but not to download it. Take the time to learn what you have, and how it can be used for your organization."

The Threat of Internal Intrusions

High-profile, organized, external attacks against organizations can be particularly disruptive and tend to make headlines. However, these incidents only account for a portion of overall attacks. The Ponemon Institute's *2011 Cost of a Data Breach* study found that malicious or negligent employees are responsible for 39 percent of data breaches, suggesting that insiders pose a greater risk than many organizations account for. "As organizations of all sizes battle an uptick in both internal and external threats," the report said, "it makes sense that having the proper security leadership in place can help address these challenges."

IBM i supports crucial internal security measures. The built-in IDS supports object-level security that controls who can access a file. However, the operating system provides limited support for user activity monitoring and different levels of access. For example, an administrator may want a particular user group to be able to read a file but not able to download it. Access control, activity monitoring, and command usage are areas in which supplemental security solutions bring a lot to the table for IBM i users.

Activity Monitoring

Activity monitoring is a critical functionality to keep tabs on employees and ensure that effective security policies are in place and function correctly. According to a 2010 SANS Institute white paper, failing to adequately monitor can result in a number of risks from malicious insiders. These include:

- **System sabotage:** When a privileged user attempts to disrupt network operation, usually by locking out other IT staff or by orchestrating a DoS attack
- **Introduction of bad code:** Either unintentionally or maliciously, IT professionals with access to corporate code or scripts can introduce poor coding practices into the system
- **Data theft:** According to SANS, this type of attack can be difficult to identify and is the most damaging type of insider threat

The SANS paper underscored the danger posed by malicious insider threats with two examples. In one case, a disgruntled employee with admin privileges injected scripts into his company's systems that would have disabled all system access if executed. In a similar case, a rogue administrator locked IT staff out of the systems that controlled San Francisco's critical infrastructure, leaving the city's network unusable for several days.

The SANS researchers also highlighted data theft as a main threat. The paper discussed an incident in which an authorized user logged in, escalated privileges and transmitted data over Secure Shell protocol. While it may seem like something easy to catch, the lack of effective monitoring solutions made it difficult for other IT staff to correlate the user's activity with the data being accessed.

As these examples show, a lack of command usage restrictions and monitoring capabilities can bring business operations to a halt. For this reason, it is important that a comprehensive IDS strategy allows for command monitoring. Software from PowerTech supports this functionality, ensuring that administrators are notified whenever a specific command is being used. This provides visibility in the event a user attempts to overstep his or her privileges. The software also allows the administrator to control who has the ability to use certain commands, preventing the types of threats outlined above.

Access Control

In general, it is a good idea to limit access as much as possible without also crippling employee productivity. IBM i supports built-in object level security, so access

to specific files can be restricted. The ability to monitor privileged user activity can be supplemented with more robust access control. Unrestricted access can pose a real risk not only to information security, but also to an organization's compliance posture.

This places a high value on ensuring that privileged users can only access the files needed to do their jobs. The SANS Institute recommends using software to segregate duties so that the risk of an "all powerful" account is eliminated. However, it's important to account for the human element in any business environment. Verizon's *2012 Data Breach Investigations Report* provided some insight into the importance of mitigating risk from multiple angles rather than relying on a single solution and assuming it is completely secure.

"One particular case illustrates the lack of individual attention that goes along with most of these attacks," the report authors wrote, describing a specific attack. "In this scenario, an online FTP server that had been misconfigured to allow anonymous FTP access, was under constant attack by brute-forcing tools (like most online systems)."

As this case underscores, a simple error in configuration can diminish the effectiveness of otherwise robust safeguards. Although Verizon's example is of an external threat agent, it's not too difficult to imagine a situation in which a minor error could open the gate to a malicious insider. For example, perhaps a user is accidentally given the ability to edit his or her own security profile—or, worse yet, other users' profiles. If the insider's intent is to steal sensitive data, the problem may go unnoticed for a significant amount of time. But if they use monitoring features like profile switch notifications, administrators are able to see exactly which access changes were made and when.

Further functionality enables date and time restrictions so that even privileged users only have access during acceptable time frames. This limits the opportunity for a disgruntled employee to compromise the system when fewer people are looking.

Compliance and Auditing

"Audit" may not be a word any administrator wants to hear, but it is a reality that most will have to face. Infrastructures for company technology are tasked with managing an ever-growing amount of sensitive data. Whether it's customer credit card data or employee Social Security numbers, auditors are tasked with ensuring that information is protected from threats both inside and out.

What may surprise more than a few admins is that one of the compliance pitfalls isn't directly related to security. In some cases, organizations fail to provide adequate documentation of policies, technological safeguards, and security safeguards. One of the necessary elements to ensuring compliance is providing a comprehensive view of activity, including attempted breaches and incident reports. Fortunately, IBM i's security audit journal offers clear visibility to external network threats, and can be further tailored to business needs through the custom threshold system.

However, IT administrators need to be able to quickly gather data relevant to specific security incidents. Storing security data across several disparate systems not only makes it more difficult to quickly track network and user activity, it leads to storage inefficiencies. Solutions like PowerTech Compliance Monitor offer benefits beyond basic compliance reporting by consolidating and compressing reports to a central system. By putting all security data in the same place, Compliance Monitor improves overall visibility and makes audits less painful.

Building a Comprehensive IDS

Especially as threats evolve and the digital ecosystem becomes a more dangerous place, businesses must adopt more effective security stances. It may sound like a monumental task, but it is really just a matter of shifting security mindsets. As many security experts have attested in recent years, there is no "silver bullet" for security, meaning that no single solution can account for all the risks companies must contend with.

This means that the best approach is one that employs an interconnected system of checks and balances. For IBM i users, this entails using the built-in IDS to develop effective security policies and, to ensure authorized users stay within accepted boundaries, supplementing those policies with monitoring solutions.

Conclusion

Intrusion detection is a key element of any organization's access control, data security, and regulatory compliance. To maintain security against a proliferating number and variety of attacks from both outside and within your organization, you must take advantage of the world-class security controls of IBM Power Systems, among them the inclusion of a network IDS in IBM i. You must also recognize where built-in features for IBM i do not provide protection—notably, from ignorant and malicious users, which requires additional IDS functionality.

To address intrusion detection inside and out, PowerTech solutions give you:

- Powerful access control and activity tracking (by user and IP address) of transactions originating from the network.
- Forensic reporting on 70+ audit events (and hundreds of security configuration metrics).
- Real-time security event notification.
- Ability to monitor and report on non-compliance.

For more information, visit **www.powertech.com/ids**.