

Managing Intrusion Detection Systems in Large Organizations, Part Two

Paul Innella 2002-04-09

Managing Intrusion Detection Systems in Large Organizations, Part Two

by *Paul Innella, Oba McMillan, and David Trout*, with assistance from *Rebecca Bace*

last updated April 9, 2002

This is the second part of a two-part series devoted to discussing the implementation of intrusion detection systems in large organizations. In the [first installment](#), we looked at some of the challenges of planning, integrating, and deploying IDSs in a large organization. In this installment, we will look at managing agents in a distributed environment, managing data from multiple IDS packages, and correlating data from distributed agents.

Managing Agents in a Distributed Environment

Global organizations face a monumental task in their efforts to implement an effective intrusion detection infrastructure. Their challenge is to manage agents residing in multiple locations around the world. There are numerous obstacles associated with geographically distributed IDS deployments, and we will highlight a few of the most significant ones.

1. Centralizing Management of Distributed IDS

The first concern relates to how the IDS hierarchy should be implemented. Organizations must determine the flow of commands and data from agents to their respective managers. Should all IDS communications and data travel from agents to distributed managers in the field, or should a single manager (or central manager farm) govern the IDS infrastructure? Data collection is a by-product of this hierarchy in that data should flow along the same path as commands. We firmly believe that, as a function of cost and logistics, it is simply smarter operational practice to house IDS managers centrally. The following are two important benefits of doing so:

1. Centralized IDS management decreases the amount of equipment (both hardware and software) that an organization must maintain and administer, resulting in:
 - o Reduced costs, due to fewer systems, minimal travel and maintenance expenses;
 - o Less administrative effort/manpower; and,
 - o More efficiency, as each manager can direct multiple devices over many networks

for numerous IDS solutions (as opposed to one manager for each subsection of the network).

2. Centralized control of IDS managers also simplifies the overall network architecture and reduces the number of vulnerable points in an organization's security infrastructure.

In short, an organizational IDS structure must be designed to ensure efficient and proper IDS data collection and agent operation, and thus we suggest using a centralized model.

2. Managing Data Communication in Distributed IDS

Another challenge in implementing and managing distributed intrusion detection systems is deciding how data should be communicated between the geographically dispersed modules. For instance, how do you send massive amounts of data from an IDS agent in Thailand back to the IDS management console in New York City? Not only does this traffic need to be encrypted and protected from unauthorized disclosure, it also needs a communication channel that will support large volumes of data. Depending on the aggressiveness of host auditing and the granularity of network monitoring, this can be a significant requirement. We have noticed some major industry trends in dealing with data transmission. The popular method is to rely on the native communication mechanism built into the IDS. Most products provide an "authenticated handshake" to establish communication. This is complemented with an encrypted channel for transmission between agents and their managers. In our opinion, the method of secure communications built into most IDS products is adequate.

It also seems to be common practice to leave communication port assignments set to their default port numbers. This presents a security vulnerability since most of us know that ISS RealSecure uses ports 901 through 903 and Symantec's NetProwler uses 5051 and 5052. Consequently, we recommend selecting random port numbers for IDSs. Some might argue that a port sweep will reveal this information, but we would counter with the notion that every small defensive measure is one more hurdle that an intruder must overcome. Careful planning needs to accompany decisions concerning agent-to-manager communications to ensure that operational requirements aren't adversely affected.

3. System-to-System Communications in a Global Network

System-to-system communications throughout a global network is another obstacle encountered in managing distributed IDSs. For example, the use of IDSs introduces the need for internal systems in London to be able to communicate with internal systems in New York

City. In order to operate effectively, these systems must transmit audit data, receive configuration updates, and exchange security alert information, often in real time. However, these systems are not always preconfigured to do so. The communications infrastructure that might include network address translation (NAT) or virtual private networking (VPN) must now incorporate these systems into their overall configuration for correct IDS deployment. For instance, an internal IDS agent in New York City may share the same private IP address as one in London. NAT would have to account for this so that the IDS manager understands which agent it is communicating with and how to interpret the data from each agent independently. Systems communications is not a trivial task; it often leads to breakdowns in a global IDS deployment.

4. Enforcing Software Updates

As discussed earlier, deploying an IDS infrastructure across a global enterprise is a demanding task. Once completed, though, worries related to the IDS itself are not over. A final challenge in managing agents from multiple locations remains: enforcing software updates and upgrades. Periodically, vendors will distribute software updates that contain valuable new functionality or patches. Organizations need to understand what is involved with implementing these updates, and what type of commitment is needed. Unfortunately, it is all too common for organizations to have a successful roll-out, only to encounter problems with intermittent software updates that cause the IDS infrastructure to become obsolete.

These obstacles highlight some of the many reasons why proper attention and planning must be given to manage IDS agents in multiple locations. However, these are not all of the concerns that a large organization must take into account in the planning and implementation of an effective intrusion detection system. There are further challenges in the proper management of the operational system.

Managing IDS Data from Multiple IDS Packages

Organizations frequently implement intrusion detection systems from several different vendors. As a result, they may be faced with the challenge of managing data coming from different sources and arriving in various formats. Though it may lead one to ask why an organization would purchase products from different vendors, there are a number of valid reasons. Primarily, IDS solutions from different vendors each have relative strengths and weaknesses. One vendor might be strong in the area of host-based intrusion detection, while another might have a

revolutionary network intrusion detection product. A third vendor's product might excel as a file integrity checker. Understandably, organizations want the best of breed in each area; thus the decision to purchase multiple products is made. Additionally, some organizations are so large that each business unit makes their own decision as to which IDS solution to purchase. Consequently, the central security division must integrate the various products. No matter what the motivation, the use of diverse IDS solutions creates the problem of dealing with data from different products.

The lack of standards in the area of intrusion detection makes integration of IDS solutions from multiple vendors a particularly complex issue. Many standards have been proposed that would allow for interoperability and the exchanging of messages amongst systems. One example of this is the [Common Vulnerabilities and Exposures™ \(CVE™\) dictionary](#), which provides a uniform name "for all publicly known vulnerabilities and security exposures." [MITRE01] Though admirable in their efforts, these standards have not enjoyed the widespread acceptance that one might have hoped for. Furthermore, although reporting interoperability such as CVE allows for consolidation of alarm mechanisms, it does not take into consideration information unrelated to intrusions, such as periodic unavailability due to scheduled software updates of an IDS agent. As such, there is no single, universally accepted way for these compartmentalized IDS systems to communicate with one another. Monitoring and analyzing data coming from multiple systems and different IDS solutions is a daunting task, even for the most seasoned security professionals.

Managed IDS Solution: Build It or Buy It?

The most viable solutions that we have seen for managing different IDS solutions are achieved in one of two ways: build it or buy it. To build such a data collection and analysis system requires a specialized knowledge and skill set that few in the industry truly possess. This system must be able to accept and interpret security messages from all deployed systems, most of which likely utilize different file formats and communication protocols. Most managed security services providers have faced this problem. Those that are able to efficiently and automatically manage varied IDS solutions are still in business while others have failed. Consequently, buying a previously developed product is sometimes the simplest solution, albeit probably the most expensive. Enterprise security management vendors such as NetIQ Corp, eSecurity Inc., and OpenService offer products that consolidate the events detected by various systems into one central repository for correlation and analysis. Many believe that this is the solution that most organizations will adopt as they struggle to collect, analyze, and respond to

security events that the various IDS tools uncover. While the decision to purchase IDS tools from multiple vendors is one that we fully support, organizations must understand the associated limitations - be prepared to either build or buy a solution that will tie the various pieces of the IDS infrastructure together.

Correlating Distributed Agent Data

Another major issue facing large organizations is the collection and correlation of IDS data from distributed agents. The first factor that effects data collection is storage. Depending on the nature of their business, organizations have a responsibility to themselves and to their customers to store IDS data for a given period of time. But how long should one retain this data? Typical periods of storage range from 90 days to two years. Keep in mind that IDS data can become copious, and in large organizations, two years of data usually equals terabytes of storage space.

Companies must also decide what level of security they should apply to the storage of IDS data. This is determined by the organizational requirements for the IDS data that is collected. If it is imperative that an organization maintain a qualified chain of evidence, then the means by which that data is stored is extremely important.

After the data is collected, organizations must make sense of it all. This may be the hardest part of the process. IDS data correlation from distributed agents is extremely difficult. The problem does not lie in identifying intrusions; rather, it stems from issues in the linkage of events across a network. For instance, the process of associating a given attack on a system in London with another one occurring in New York City is complex, if not impossible. It is tricky to discover attack patterns on a single system when the attacks are spread out over time and methodically executed, let alone on a distributed IDS.

There are very few intrusion detection systems that provide organizations with a heuristic analysis of combined data or even a static-state assessment of correlated intrusions from separate IDS agents. Furthermore, we have yet to see any tool that will actually analyze and provide useful diagnostic information working with different IDS vendors' data. However, there are some high-bandwidth network-based IDS systems that allow a customer to monitor several points in their network and correlate the results. The results of such correlation are so far encouraging.

Perhaps the most significant progress in this area has come from the managed security services providers who have developed proprietary solutions for correlating IDS data from distributed agents and different IDS solutions. As an example, Ripstech's CaltarianSM technology processes IDS data from numerous customers with different IDS solutions installed in their environments. They are capable of recognizing attacks that originate in different areas, but target one victim. Additionally, they can link distributed attacks such as those recently launched from China. In order to bring IDS to bear, the most important task to accomplish is correlating IDS data. It is imperative that we evolve from simply detecting single intrusions, to actually understanding the relationship between distributed attacks against different systems and environments.

Conclusion

In conclusion, there is a definite need for intrusion detection systems in an overall security infrastructure; however, in large organizations, this may be easier said than done. There are many obstacles to the deployment, configuration, management, and data handling of the various IDS solutions. Many of these challenges have already been faced and overcome by organizations that have successfully implemented IDSs. Fortunately, we can learn a lot from their struggles. Our own experiences in implementing IDS solutions in major organizations have revealed many of these problems, as well as the associated solutions. We offered these experiences in the hopes that you might properly prepare for the successful implementation of enterprise-wide IDS protection.

References

[MITRE01] The Mitre Corporation, 2001, Common Vulnerabilities and Exposures Web site:
<http://www.cve.mitre.org>

Paul Innella, CISSP, is President of [Tetrad Digital Integrity \(TDI\)](#), a Washington DC based information security services company offering [one day seminars on intrusion detection](#). Mr. Innella has nearly ten years of experience in the computer industry working at several commercial and government companies where he held the role of security engineer, developer, integrator, systems administrator, program manager, and sales engineer.

Mr. Oba McMillan, CISSP, is Vice President of TDI. He has many years of network security and IT related work experience. Mr. McMillan has published a number of articles on information security and recently wrote an article for Secure Computing Magazine Opinion newsletter entitled "PDA policies – Worth Their Weight In Gold." He also contributed to a SC Magazine article entitled "Road Warriors: Be Careful Out There" about wireless security.

David Trout, CISA, CISSP, is President of [SecureIT Consulting Group](#), a Northern Virginia based IT audit and information security consulting firm [that offers one day seminars in intrusion detection](#). Mr. Trout brings over seven years experience in the areas of IT Audit and IT Security Consulting.

Relevant Links

[Managing IDS Solutions in Large Organizations, Part One](#)

Paul Innella, Oba McMillan, and David Trout, SecurityFocus

[Privacy Statement](#)

Copyright 2006, SecurityFocus