Best Practices in Intrusion Detection System Implementation

by Erin Buxton

October 2002

A "one-size-fits-all" approach to Intrusion Detection Systems (IDS), while sometimes touted by vendors to convey the ease of deployment of their products, will most likely not suit the needs of a company and can even be dangerous. There are a number of challenges and pitfalls to overcome. A successful solution that fulfills requirements and budget constraints can only be achieved by carefully considering each factor and component.

Common Issues for Deployment and Maintenance of an IDS

- **Implementation of some vendor features without careful review may cause serious unintended effects**
- **Traffic "sniffing" is hard to deploy cost-effectively**
- **It is hard to discern normal behavior, which is different for each organization, from ill-configured servers and hacker activity**
- **An excessive number of alarms can create a self-induced denial of service**
- **Logs capture essential evidence, but IDS products may not have adequate log databases**
- **Additional pitfalls need to be addressed through proper planning and budgeting, and communication with all stakeholders**

"Sniffing" Configuration Issues

The ability to perform intrusion detection should be considered during the initial design of a network. However, this is not always the case, and IDS vendors do not always introduce the subject. "Sniffing," or monitoring and analyzing traffic across a network segment, is possible using hubs, but this solution is typically no longer used due to the number of collisions that occur with such a design. Switched networks are most frequently used, as they segregate traffic and send packets only to the ports for which they are intended. Unfortunately, this renders sniffing impossible without additional configuration. There are several solutions, each with pros and cons.

One solution is to use the "span port" capability of the switch, which mirrors traffic. The IDS system can then be plugged into the span port to sniff data. However, some companies do not allow spanning. Also, mirroring multiple ports may not be possible due to high traffic levels or limitations of the IDS sensor or switch.

Network taps can also be utilized through placement between two network devices. But taps present issues as well. They can be expensive, may not support all features of the IDS, require the IDS to support a promiscuous-mode Network Interface Card (NIC), and only capture network traffic in one direction—usually received (RX) data. Therefore, some attacks may not be detected. To overcome this last point, two IDS sensors are necessary to capture network traffic in both directions, which increases hardware and software licensing costs. If the traffic in each direction is not at the upper limit of the sensor capabilities, the sensors will not be fully utilized, thus reducing their value.

To address this concern, taps may be consolidated through a switch. A span port would then be set up on the switch and connected to a sensor. This would reduce the number of IDS sensors, hence reducing the cost while increasing the value of the sensor by maximizing utilization. Yet this solution is not itself devoid of drawbacks; span ports can be overloaded, and depending on the sensor capabilities, it may only be advantageous when consolidating links with low utilization.

Using a load-balancing device specialized for intrusion detection devices would resolve many of these drawbacks. These switches allow for traffic load balancing across multiple span ports, provide for a diversity of port speeds and allow for load balancing to two or more sensors to maximize usage. This is a more comprehensive solution, however it is expensive and fairly new. In addition, the company may not support this type of switch; administration, maintenance and training may be costly; and the sensors still need to work with a promiscuous mode NIC.

Regardless of the solution, the customer should verify that the vendor will support the products selected.

Monitoring Alarms

When using a network-based IDS, it is important to understand what protocols and traffic are traveling over the network. Practice and time with the network IDS will improve recognition and understanding of normal traffic, violations and attacks. One can also initiate attacks using freeware security scanners, which can help train administrators and test IDS configurations and automatic alert systems. Stress tests can be run to observe alerts generated and stress imposed on the network, e-mail server, network management tool and firewall, depending on what types of alarms are being used. It is important to do this in a lab environment so as not to disrupt traffic on the network. When doing this work for a separate company, written permission is usually required before performing any scans on the network, and network administrators should be informed of these scans to prevent downtime.

Many network-based IDS systems will give details on each violation type and the standard priority level given. Ultimately, the administrator will have to make the call on traffic legitimacy and response methods.

Scaling Alert Notifications

Although many vendors offer a range of features, including automatic notification and automatic device reconfiguration, one should be wary of implementing these on a large scale, on a high-traffic network or with an IDS that alerts administrators each time a violation is discovered. While these features appear exciting and useful, they can cause more harm than good. In the beginning, receiving numerous alerts is normal until the IDS policy is refined to fit the network. It is important to test the IDS with e-mail and network administrators; if they can quickly remedy issues, this will prevent embarrassment and costly downtime. Some network management tools can reduce the number of alarms by alerting administrators only after a certain number of violations have been discovered.

For e-mail alerts, the information sent should not violate company security policies, such as giving away internal IP addresses. Security policies within the critical network segments may need to be stricter due to the sensitivity of the information on those systems.

Managing the IDS Log Database

Although IDS vendors include a database in the product, many IDS products need another database application to handle the logs and alarms created. This can increase the cost of the network and create new database administration issues. For very small networks, another database may not be necessary. However, depending on the configuration of the IDS and level of attacks, logs may quickly increase, causing database corruption and information loss. Therefore, in order to forecast any additional costs and administrative issues, one should investigate which database servers are supported by vendor products.

Other Potential Pitfalls

Common non-technical issues include financial and personnel constraints, over-reaction and apathy.

Creating hardware and software lists with cost justification and/or risk analysis will assist with budgeting. However, it is only after lab testing that final costs can be calculated.

If the company does not have IDS expertise, seeking help outside the organization will improve the chances of success. It can be costly and dangerous to the company's reputation to let untrained personnel use the system. Ongoing product and security training must be included in the budget. Administrator duties should be separated between groups in order to allow individuals sufficient time to perform duties and to crosscheck security policies and behavior.

Figure 1: Many components factor into the design and maintenance for an IDS solution.

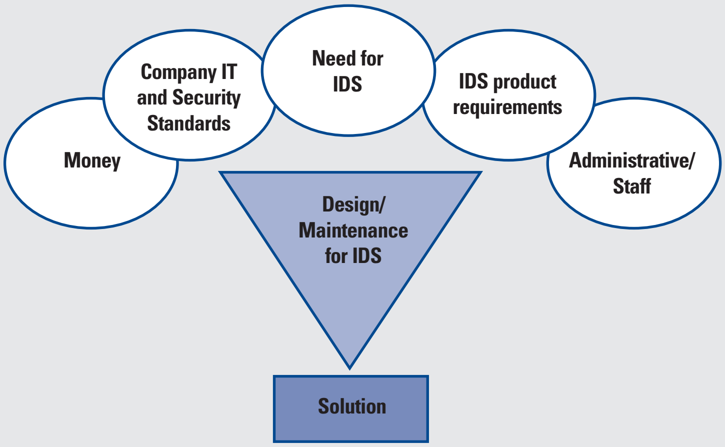

In addition, lack of training can result in over-reaction to alerts. Administrators need to understand the network and what protocols run on it. If administrators have not been trained on how to research the alerts (i.e., causes, false positives and network functionality), they may block legitimate traffic, thus causing a self-induced denial of service. This can also create a "crying wolf" administrator, whom no one believes when an actual attack occurs.

Finally, over-reaction may be followed by apathy. If an IDS is configured to notify the administrator too often, the administrator will become apathetic to alarms. The IDS can still be configured to log all suspicious behavior that can be reviewed for patterns and auditing purposes, but notifications should be selective enough to reduce potential apathy. This is a difficult balance to achieve and requires familiarity with the network and security, which usually comes after months of watching over the IDS.

Involve Other People to Avoid Surprises

All groups impacted by the design or operation of the network and security should be consulted to ensure that no conflicts exist. Although it may be frustrating to consult many people about the design, involving others early will save time, frustration and money. With written approval of a design, awareness will be heightened and conflicts avoided.

Know The Enemy

The obvious enemy is the outsider—the hacker. But is that really the only enemy? Often, internal attacks pose the greatest threat; insiders typically know which systems to attack and may be authorized to access them, making these intrusions difficult to track and prevent. With knowledge of the enemy, reviews of network-based and host-based IDS configurations should be conducted to verify that internal threats are being considered. All suspicious behavior, external and internal, should be logged to avoid being surprised by internal attacks. Security policies should be reviewed to ensure that they include administrator accountability and crosschecking.