SECURITY

Security Operations Center

Building, Operating, and Maintaining Your SOC

Joseph Muniz
Gary McIntyre
Nadhem AlFardan, CCIE No. 20519

ciscopress.com

Security Operations Center

Joseph Muniz
Gary McIntyre
Nadhem AlFardan

Cisco Press
800 East 96th Street
Indianapolis, Indiana 46240 USA

Some SOC tools require access to the Internet or other external networks. This includes, for example, access to threat intelligence information or downloading signatures and reputation information. Such access should be provisioned through a proxy or application layer firewall limiting access to necessary services only. This can be accomplished using a whitelist of destination domains and IP addresses, allowing the SOC tools to access only a predefined set of Internet destinations and protocols.

Access to Systems

The SOC enclave will be connected either in-band or using an out-of-band network with direct or indirect access to systems monitored by the SOC. It is recommended that devices be monitored using an out-of-band network allowing secure access that is not impacted by the in-band traffic, especially in the case of an attack. Today, most enterprise-level devices are provisioned with one or more management ports that can be connected to the logical or physical out-of-band management network. Access to this management network should be available only to authorized entities such as NOC and SOC members and devices.

Access to the network components should be authorized, preferably through a centralized authentication, authorization, and accounting (AAA) service and using network protocols such as Remote Authentication Dial-In User Service (RADIUS) or Terminal Access Controller Access Control System (TACACS). Using an AAA server guarantees that all access attempts are logged with an audit trail that can be retrieved during an incident investigation. Best practice is using as least two factors for authenticating users to SOC systems. Multifactor authentication solutions typically are a mix of something you know (password), something you have (access cards, hard or soft tokens), and something you are (fingerprint scanner). Communication between devices, such as from a user desktop to a SOC device, should leverage some form of encryption such as Secure Shell (SSH), Secure Sockets Layer (SSL), or Transport Layer Security (TLS).

Security

Protecting the SOC network, hosts, and applications is of topmost priority. Compromising the SOC network and systems undermines the objective of having a SOC and will negatively impact detection capabilities, organizational trust, and justification for future investments. Developing a SOC enclave using network segmentation and security controls can prevent a breach from happening.

Developing a SOC enclave starts with segmenting different trust levels of traffic using either physical separation or different variations of virtual network segmentation such as VLANs, access lists, or security group tags (SGT). Inbound and outbound connections from and to the SOC network segments should be monitored and controlled in support of the least privilege principle, in which access is limited to what a service or a user requires to operate.

When firewalls are used, firewall segmentation rules should provide standard network-based access control that applies to incoming, outgoing, and intra-SOC traffic. The rules should be managed and maintained by the SOC team and should follow a standard change management control process. Maintaining the firewall rule base should include regular reviews and updates. For example, expired rules should be manually or automatically disabled or deleted. All rules should also be properly documented. Each rule should be associated with a description that includes information such as the following:

- Requestor
- Purpose
- Creation and expiry dates
- Change management approval reference number

Intrusion prevention control should be integrated with the SOC enclave design such that traffic is inspected inline wherever possible. The intrusion prevention system (IPS) can be implemented as a standalone solution or can be integrated with the firewall. The IPS mode of operation can be set to fail-close or fail-open. When configured in fail-open, an IPS failure would result in the IPS acting as a network bypass, allowing traffic through with no inspection. Your security policy should identify the acceptable mode of operation based on understanding the risks associated with both options.

The IPS signatures and rules should be tuned such that irrelevant ones are disabled, while relevant signatures are manually or automatically enabled and tuned. It is also extremely important to consider the performance impact resulting from enabling too many rules and from operating the IPS and the firewall functions in the same system. Make sure to consult with your vendor on the performance impact before enabling multiple security features.

Note Some IDS/IPS solutions such as Cisco Firepower Management Center (FMC) offer autotuning or recommendations of IPS rules to enable or disable to assist with IPS rule management.

There might be a requirement to inspect traffic within the same network segment such as the same VLAN. One way to accomplish this is by implementing an IDS so that traffic from one or more VLANs is copied to a switch port configured where the IDS is connected, thus enabling you to monitor and inspect intra-VLAN traffic. In many cases, the same system can be configured in IPS and IDS modes of operation.

Figure 5-6 shows the implementation of the IPS and IDS functions within the SOC enclave shown in Figure 5-5.

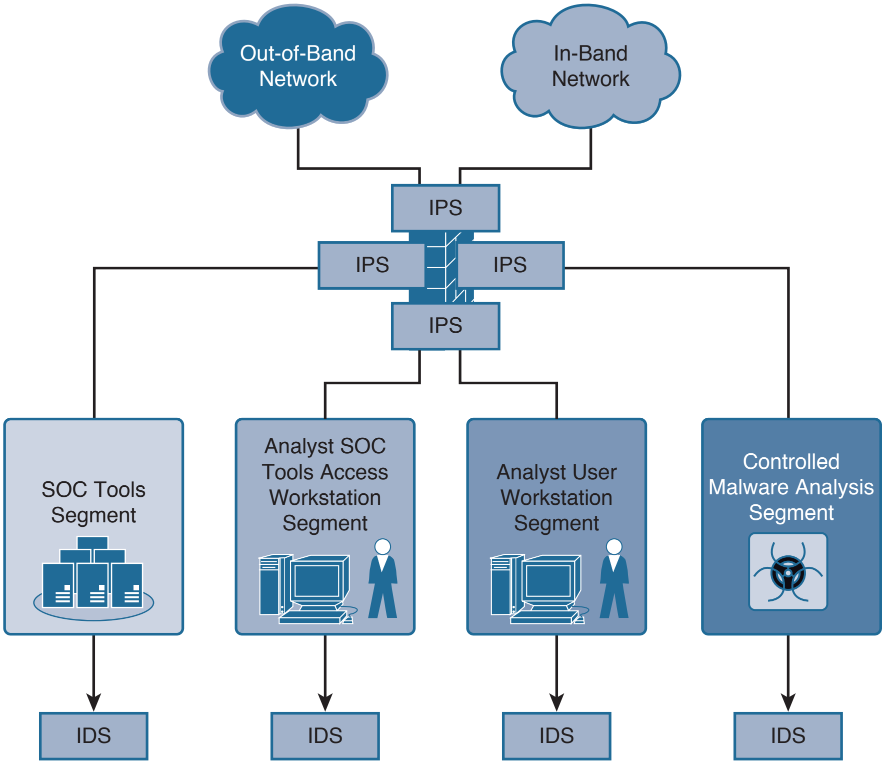
Out-of-Band Network
In-Band Network
IPS
IPS
IPS
IPS
SOC Tools Segment
Analyst SOC Tools Access Workstation Segment
Analyst User Workstation Segment
Controlled Malware Analysis Segment
IDS
IDS
IDS
IDS

Figure 5-6 *Integrating IPS and IDS Functions in the SOC Enclave*

Another method to identify intra-VLAN traffic along with other areas of the network is by leveraging network telemetry, also known as NetFlow. For example, Lancope's StealthWatch can either digest NetFlow from existing networking equipment or place a StealthWatch NetFlow sensor that converts raw packets into NetFlow.

Additional network-based security tools such as breach detection can be used to identify network and security events within the SOC network. Chapter 9, "The Technology," provides more information about the technology options that are leveraged by and used to secure the SOC.

Another focus area for SOC security is protecting host systems and applications. For example, a security information and event management (SIEM) tool should have a web-based interface that is developed with the use of components like Tomcat,[1] JBOSS,[2] and OpenSSL.[3] Exploiting vulnerabilities in these components may result in a compromised SIEM that could provide an attacker with privileged access to other applications and systems in and out of the SOC.

- The first step to avoid host system compromise is to enforce security configuration benchmarks for operating systems and applications. Examples of secure configuration benchmarks include the U.S. Department of Defense (DoD) DISA Security Technical Implementation Guides (STIG)[4] and the Center for Internet Security (CIS)[5] set of benchmarks. You may also refer to the system provider for specific hardening recommendations.
- It is critical to stay on top of implementing the required and recommended security patches and fixes for operating systems, applications, and databases. Typically, the cycle of testing and implementing critical security patches can take a shorter time in the case of SOC when compared to other environments such as IT. Why is this so important? According to the 2015 Verizon Breach Investigation Report (VBIR),[6] studies showed that 99.9 percent of the exploited vulnerabilities were compromised more than a year after the CVE (method to patch) was published. This means that the majority of businesses that had a vulnerability exploited had over a year to fix it based on when the vulnerability was publicly announced. The same report showed that the most-exploited vulnerabilities for 2015 have been known since 2007. Chapter 7, "Vulnerability Management," covers best practices for vulnerability management.
- All host systems and servers should have protection software with signature-based and behavior-based protection features. Most vendors in this market space offer feature bundles that can include antivirus, antimalware, content filtering, host-based IPS, and so on.
- Another good practice is to enable process whitelisting or to use breach-detection applications to prevent unauthorized programs from running. These solutions offer capabilities beyond signature based antivirus by monitoring and controlling what processes are used by applications and processes and by looking for unusual endpoint behavior. Cisco Advanced Malware Protection and Bit9 are examples of breach-detection solutions.

Compute

Compute generally refers to the processing resources available on a system, represented in terms of processors, cores, and speed in gigahertz (GHz). For example, a quad-core 2-GHz processor contains four independent 2-GHz cores, where a core is a CPU that executes machine instructions. Common CPU architectures include x86 (32 or 64 bit), SPARC, and PowerPC. Note that the compute discussion is always coupled with the available system memory resources. Compute can also be carved up using virtualization technology to maximize the investment in the system. There are pros and cons to using virtualization versus dedicated compute.

Dedicated Versus Virtualized Environment

With today's technology, you have the choice of provisioning compute resources in dedicated and virtualized modes. *Dedicated* refers to deploying operating systems directly on bare-metal hardware. *Virtualization* enables you to host multiple guest machines that share the same underlying hardware, optimizing the allocation of resources in terms of processing power, memory, network connections, and so on. Examples of virtualization software include commercial products such as VMware vSphere,[7] Microsoft Hyper-V,[8] and open source products such as KVM,[9] Xen,[10] and VirtualBox.[11] Figure 5-7 is an example of how system virtualization can run multiple separate virtual systems on the same hardware.

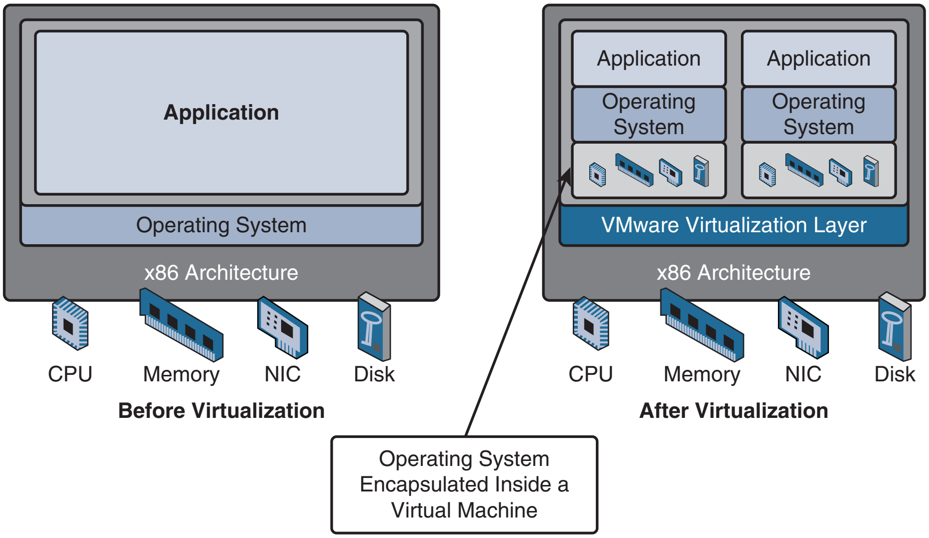

Figure 5-7 *System Virtualization*

Applications tend to perform slightly slower when operated in a virtualized environment. For example, according to the Splunk capacity-planning manual, the Splunk application can index data about 10 percent to 15 percent slower on a virtual machine than it does on a standard bare-metal machine. Similar performance drops are expected with other open source and commercial software.

When considering virtualized technology, it is important to verify that all applications can operate and are supported for virtualized environments. This includes identifying the performance impact that results when deploying the SOC application in the virtual environment, which may vary based on available resources on the hardware hosting the virtual services. Best practice is also to include logical segmentation of SOC and non-SOC applications and users in a virtualized environment to ensure data protection.

The amount and type of compute resources required to operate the SOC infrastructure depend on a number of factors such as the SOC services; the model of data collection; the location, type, and amount of data to collect, process, store; and the technologies of choice. Remote sites generally connect to the rest of the network using what is considered low-bandwidth links, limiting the amount of nonproduction traffic that can be sent to a central collection location. This typically forces you to deploy compute resources near the data for collection and possibly analysis purposes.

The amount of data can be expressed in terms of events per second (EPS) in the case of collecting logging messages, flows per second in the case of collecting network flows, and (K/M/G)bps in the case of capturing network packets. The amount of data to collect and process has a direct impact on how you design your compute infrastructure in terms of the number of collection engines to deploy to distribute the workload of processing incoming events and the number and speed of CPUs and the amount of memory to deploy per machine.

Let's take Splunk as an example. Figure 5-8 estimates the number of Splunk search heads and indexers to deploy based on the daily data volume processed by the indexers and the number of system administrators accessing the search head web interface. A Splunk search head is an instance with a web interface that accepts users and directs search requests to place the results into an index.

	Daily Indexing Volume						
	<2GB/day	2GB to 250GB/Day	250GB to 500GB/Day	500GB to 750GB/Day	750GB to 1TB/Day	1TB to 2TB/Day	2TB to 3TB/Day
Total Users: Fewer Than 4	1 Combined Instance	1 Search Head 1 Indexer	1 Search Head 2 Indexers	1 Search Head 3 Indexers	1 Search Head 4 Indexers	1 Search Head 8 Indexers	1 Search Head 12 Indexers
Total Users: Up To 8	1 Combined Instance	1 Search Head 1 Indexer	1 Search Head 2 Indexers	1 Search Head 4 Indexers	1 Search Head 5 Indexers	1 Search Head 10 Indexers	1 Search Head 15 Indexers
Total Users: Up To 16	1 Search Head 1 Indexer	1 Search Head 1 Indexer	1 Search Head 3 Indexers	1 Search Head 4 Indexers	2 Search Heads 6 Indexers	2 Search Heads 12 Indexers	2 Search Heads 18 Indexers
Total Users: Up To 24	1 Search Head 1 Indexer	1 Search Head 2 Indexers	2 Search Heads 3 Indexers	2 Search Heads 4 Indexers	2 Search Heads 6 Indexers	2 Search Heads 12 Indexers	2 Search Heads 18 Indexers
Total Users: Up To 48		1 Search Head 2 Indexers	2 Search Heads 3 Indexers	2 Search Heads 4 Indexers	3 Search Heads 8 Indexers	3 Search Heads 16 Indexers	3 Search Heads 24 Indexers

Figure 5-8 *Splunk Recommended Configuration*