Cisco Protects Data Center Assets with Network-Based Intrusion Prevention System

Cisco Computer Security Incident Response Team (CSIRT) detects and mitigates network threats before the onset of data loss or service interruption.

Cisco IT Case Study / Network Security / Network-based Intrusion Prevention: This case study describes why and how Cisco® augmented its perimeter-based intrusion prevention system (IPS) deployment with network-based IPS in data centers. Although perimeter-based IPS sensors detect malicious traffic that traverses the company firewalls, they do not detect suspicious traffic that both originates and terminates within the company. Network-based IPS sensors at Cisco help protect the company's most important assets, which reside within data centers. Cisco customers can draw on Cisco IT's real-world experience in this area to help support similar enterprise needs for intrusion prevention.

"Network-based IPS enables us to detect and mitigate internal security events before users experience a secondary impact, such as… service disruption, loss of intellectual property, or infection."

Gavin Reid, Manager of Computer Security Incident Response Team, Cisco

Business Challenge

To protect network availability and intellectual property, Cisco constantly monitors its network for a large assortment of threats, including insider threats; unauthorized access to data center assets; policy violations, such as privileged account logins with cleartext protocols; and botnets, which are used by miscreants and organized criminals for phishing and stealing intellectual property. Malicious activity can originate from outside the company walls or from within, usually from infected or compromised hosts.

Intrusion prevention systems (IPS) are an important part of the defense-in-depth strategy at Cisco, complementing firewalls; access control lists (ACLs); Cisco Security Agent, a host-based IPS; and anomaly detection based on analysis of NetFlow statistics. When Cisco first deployed a Cisco Intrusion Detection System (IDS) sensor, in 2000, the main objective was to detect threats traveling across the network perimeter, either from the outside in or the inside out. Therefore, the Cisco Computer Security Incident Response Team (CSIRT) deployed IPS sensors in the Cisco DMZ and at service provider points of presence (POPs).

But while perimeter-based IPS deployment did what it was supposed to, it did not provide visibility into security events inside data centers, where Cisco's most valuable assets reside. "Perimeter-based IPS does not detect malicious events that are contained in the trusted cloud within Cisco," says Jayson Mondala, IT security engineer. "The sensors can only see network conversations that traverse the perimeter between the intranet and Internet."

EXECUTIVE SUMMARY
BACKGROUND • Cisco constantly monitors its network for threats. • Company previously deployed perimeter-based Intrusion Prevention Systems (IPS)
CHALLENGE • Protect data center assets • Avoid false positives and false negatives
SOLUTION • Deployed Cisco IPS in data centers, engineering server rooms, and offshore development centers
RESULTS • Protected data center assets • Enabled early detection and remediation • Helped prevent data loss or service interruption
LESSONS LEARNED • Plan for IT resource requirements. • Include a team member with in-depth understanding of IP subnetting and internal networks • Remember that IPS does not replace preventive security methodologies.

Lacking a view into data center traffic, the Cisco CSIRT only became aware of unwanted network traffic after it had caused a problem. For example, only after users reported sluggish performance on a WAN link from a branch office to headquarters would the Cisco IT networking team discover that the link was saturated with malicious traffic. In 2004, Cisco CSIRT discovered 8 percent of security incidents; the remaining 92 percent were reported by Cisco users or external agencies. This ratio was far from ideal: “The CSIRT would prefer to be the first to know of an incident so that it can take corrective action before harm occurs,” says Gavin Reid, manager of the Cisco Computer Security Incident Response Team (CSIRT).

Network-based IPS deployment would provide high value to Cisco by furnishing actionable information about threats to the company’s most valuable assets. “Companies the size of Cisco receive millions of attacks from the Internet daily, but only a fraction of the attacks are dangerous or can be acted upon,” says Reid. “In contrast, attacks that originate from inside Cisco are always significant. Not only are they specifically targeted at Cisco, but we have the power to remediate the systems initiating the attacks because we own them.”

The shift in emphasis from perimeter-based to internal IPS at Cisco also reflects recent changes in the threat landscape. “When we originally deployed Cisco IPS on the perimeter, we were protecting ourselves against noisy worms, such as Nimda and Code Red,” says Mondala. “The IPS sensor detected worms leaving the company, helping Cisco to be a good ‘netizen.’” Today, however, this type of virus is less prevalent, and the DMZ is less relevant because it is just one of many ways to damage the company. This has freed up the Cisco CSIRT to focus its efforts on more targeted threats, which are also more dangerous. “In 2002, a typical threat involved 1000 computers launching a concerted attack on one server,” says Mondala. “In 2007, a larger concern is preventing a hacker from quietly gaining access to our data. Rather than simply protecting Cisco’s reputation by helping ensure that our hosts are not propagating worms, we are also deploying Cisco IPS sensors in our data centers to protect intellectual property.”

Network Solution

In 2005, CSIRT began augmenting the Cisco IPS deployment in the Cisco DMZ with additional Cisco IPS sensors on the Cisco network: data centers, engineering server rooms, and offshore development centers. “Network-based IPS enables us to detect and mitigate internal security events before users experience a secondary impact, such as a Cisco server attacking an outside server, service disruption, loss of intellectual property, or infection,” says Reid.

Network-based IPS at Cisco involves monitoring and analysis, investigation, tuning and custom signatures, and deployment.

Monitoring and Analysis

Cisco IPS sensors collect information about network security events that match threat signatures. The product comes with more than 2000 signatures for known threats to applications and operating systems, which Cisco updates constantly. In addition, CSIRT continually develops custom signatures to meet the company’s unique needs. Cisco CSIRT uses alert-management software, such as Cisco Security Monitoring, Analysis, and Response System (MARS) and netForensics, to generate reports from the IPS sensor data at regular intervals or on demand. A team of

security engineers reviews the reports to identify potentially significant events. If a report is deemed important, the engineer takes corrective action or sends a page to a 24-hour response team.

Investigation

When command-and-control servers are detected, the response team temporarily applies a Border Gateway Protocol (BGP) blackhole. This prevents outside servers from controlling the affected desktop computer until the desktop support team remediates it.[1] "During one incident, the monitoring team saw the alert fire again after the BGP blackhole had been implemented," says Chris Fry, network engineer. "That should not have happened. We investigated it and discovered that the BGP mesh on the DMZ had failed for around 30 minutes. Without the Cisco IPS, we would not have known about the failure and been alert to possible consequences."

If the alert-management software shows a threat with questionable risk, an IT security engineer turns on the "produce verbose" alert feature of the Cisco IPS, which is a miniature packet-capture utility. The feature saves the "triggering" packet and performs deep packet inspection to provide more information. "The expression join # can indicate either IRC traffic, which is a potentially dangerous, or a database table join, which is benign," says Jeff Bollinger, IT security engineer. "The produce verbose alert can distinguish between the two by using deep packet inspection."

Tuning and Custom Signatures

Using the signatures with which it is shipped, an untuned Cisco IPS sensor deployed on the Cisco DMZ or in a data center would report 2 million or 10 million events daily, respectively. Therefore, Cisco CSIRT staff tunes the sensors to report only the events that pose significant threats to important assets. "The goal of tuning is to reduce false positives while not generating false negatives," says Bollinger.

Tuning requires in-depth knowledge of the network. "Our security posture improved markedly when we engaged people from our network team, who truly understand the unique workings of the Cisco network," says Reid. "They understand how the Cisco IPS interprets events such as a misconfigured router, which might generate 30,000 alerts. With this knowledge, we can tune the sensors to provide more useful information."

To prioritize alert response, it is important to know the location of important assets. "If an attack is under way on a data center subnet, it is much easier for the event-analysis teams to focus in on this activity if they know the affected subnet," says Mondala. "It's even better if the alert has a built-in source/destination tag." Cisco's monitoring teams use Cisco products that can use network information on networks, hosts, and protocols to enhance the quality of signature alerts.

Tuning often involves applying different policies based on the sensor's location in the network. "The Cisco IPS sensors in our DMZ are tuned broadly, while the sensors in the data center are tuned more narrowly, to exclude what we know to be valid traffic," says Mondala. Valid data center traffic, for example, includes network backup traffic, network file sharing, and Secure Shell (SSH). Outbound peer-to-peer file sharing, in contrast, is considered dangerous. Examples of tuning that Cisco has performed for the IPS sensors deployed in data centers include:

- Excluding legitimate management pings: The Cisco IT group frequently pings all of its data center equipment to check availability. The sensor needs to be able to distinguish these legitimate pings from those issued by malicious hosts. To accomplish this, Bollinger used the network-locale feature of the Cisco IPS software to enumerate all the addresses of Cisco management systems that might issue pings, and created a variable called "management servers." Then, he added a single tuning statement, "Drop alert from management systems," to apply the policy to hundreds of management systems.

1 To read a white paper on remotely triggered blackhole filtering, visit: http://www.cisco.com/application/pdf/en/us/guest/products/ps6642/c1244/cdccont_0900aecd80313fac.pdf.

- Excluding scans originating from the Internet: All enterprises receive a daily deluge of random scans from the Internet, which are not especially risky. Scans originating from within the Cisco network, however, indicate that the host performing the scan is infected and needs immediate remediation. Therefore, Cisco tuned the signature to only fire if the source of a scan is a data center host.
- Excluding scans originating from vulnerability assessment tools: Cisco regularly performs vulnerability assessments of data center systems using outside services as well as internally developed network scanners. Scans from systems on the list of vulnerability assessment tools are ignored, while scans from systems not on the list are captured and escalated for action.

Fry notes that a network-based IPS deployment in a data center must be tuned for many more protocols and applications than a perimeter-based deployment—perhaps dozens compared to half a dozen. Therefore, the person who performs tuning needs a good understanding of protocols, applications, IP subnetting, and internal networks. “The protocols and applications used on the perimeter are typically limited to e-mail, DNS [Domain Name Server], Web, and SSH,” Fry says. “The data center, in contrast, introduces many more potential false positives, including Windows protocols, SIP [Session Initiation Protocol] for IP telephony, protocols used between data centers, and protocols used between users and data centers.”

Bollinger adds, “Cisco IPS is not something to set up and forget. Its value is proportionate to the time that we spend customizing it to our unique network environment. It is critical to dedicate a resource to tuning at least part time, to help ensure the data that the sensor provides is useful.”

Cisco CSIRT further customized the Cisco IPS sensors by developing custom signatures. The signatures that are included with Cisco IPS detect threats that are common to all organizations, such as packets containing a regular expression that is an attack directed to Windows servers. The Cisco IPS product team releases new signatures whenever new threats are discovered. CSIRT augments these general signatures with signatures that it develops itself, for vulnerabilities and threats specific to the Cisco network. “Every organization’s network has different vulnerabilities,” says Bollinger, who spends about one-third of his time developing custom signatures and signature tuning. Some of the signatures that Bollinger has developed for Cisco CSIRT identify:

- Hosts inside the network that are communicating with command-and-control servers on the Internet that run Internet Relay Chat (IRC) botnets
- IRC traffic, characterized by a username followed by a nickname, that is traveling over nonstandard ports
- HTTP uploads from certain sites. This signature is intended to prevent intellectual property from leaving Cisco

Writing a custom signature is fairly easy, according to Bollinger, and is based on pattern matching using regular expressions. “The tricky part is not writing the signature, but rather knowing what to look for,” he says.

After developing a custom signature, Bollinger deploys it in a lab environment with one Cisco IPS sensor and a couple of Cisco routers and switches, and tunes the signature until it works properly. “Signature development is an ongoing activity at Cisco because the attack landscape is always changing,” he says.

Deployment

The Cisco team developed two designs for network-based IPS, based on aggregate bandwidth utilization. Large data centers, which use 600 Mbps bandwidth or more, have a Cisco Catalyst 6504E Switch, which provides load balancing among four or eight Cisco IPS 4260 sensors. The engineering data center at company headquarters in San Jose, California, has eight sensors on a 10-Gbps backbone.

Cisco wanted to help ensure that the solution design and implementation did not interfere with data center production activities. Therefore, both the large and small data center designs are based on passive monitoring, which has no impact on network performance, rather than inline monitoring. The switches receive a copy of network traffic, which is copied into a Remote Span (RSPAN) VLAN and filtered by a VLAN access control list (VACL). The VACL filters out encrypted traffic, which is not inspected. All other traffic is load balanced with EtherChannel to the Cisco IPS sensors. “Filtering out encrypted and uninteresting traffic maximizes our investment by enabling us to handle more traffic with the same number of devices and reducing hardware support costs,” says Fry.

Figure 1 Design for Large Data Centers

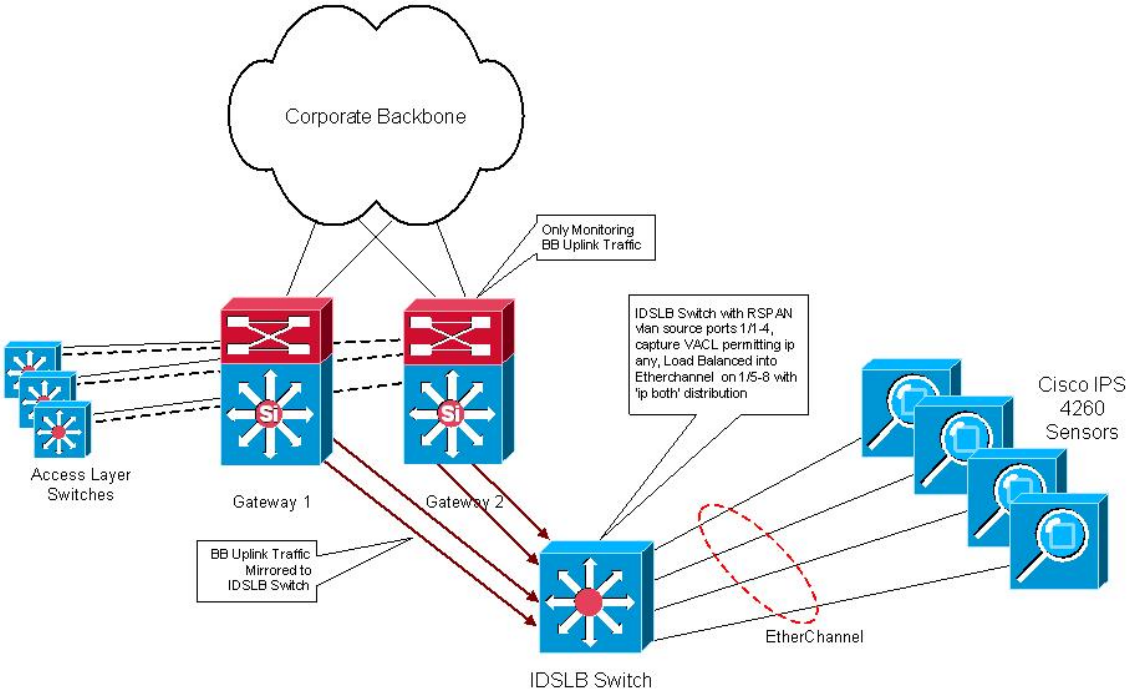

Cisco sites with less than 600 Mbps bandwidth use a single Cisco 4255 Sensor. These sites, which do not have onsite IT staff, include small Cisco data centers, engineering server rooms, and offshore development centers. Centralized CSIRT personnel manage the sensors remotely, using out-of-band access on the Cisco 4255 Sensor.

Figure 2 Small Data Center Design

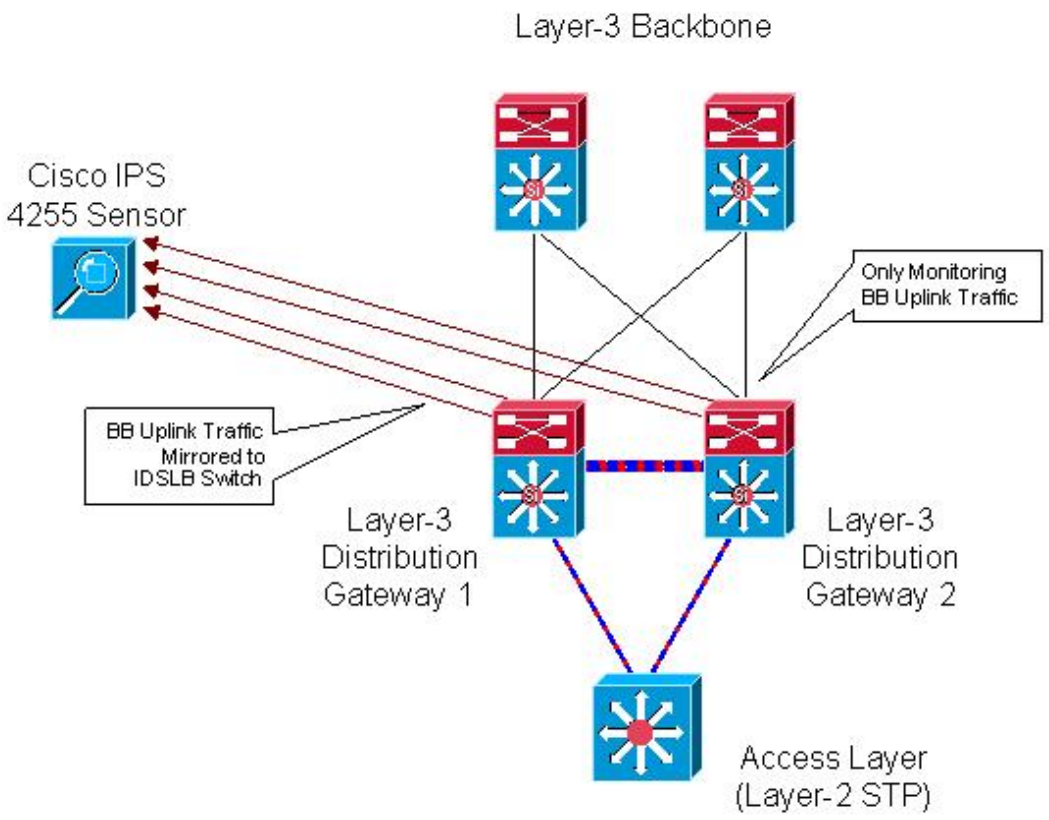

"Deploying the gear is easy," says Fry. "The real work, and the way we get value from a network-based IPS deployment, is the tuning."

Business Results

In 2004, the CSIRT team discovered 8 percent of security incidents; the remaining 92 percent were reported by Cisco users or external agencies. In 2007, the ratio had reversed: the CSIRT team discovered 97 percent of security incidents. The improvement is directly attributable to Cisco IPS. "Cisco IPS is now the primary source of network-security event data at Cisco," says Reid.

One success story for network-based IPS involves the Rinbot virus, which appeared in March 2007. On the day that the virus appeared, the Cisco IPS sensor generated an alert for a generic Windows buffer overflow attack. The team's investigation revealed that a new variant of virus was taking advantage of a recently published vulnerability in Microsoft Windows. "We analyzed the malware and discovered that it was looking up three different Internet domain names," says Fry. The CSIRT team deployed a custom signature that looked for any PC that performed a DNS lookup on one of the domains associated with the virus. "This created a high-fidelity signature telling us that a PC was infected," says Fry. "When we pushed out the signature to the Cisco IPS sensors, we identified a few hundred affected lab systems. Early detection and remediation prevented further damage."

Also in 2007, the monitoring team saw an IPS alert for a malware-related signature. They escalated it to the response team, which examined the signature and payload data contained in the alert to determine that the desktop was infected with a key logger. The team then logged onto the Cisco Security Agent management console and discovered that the employee had ignored an alert from Cisco Security Agent that the downloaded software might be malware, and downloaded it anyway. Armed with this knowledge, the team remediated the compromised client.

The network-based IPS system proves its value daily. Every week between March and June 2007, Cisco CSIRT found and mitigated ten new botnet command-and-control servers that were being used to remotely control systems within Cisco. “The ability to quickly find and mitigate malware has a significant business impact for Cisco,” says Fry. “We have become very proactive in mitigating network threats before we experience data loss or service interruption.”

PRODUCT LIST

Routing and Switching
- Cisco Catalyst 6504E Switch

Security and VPN
- Cisco 4260 Sensors (large data centers)
- Cisco 4255 Sensors (small data centers)
- Cisco Security Monitoring, Analysis, and Response System (MARS)

Lessons Learned

The Cisco CSIRT team offers the following suggestions for other organizations that deploy network-based IPS:

- Be aware that IPS relies on people as much as it does on technology. Time requirements include IPS analysis and monitoring, investigation, deployment and maintenance, and tuning. These functions can be performed by different groups. Cultivate good relationships with other IT teams because they can tell you which types of traffic are benign.
- If you suspect an event is a false positive, contact the application and operating system owners within the IT organization.
- Include on the team someone who has in-depth understanding of IP subnetting and internal networks. “Knowing your network is a large investment—and worth every dollar,” says Fry.
- Separate case investigation, which involves stopping malicious traffic and remediation, from event monitoring, which involves looking for malicious activity. The goal is to minimize monitoring time. For Cisco, the goal is twice-a-day monitoring or less.
- Plan for IT resource requirements. In general, the more open the security policy, the more IT resources are required for tuning and monitoring. Financial institutions and military organizations, for example, require less tuning than universities and enterprises such as Cisco. More personnel are required for monitoring at the beginning of the deployment, and resource requirements decrease after false positives have been reduced.
- IPS is a good reactive methodology. Keep in mind that it does not replace preventive methodologies like patching, antivirus, security ACL configuration, host-based IPS such as Cisco Security Agent, and other security in-depth practices.
- Identify the most important assets in order to decide where to deploy IPS sensors.

“Cisco IPS is our view into aberrant events,” says Reid. “It is the core element of a complete program, including processes and procedures for whom to contact, and remediation paths.”

For More Information

To read the entire case study or for additional Cisco IT case studies on a variety of business solutions, visit Cisco on Cisco: Inside Cisco IT www.cisco.com/go/ciscoit

To read more about Cisco IPS, visit: www.cisco.com/go/ips

Note

This publication describes how Cisco has benefited from the deployment of its own products. Many factors may have contributed to the results and benefits described; Cisco does not guarantee comparable results elsewhere.

CISCO PROVIDES THIS PUBLICATION AS IS WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING THE IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

Some jurisdictions do not allow disclaimer of express or implied warranties, therefore this disclaimer may not apply to you.

Americas Headquarters
Cisco Systems, Inc.
170 West Tasman Drive
San Jose, CA 95134-1706
USA
www.cisco.com
Tel: 408 526-4000
800 553-NETS (6387)
Fax: 408 527-0883

Asia Pacific Headquarters
Cisco Systems (USA) Pte. Ltd.
168 Robinson Road
#28-01 Capital Tower
Singapore 068912
www.cisco.com
Tel: +65 6317 7777
Fax: +65 6317 7799

Europe Headquarters
Cisco Systems International BV
Haarlerbergpark
Haarlerbergweg 13-19
1101 CH Amsterdam
The Netherlands
www-europe.cisco.com
Tel: +31 0 800 020 0791
Fax: +31 0 20 357 1100

Cisco has more than 200 offices worldwide. Addresses, phone numbers, and fax numbers are listed on the Cisco Website at **www.cisco.com/go/offices.**

Printed in USA

C00-000000-00 08/07