CHAPTER 3

Identifying and Classifying Security Threats

Worms and denial of service (DoS) attacks are used maliciously to consume the resources of your hosts and network that would otherwise be used to serve legitimate users. In some cases, misconfigured hosts and servers can send traffic that consumes network resources unnecessarily. Having the necessary tools and mechanisms to identify and classify security threats and anomalies in the network is crucial. This chapter presents several best practices and methodologies you can use to successfully and quickly identify and classify such threats.

Most people classify security attacks into two separate categories: *logic attacks* and *resource attacks*. Logic attacks exploit existing software deficiencies and vulnerabilities to cause systems to crash, to substantially degrade their performance, or to enable attackers to gain access to a system. An example of this type of attack is the exploit of the Microsoft PnP MS05-039 Overflow Vulnerability, in which the attacker exploits a stack overflow in the Windows "plug and play" (PnP) service. You can exploit this vulnerability on Windows 2000 without a valid user account. Another example is the famous and old *ping-of-death*, whereby an attacker sends the system Internet Control Message Protocol (ICMP) packets that exceed the maximum legal length (65535 octets). You can prevent most of these attacks by either upgrading the vulnerable software or by filtering particular packet sequences.

The second category of attacks is referred to as resource attacks. The goal with these types of attacks is to overwhelm the victim system/network resources, such as CPU and memory. In most cases, this is done by sending numerous IP packets or forged requests. An attacker can build up a more powerful attack with a more sophisticated and effective method of compromising multiple hosts and installing small attack daemon(s). This is what many call *zombies* or *bot* hosts/nets. Subsequently, an attacker can launch a coordinated attack from thousands of zombies onto a single victim. This daemon typically contains both the code for sourcing a variety of attacks and some basic communications infrastructure to allow for remote control. A zombie attack is illustrated in Figure 3-1.

Figure 3-1 *Zombies and Bots*

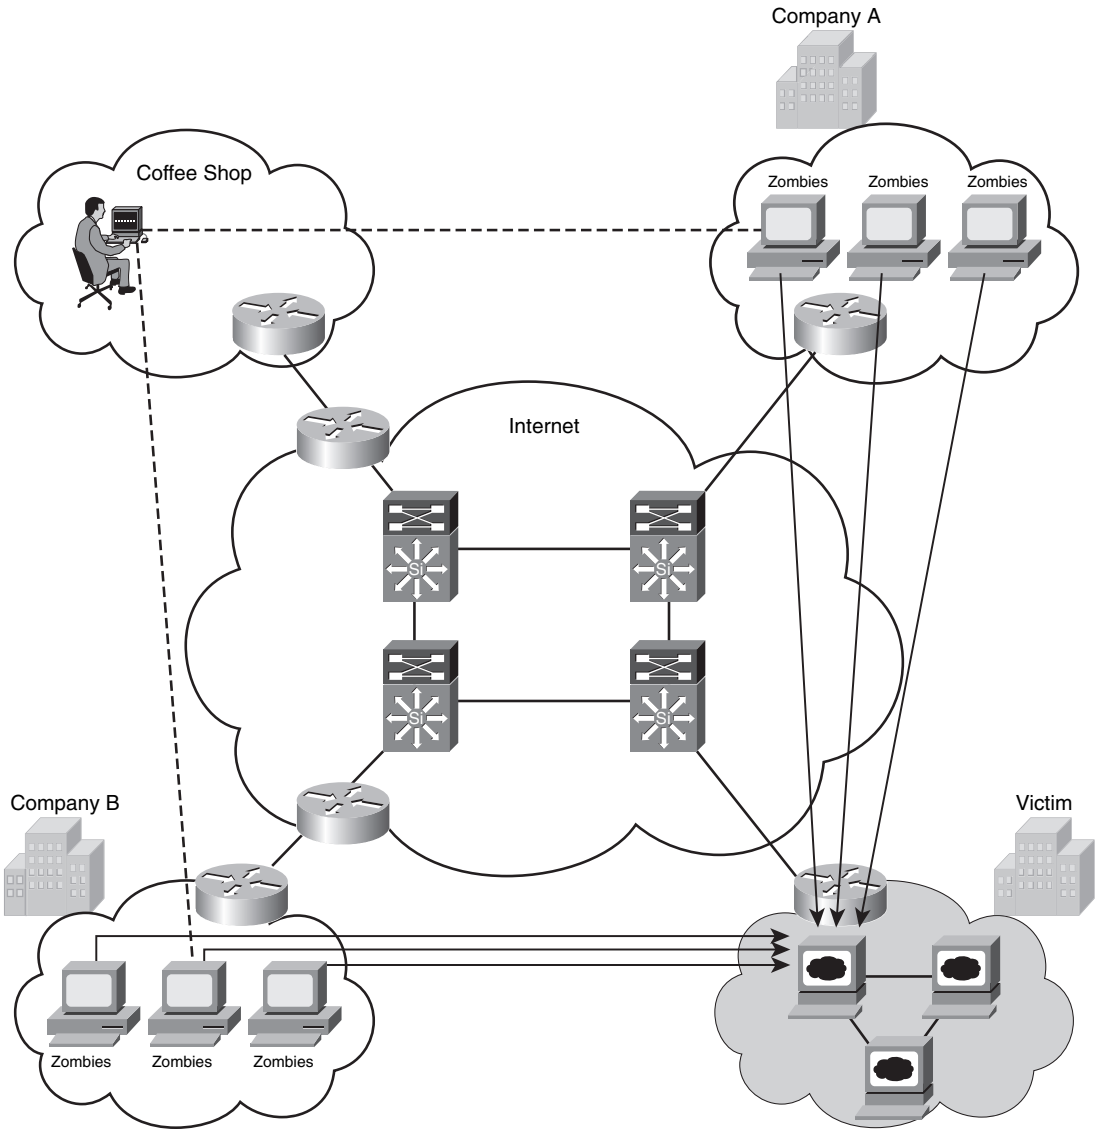

In Figure 3-1, an attacker controls compromised hosts in Company A and Company B to attack a web server farm in another organization.

You can use different mechanisms and methodologies to successfully identify and classify these threats/attacks depending on their type. In other words, depending on the threat, you can use specific techniques to identify and classify them accordingly. Following are the most common methodologies:

- The use of anomaly detection tools
- Network telemetry using flow-based analysis
- The use of intrusion detection and intrusion prevention systems (IDS/IPS)
- Analyzing network component logs (that is, SYSLOG from different network devices, accounting records, application logs, Simple Network Management Protocol (SNMP), and so on)

Complete visibility is one of the key requirements when identifying and classifying security threats. The following sections explain best practices for achieving complete network visibility and the use of the previously mentioned tools and mechanisms.

Network Visibility

The first step in the process of preparing your network and staff to successfully identify security threats is achieving complete network visibility. You cannot protect against or mitigate what you cannot view/detect. You can achieve this level of network visibility through existing features on network devices you already have and on devices whose potential you do not even realize. In addition, you should create strategic network diagrams to clearly illustrate your packet flows and where, within the network, you may enable security mechanisms to identify, classify, and mitigate the threat. Remember that network security is a constant war. When defending against the enemy, you must know your own territory and implement defense mechanisms in place. Figure 3-2 illustrates a fairly simple high-level enterprise diagram.

Figure 3-2 *High-Level Enterprise Diagram*

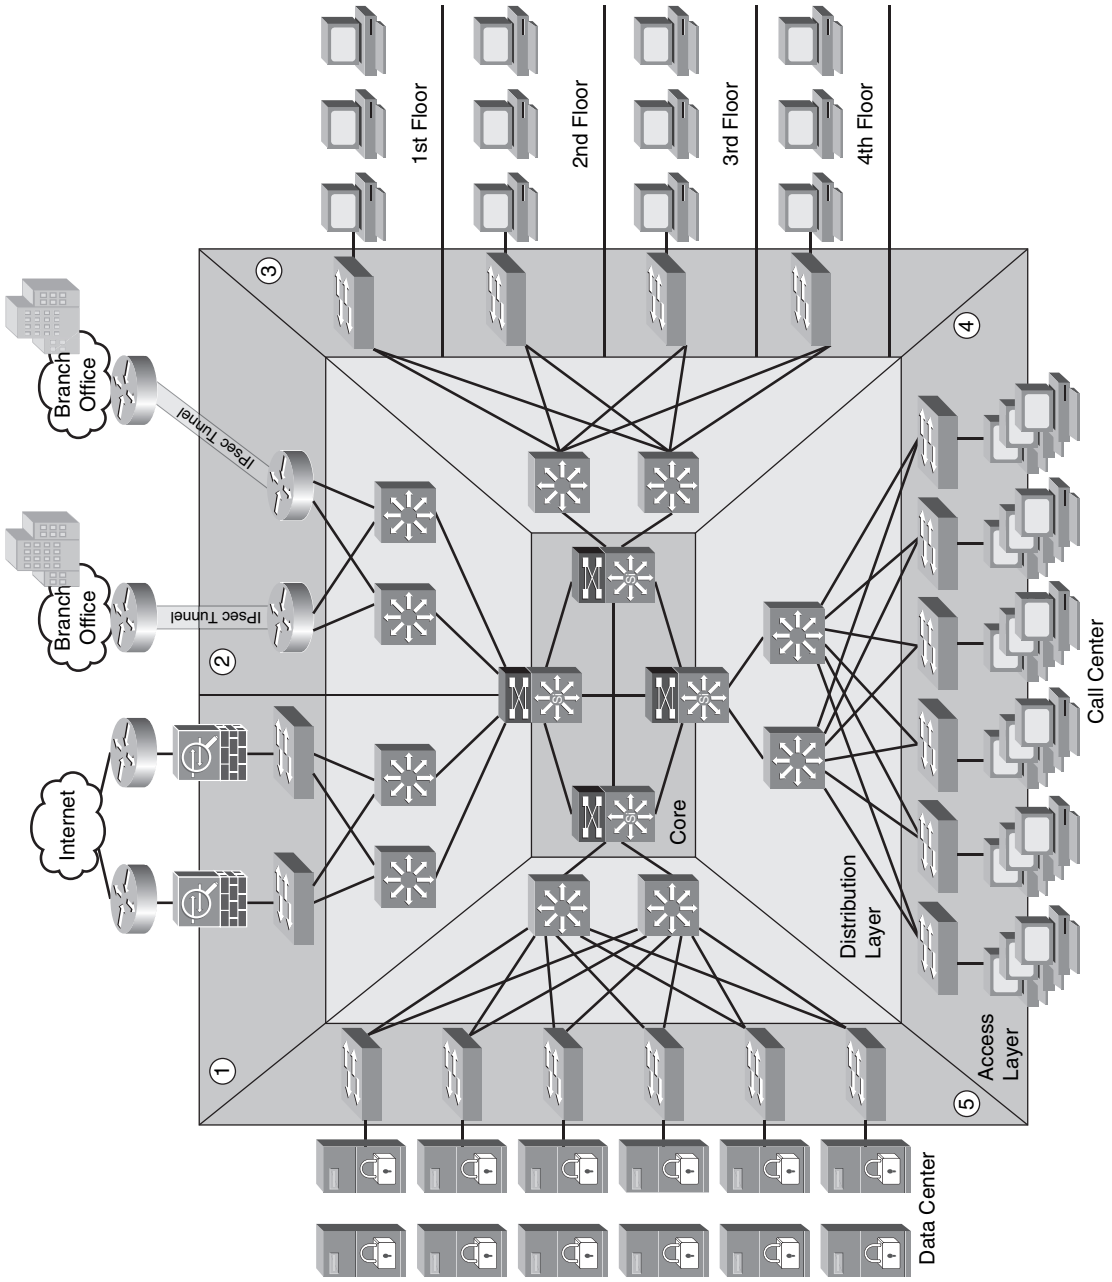

In Figure 3-2, the following sections are numbered:

1. **The Internet edge:** In this example, the enterprise headquarters is connected to the Internet via redundant links. Two Cisco Adaptive Security Appliances (ASA) are configured to protect the infrastructure.
2. **Site-to-Site VPN:** The headquarters office is connected to two branches via IPsec site-to-site VPN tunnels terminated on two Cisco IOS routers.
3. **End users:** The headquarters building has its sales, finance, engineering, and marketing departments on four separate floors.
4. **Call center:** There is a call center with more than 100 agents on the 5th floor.
5. **Data center:** The data center includes e-commerce, e-mail, database, and other application servers.

You can create this type of diagram not only to understand the architecture of your organization but also to strategically identify places within the infrastructure where you can implement telemetry mechanisms like NetFlow and identify choke points where you can mitigate an incident. Notice that the access, distribution, and core layers/boundaries are clearly defined.

Look at the example illustrated in Figure 3-3. A workstation at the call center usually communicates over TCP port 80 (HTTP) to a server in the data center. This traffic is allowed within the access control lists because it is legitimate traffic to the server. However, the traffic from this specific workstation increased more than 400 percent over normal. Subsequently, performance on the server is degraded, and the infrastructure is congested with unnecessary packets.

In this case, NetFlow was configured at the distribution layer switch, and the administrator was able to detect the anomaly. The administrator then configures a host-specific ACL to deny the traffic from the call center workstation, as shown in Figure 3-4. In more sophisticated environments, you can even implement remotely triggered black hole (RTBH) routing to mitigate this incident.

In the example illustrated in Figure 3-4, the problem was a defect within the call center workstation application. The administrator was able to perform detailed analysis and patch the machine while preventing disruption of service.

Figure 3-3 *NetFlow at the Distribution Switch*

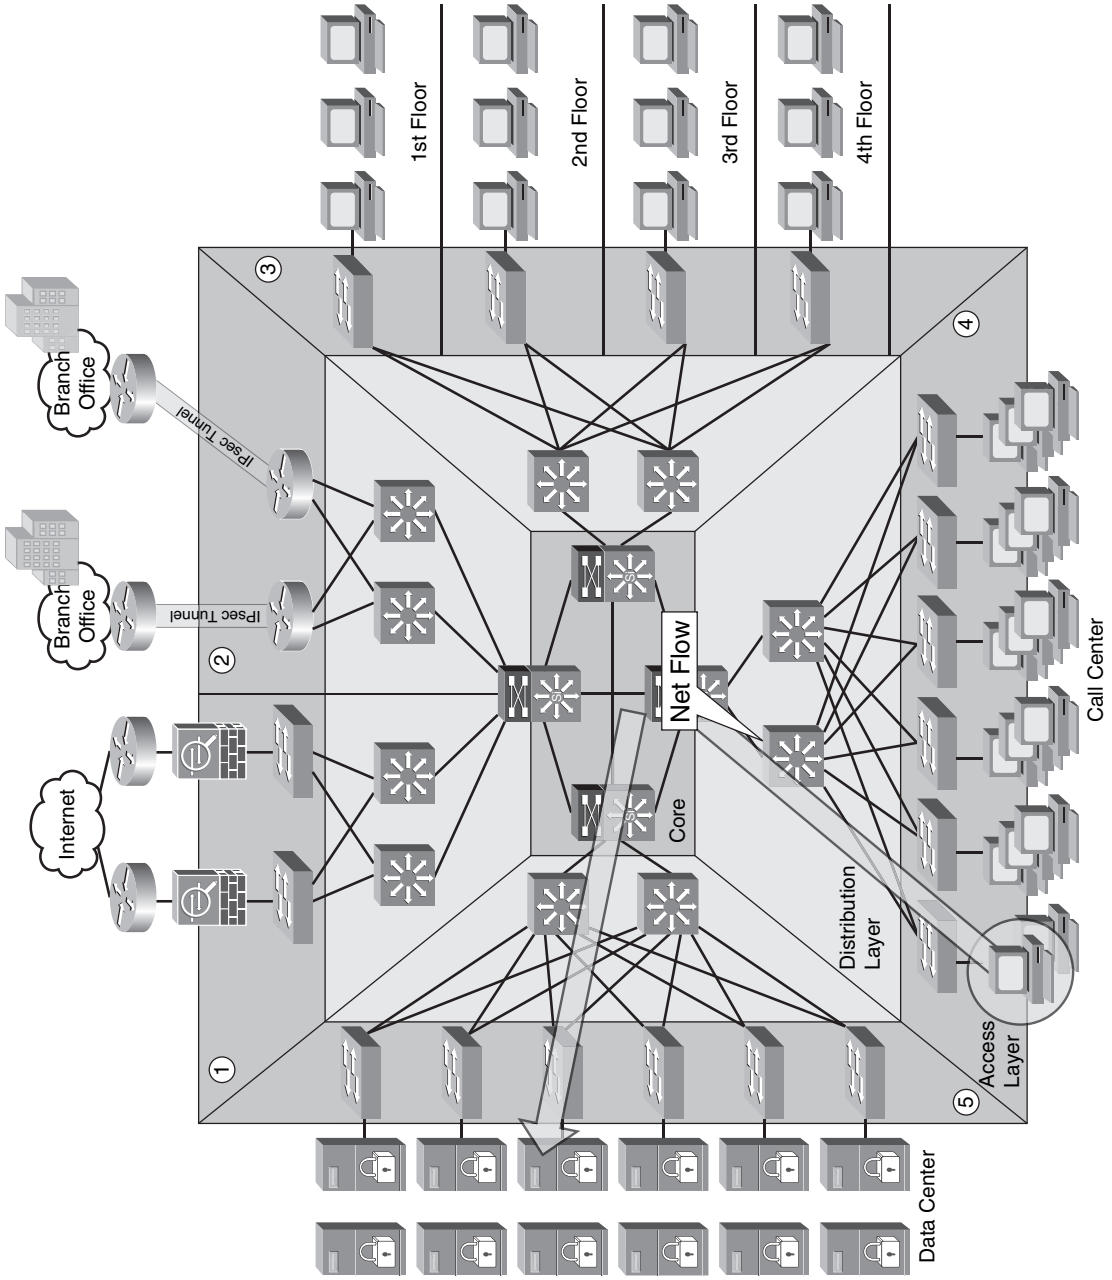

Figure 3-4 *Abnormal Traffic Stopped*

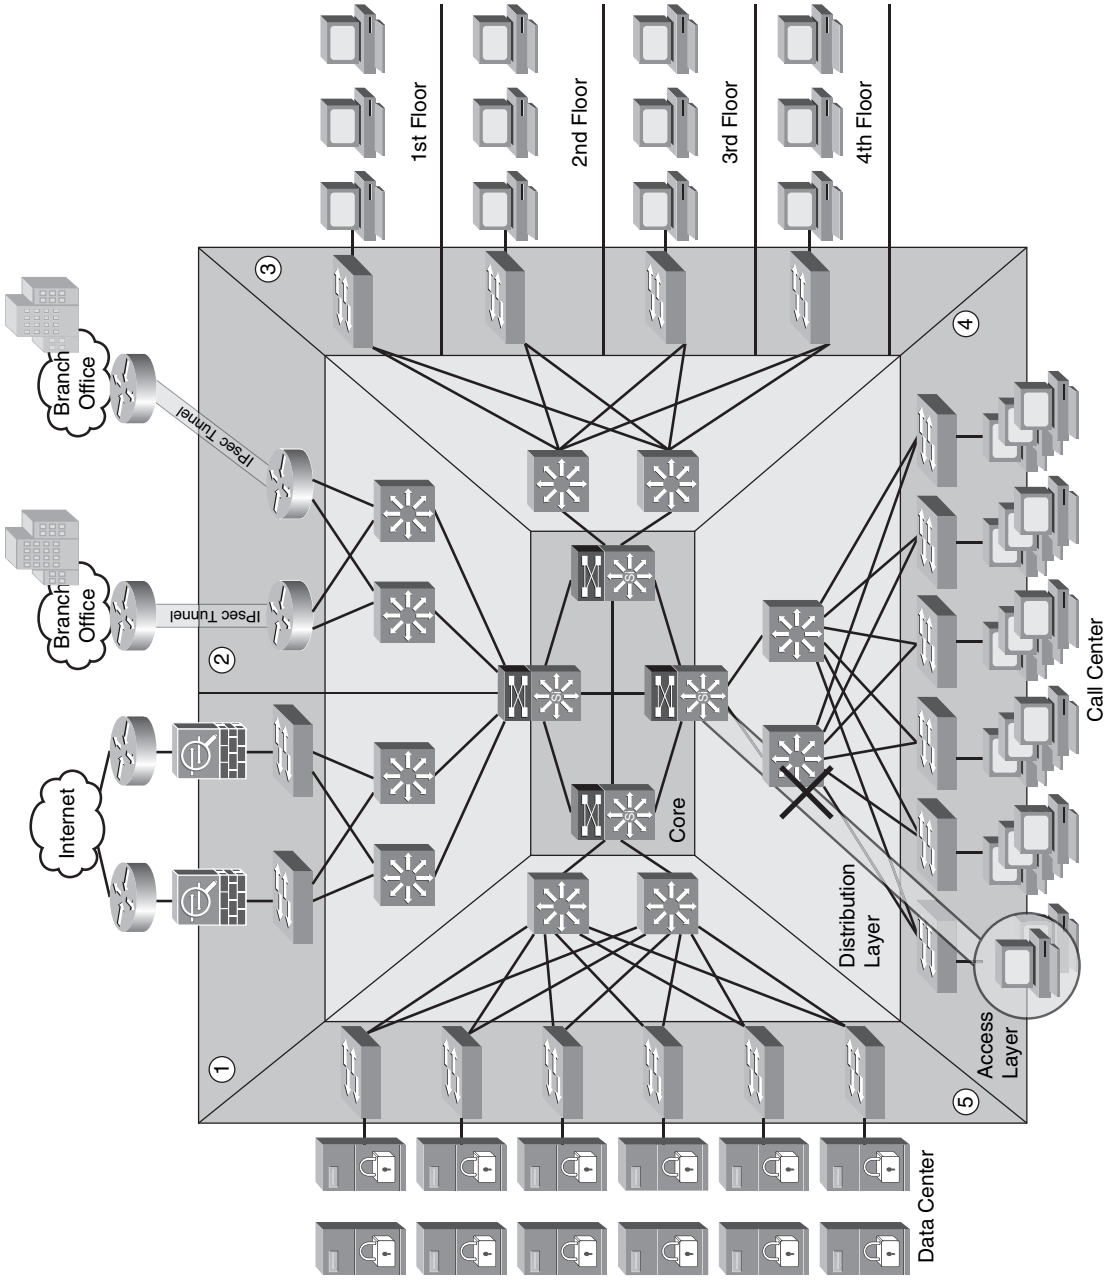

TIP To detect abnormal and possibly malicious activity, you must first establish a baseline of normal network activity, traffic patterns, and other factors. NetFlow, as well as other mechanisms, can be enabled within your infrastructure to successfully identify and classify threats and anomalies. Prior to implementing an anomaly-detection system, you should perform traffic analysis to gain an understanding of general traffic rates and patterns. In anomaly detection systems, learning is generally performed over a significant interval, including both the peaks and valleys of network activity. Anomaly detection and telemetry are covered in detail later in this chapter.

You can also develop a different type of diagram to visualize operational risks within your organization. These diagrams are based on device roles and can be developed for critical systems you want to protect. For example, identify a critical system within your organization and create a layered diagram similar to the one in Figure 3-5. In this example, a database called ABC is the most critical application/data source for this company. The diagram presents ABC Database Server in the center.

Figure 3-5 *Layered Diagram for Visualizing Risk*

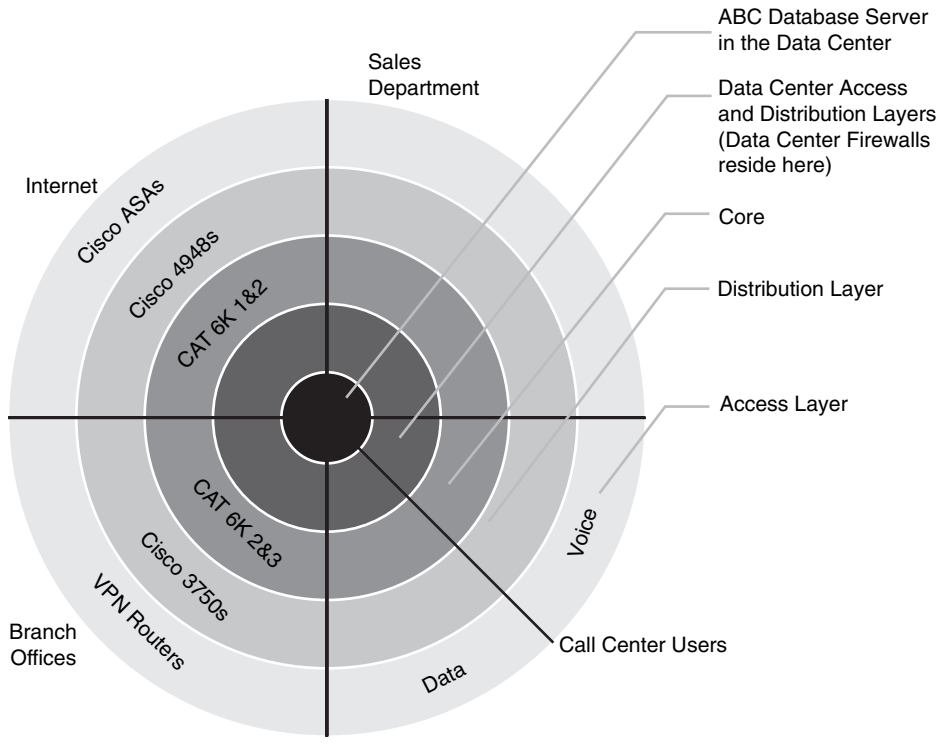

You can use this type of diagram to audit device roles and the type of services they should be running. For example, you can decide in what devices you can run services like Cisco NetFlow or where to enforce security policies. In addition, you can see the life of a packet within your infrastructure depending on the source and destination. An example is illustrated in Figure 3-6.

Figure 3-6 *Illustrating a Packet Flow*

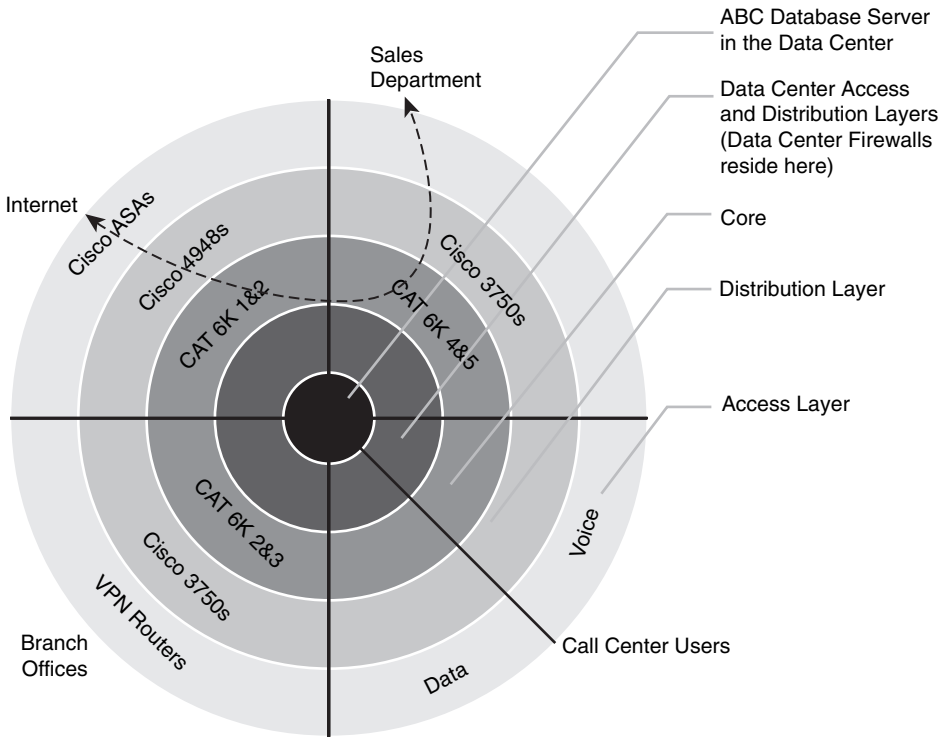

Figure 3-6 shows the packet flow that occurs when a user from the sales department accesses an Internet site. You know exactly where the packet is going based on your architecture and your security and routing policies. This is a simple example; however, you can use this concept to visualize risks and to prepare your isolation policies.

NOTE Additional examples and techniques are covered in Chapter 7, "Proactive Security Framework."

Telemetry and Anomaly Detection

Anomaly detection systems passively monitor network traffic, looking for any deviation from "normal" or "baseline" behavior that may indicate a security threat or a misconfiguration. You can use several commercial tools and even open source tools to successfully identify security threats within your network. These tools include the following:

- Cisco NetFlow
- Cisco Security Monitoring, Analysis and Response System (CS-MARS)
- Cisco Traffic Anomaly Detectors and Cisco Guard DDoS Mitigation Appliances
- Cisco IPS sensors (Version 6.x and later)
- Cisco Network Analysis Module (NAM)
- Open Source Monitoring tools

The following are other technologies and tools you can use to achieve complete visibility of what is happening within your network:

- Syslog
- SNMP

NetFlow

Cisco NetFlow was initially introduced as a packet accounting system for network administration and, in some cases, for billing. However, today you can use NetFlow to listen to the network itself, thereby gaining valuable insight into the overall security state of the network. This is why it is classified as a form of telemetry that provides information about traffic passing through or directly to each router or switch.

NetFlow is supported in the following Cisco platforms:

- Cisco 1700
- Cisco 1800
- Cisco 2800
- Cisco 3800
- Cisco 4500
- Cisco 7200
- Cisco 7300
- Cisco 7500

- Cisco 7600/6500 (hybrid and native configurations)
- Cisco 10000
- Cisco 12000

NOTE Indicated models have platform-specific considerations. Please refer to http://www.cisco.com/go/netflow for more compatibility information.

The word *netflow* is a combination of *net* (or network) and *flow*. What is a *flow*? An individual flow comprises, at a minimum, the following elements:

- Source IP address.
- Destination IP address.
- Protocol.
- Source port number. (With certain protocols, this can be a type/code or any other construct—for example, ICMP.)
- Destination port number. (With certain protocols, this can be a type/code or any other construct—for example, ICMP.)

NetFlow also can give you information about network traffic. This information varies somewhat depending on what version of NetFlow Data Export (NDE) you run. The most commonly deployed versions are Versions 5 and 9. Following is some of the additional information you can obtain from a flow in NetFlow Version 5:

- Start time of the flow.
- End time of the flow.
- Number of packets in the flow.
- Amount of data transferred in the flow.
- Type of Service (ToS) bits present in the flow or Differentiated Services Code Point (DSCP) type.
- Logical OR of all TCP flags present in TCP-based flows (platform-specific caveats apply).
- Input interface ifIndex.
- Output interface ifIndex.
- Origin-AS or destination-AS information, if Border Gateway Protocol (BGP) is enabled on the routers/Layer 3 switches in question. (The selection of origin- or destination-AS reporting is made during the configuration of NetFlow on each device.)

- BGP next-hop information, if BGP is enabled on the routers/Layer 3 switches in question.
- Fragmentation information (known as *fragmentation bit*).

All this information can be exported to monitoring systems for further analysis. NetFlow Version 9 supports the same reporting capabilities as NetFlow Version 5 with some additional information. One of the biggest advantages of NetFlow Version 9 is its ability to be configured by the use of templates to use various features to export additional or different information to external systems. In NetFlow Version 5 and earlier, you can export the flow data over UDP. NetFlow Version 9 supports NDE via TCP and SCTP, as well as the classic UDP mode.

NOTE All new NetFlow development is based on NetFlow Version 9.

In NetFlow Version 9, you can use a template describing the NDE fields within the flow information. This template information is contained in the first NetFlow Version 9 NDE packets sent to the NDE destination (monitoring system) after NDE is enabled on the router or switch. This information is also periodically retransmitted. When the configuration of NDE fields is changed on the router or switch, the updated template is immediately transmitted.

The IETF Internet Protocol Flow Information eXport (IPFIX) working group (WG) has been tasked with developing a common standard for IP-based flow export. This working group has selected Cisco NetFlow Version 9 as the technology of choice.

NOTE The IPFIX requirements are defined in RFC 3917. RFC 3954 explains the evaluation of NetFlow Version 9 in IPFIX. The actual outcome and the criteria for the selection of NetFlow Version 9 as the basis for the IPFIX standard are defined in RFC 3955.

It is recommended that you use an isolated out-of-band (OOB) management network to allow you to access and control NetFlow-enabled devices over the network, even when you are under attack or during any security incident or network malfunction. When you transmit network telemetry over the OOB network, you reduce the chance for disruption of the information that provides insightful network visibility.

Enabling NetFlow

Typically, enabling NetFlow on software-based platforms consists of one or two steps:

- Enabling NetFlow on the relevant physical and logical interfaces
- (Optional) Enabling the device (NDE) to export the flow information from the device to an external monitoring system

When you configure NetFlow, you must decide between ingress or egress NetFlow for each device. This decision depends on the use and the topology. You can also enable NetFlow for both ingress and egress.

NOTE Egress NetFlow is dependent on the version of Cisco IOS you are running. For more information, go to http://www.cisco.com/go/fn.

The following example shows how you can enable *ingress* NetFlow on a particular interface (GigabitEthernet0/0 in this case):

```
myrouter#configure terminal
myrouter(config)#interface GigabitEthernet0/0
myrouter(config-if)#ip flow ingress
```

To enable egress NetFlow, use the **ip flow egress** interface subcommand as follows:

```
myrouter(config)#interface GigabitEthernet0/0
myrouter(config-if)#ip flow egress
```

NOTE Ingress NetFlow is the most commonly used method. Egress NetFlow is more commonly used with MPLS VPN. The MPLS Egress NetFlow Accounting feature allows you to capture IP flow information for packets undergoing MPLS label disposition. In other words, it captures packets that arrive on a router as MPLS packets and are transmitted as IP packets. Egress NetFlow accounting might adversely affect network performance because of the additional accounting-related computations that occur in the traffic-forwarding path of the router.

The following example shows how to configure the NetFlow-enabled device to export the flow data to a monitoring system:

```
myrouter(config)#ip flow-export version 5
myrouter(config)#ip flow-export source loopback 0
myrouter(config)#ip flow-export destination 172.18.85.190 2055
```

In this example, NDE Version 5 is used. All NetFlow export packets are sourced from a loopback interface configured in the router (loopback 0). The destination is a Cisco Secure Monitoring and Response System (CS-MARS) box with the IP address 172.18.85.190 and the destination UDP port 2055.

It is recommended that you alter the setting of the active flow timeout parameter from its default of 30 minutes to the minimum value of one minute. This helps you achieve an environment that is closer to real time. You can do this with the **ip flow-cache timeout active** command, as shown here:

```
myrouter(config)#ip flow-cache timeout active 1
```

NOTE The default value for the number of minutes that an active flow remains in the cache before it times out is 30.

The default value for the number of seconds that an inactive flow remains in the cache before it times out is 15.

Collecting NetFlow Statistics from the CLI

To view the basic NetFlow information from the CLI, you can use the **show ip cache flow** command, as shown in Example 3-1:

Example 3-1 *Output of the* **show ip cache flow** *Command*

```
myrouter#show ip cache flow
IP packet size distribution (9257M total packets):
   1-32   64   96  128  160  192  224  256  288  320  352  384  416  448  480
   .088 .314 .011 .011 .027 .001 .007 .001 .013 .016 .002 .002 .000 .001 .000

    512  544  576 1024 1536 2048 2560 3072 3584 4096 4608
   .000 .001 .002 .043 .452 .000 .000 .000 .000 .000 .000

IP Flow Switching Cache, 4456704 bytes
  43 active, 65493 inactive, 884110623 added
  3341579080 ager polls, 0 flow alloc failures
  Active flows timeout in 30 minutes
  Inactive flows timeout in 15 seconds
  last clearing of statistics never
Protocol         Total    Flows   Packets Bytes  Packets Active(Sec) Idle(Sec)
--------         Flows     /Sec     /Flow  /Pkt     /Sec     /Flow     /Flow
TCP-Telnet     1072696      0.2        17   578      4.4       9.8      15.3
TCP-FTP          33386      0.0      2392    57     18.6     697.2       7.6
```

Example 3-1 *Output of the* **show ip cache flow** *Command (Continued)*

```
TCP-FTPD          2967      0.0      2869  1049      1.9       4.3      15.2
TCP-WWW        9091735      2.1       222   904    470.3       6.0       5.6
TCP-SMTP        538619      0.1         1    59      0.2       6.9      15.9
TCP-X             3246      0.0        44   909      0.0       0.1      13.4
TCP-BGP         280550      0.0         2    44      0.1       7.2      15.8
TCP-NNTP          2306      0.0         1    46      0.0       0.0      18.1
TCP-Frag             7      0.0        19   152      0.0       8.8      15.4
TCP-other     48037166     11.1       115   887   1289.2       4.5       6.2
UDP-DNS         1043579      0.2         2    74      0.4       3.9      15.9
UDP-NTP          891663      0.2         1    79      0.2       0.0      15.5
UDP-TFTP         138376      0.0         7    55      0.2      21.2      15.5
UDP-Frag           9736      0.0       182  1366      0.4      22.1      15.4
UDP-other     816395802    190.0         1   109    316.9       0.1      18.8
ICMP            6533952      1.5        13    95     20.5       8.3      15.5
GRE                 239      0.0        41    97      0.0      66.9      15.2
IP-other          34558      0.0      3907   156     31.4      66.1      15.0
Total:        884110583    205.8        10   750   2155.4       0.5      17.9
SrcIf         SrcIPaddress    DstIf         DstIPaddress    Pr SrcP DstP  Pkts
Fa1/1         14.38.1.9       Null          255.255.255.255 11 0044 0043     1
Fa1/1         0.0.0.0         Null          255.255.255.255 11 0044 0043   209
Fa0/0         172.18.173.68   Fa1/0         14.36.1.208     06 05BC 01BB   452
Fa0/0         172.18.173.68   Fa1/0         14.36.1.186     06 0631 01BB   388
Fa1/0         14.36.1.120     Null          14.36.255.255   11 008A 008A     3
Fa0/0         14.36.1.120     Null          14.36.255.255   11 008A 008A     3
Fa0/0         172.18.124.223  Fa1/0         14.36.197.213   06 8107 2323  1547
Fa0/0         172.18.124.66   Null          14.36.1.184     06 EC83 01BB     1
Fa1/0         14.36.8.48      Fa0/0         172.18.124.154  06 15FE 0FA5     1
Fa1/0         14.36.8.48      Fa0/0         172.18.124.154  06 15FF 0FA5     1
Fa1/0         14.36.8.48      Fa0/0         172.18.124.154  06 15FD 0FA5     1
Fa1/0         14.36.1.3       Fa0/0         172.18.123.69   01 0000 0303     3
Fa1/0         14.36.8.36      Fa0/0         172.18.124.66   11 0202 0202     4
Fa1/0         14.36.99.77     Fa0/0         172.18.124.225  06 01BB 137C    85
Fa1/0         14.36.197.213   Fa0/0         172.18.124.223  06 2323 8107   780
Fa0/0         172.18.124.223  Fa1/0         14.36.1.203     06 8105 2323  19992167
Fa0/0         172.18.85.169   Local         14.36.1.1       06 8E5E 0017    97
Fa0/0         172.18.124.225  Fa1/0         14.36.99.77     06 137C 01BB    85
Fa0/0         172.18.124.128  Fa1/0         14.36.1.128     06 916E 2323   138
Fa0/0         172.18.124.128  Fa1/0         14.36.1.128     06 916D 2323    54
Fa1/0         14.36.1.208     Fa0/0         172.18.173.68   06 01BB 05BC   678
```

In the highlighted line, you can see that a host (172.18.124.223 is sending 19,992,167 packets to 14.36.1.203. This may be abnormal behavior or an infected machine. The protocol is 06 (TCP), the source port is 33029 (Hex 8105), and the destination port is 8995 (Hex 2323).

You can also obtain export flow information using the **show ip flow export** command, as shown in Example 3-2:

Example 3-2 *Output of the* **show ip flow export** *Command*

```
myrouter#show ip flow export
Flow export v5 is enabled for main cache
  Exporting flows to 172.18.85.190 (2055)
  Exporting using source IP address 172.18.124.47
  Version 5 flow records
  884111088 flows exported in 31352026 udp datagrams
  0 flows failed due to lack of export packet
  4 export packets were sent up to process level
  0 export packets were dropped due to no fib
  0 export packets were dropped due to adjacency issues
  0 export packets were dropped due to fragmentation failures
  0 export packets were dropped due to encapsulation fixup failures
```

In Example 3-2, you can see that the router is exporting the NetFlow information to the 172.18.85.190 device (a CS-MARS in this case) over UDP port 2055. The source IP address is 172.18.124.47. A total of 884,111,088 flows have been exported in 31,352,026 UDP datagrams. Please note that all protocol numbers, source ports, and TCP/UDP destination ports are shown in hexadecimal. ICMP packets are represented with the source port field set to 0000, the first two bytes of the destination field set to the ICMP type, and the second two bytes to the ICMP code. If you are using features such as policy-based routing (PBR), Web Cache Communications Protocol (WCCP), Network Address Translation (NAT), or Unicast Reverse Path Forwarding (uRPF) ACLs, you will see a (DstIf) value of *Null*. To see packet drops caused by ACLs, uRPF, PBR, or null routes, use the **show ip cache flow** with the **include Null** option, as shown in Example 3-3:

Example 3-3 *Output of the* **show ip cache flow | include Null** *Command*

```
myrouter#show ip cache flow | include Null
Fa1/0         14.36.1.8       Null          255.255.255.255 11 0044 0043     1
Fa1/1         0.0.0.0         Null          255.255.255.255 11 0044 0043   891
Fa0/0         172.18.124.66   Null          14.36.1.184     06 80AC 01BB     3
Fa0/0         14.1.17.111     Null          14.38.201.1     06 51CD 00B3     2
Fa1/0         172.18.124.11   Null          172.18.124.255  11 0089 0089    18
Fa1/0         172.18.124.153  Null          172.18.124.255  11 008A 008A     3
```

To see flows that contain thousands or millions of packets, you can use **show ip cache flow | include K** or **show ip cache flow | include M** commands, respectively.

The Cisco Catalyst 6500 switches and Cisco 7600 router obtain NetFlow information via the Multilayer Switching (MLS) cache. In addition, the amount and type of data recorded in the table must be selected. The **mls flow ip interface-full** command provides the most useful information and can be configured as follows:

```
CAT6k(config)# mls flow ip interface-full
CAT6k(config)# mls nde interface
```

TIP If your NetFlow table has too many entries, you can try to reduce the MLS aging time. For PFC2, set the aging time high enough to keep the number of entries within the 32,000 flow range of the PFC2. For PFC3, set the aging time high enough to keep the number of entries within the 64,000 flow range of the PFC3.

Make sure you set the aging time to 1 second when using bridged-flow statistics with a Supervisor Engine 2 (SUP2). If some protocols have fewer packets per flow running, reduce the MLS fast aging time.

The following site includes detailed configuration and design information for NetFlow on Catalyst 6500 switches:

http://www.cisco.com/en/US/partner/products/hw/switches/ps708/products_configuration_guide_chapter09186a0080207758.html

SYSLOG

System logs or SYSLOG provide you with information for monitoring and troubleshooting devices within your infrastructure. In addition, they give you excellent visibility into what is happening within your network. You can enable SYSLOG on network devices such as routers, switches, firewalls, VPN devices, and others. This section covers how to enable SYSLOG on routers, switches, the Cisco ASA, and Cisco PIX security appliances.

Enabling Logging (SYSLOG) on Cisco IOS Routers and Switches

The logging facility on Cisco IOS routers and switches allows you to save SYSLOG messages locally or to a remote host. By default, routers send logging messages to a logging process. The logging process controls the delivery of logging messages to various destinations, such as the logging buffer, terminal lines, a SYSLOG server, or a monitoring event correlation system such as CS-MARS. You can set the severity level of the messages to control the type of messages displayed, in addition to a time stamp to successfully track the reported information.

TIP It is extremely important that your SYSLOG and other messages are time-stamped with the correct date and time. This is why the use of NTP is strongly recommended (*see the NTP example in Chapter 2, "Preparation Phase"*).

The following example shows the commands necessary to configure SYSLOG on Cisco IOS devices:

```
myrouter#configure terminal
myrouter(config)#logging on
myrouter(config)#logging host 172.18.85.190
```

In this example, the router is configured to send the SYSLOG messages to a host with IP address 172.18.85.190. (This is the CS-MARS used in the examples of the previous sections.)

On Cisco IOS routers, the log messages are not time-stamped by default. To enable time stamping of log messages, use the **service timestamps log datetime** command. The following example shows the different options of this command:

```
myrouter(config)#service timestamps log datetime ?
  localtime      Use local time zone for timestamps
  msec           Include milliseconds in timestamp
  show-timezone  Add time zone information to timestamp
  year           Include year in timestamp
```

You can specify the severity level of the SYSLOG messages. The following are the different levels you can configure:

- **Level 0:** Emergencies
- **Level 1:** Alerts
- **Level 2:** Critical
- **Level 3:** Errors
- **Level 4:** Warnings
- **Level 5:** Notifications
- **Level 6:** Informational
- **Level 7:** Debugging

To set the severity level of log messages sent to a SYSLOG server, use the **logging trap** command. The following example shows the options of this command:

```
myrouter(config)#logging trap ?
  <0-7>          Logging severity level
  alerts         Immediate action needed           (severity=1)
  critical       Critical conditions               (severity=2)
  debugging      Debugging messages                (severity=7)
  emergencies    System is unusable                (severity=0)
  errors         Error conditions                  (severity=3)
  informational  Informational messages            (severity=6)
  notifications  Normal but significant conditions (severity=5)
  warnings       Warning conditions                (severity=4)
```

It is recommended that you send SYSLOG messages over a separate management segment, just as you learned to do earlier in this chapter in the "NetFlow" section.

Enabling Logging Cisco Catalyst Switches Running CATOS

To enable the logging of system messages to a SYSLOG server on Cisco Catalyst switches running Catalyst Operating System (CATOS), use the following commands:

```
set logging server enable
set logging server syslog server 172.18.85.190
set logging timestamp enable
set logging server severity 4
```

In this example, the switch is configured to send the SYSLOG messages to the host with IP address 172.18.85.190. Time stamp is enabled, and the severity level of the messages sent to the external server is set to 4 or warnings. Setting logging to the debugging level can cause performance problems. A good rule of thumb is to set the logging severity to 4 or warnings.

NOTE A good whitepaper describing best practices when managing Cisco Catalyst switches running CATOS is located at http://www.cisco.com/en/US/products/hw/switches/ps663/products_tech_note09186a0080094713.shtml.

Enabling Logging on Cisco ASA and Cisco PIX Security Appliances

The commands used to enable logging and to send SYSLOG messages to a SYSLOG server are the same on the Cisco ASA and the Cisco PIX security appliances. To enable logging, use the **logging on** command. To configure the ASA or PIX to send logs to a SYSLOG server, use the **logging host** command, and to change the log severity level, use the **logging trap** command. The following example demonstrates the use of these commands.

```
ciscoasa(config)# logging on
ciscoasa(config)# logging host inside 172.18.85.190
ciscoasa(config)# logging trap informational
```

In this example, the Cisco ASA is configured to send its logs to the host with IP address 172.18.85.190, and the severity level is set to informational.

On the Cisco ASA and Cisco PIX security appliances, all SYSLOG messages begin with a percent sign (%) and are designed as follows:

```
%PIX|ASA  Level Message_number: Message_text
```

The following is an example of a SYSLOG message.

```
Apr 09 2007 07:35:56: %ASA-6-302021: Teardown ICMP connection for faddr
192.168.202.22/0 gaddr 192.168.202.40/0 laddr 192.168.202.40/0
```

- **PIX|ASA:** A static value indicating that the log message is generated by a Cisco ASA or Cisco PIX.
- **Level:** The severity level (1–7). For most environments, it is recommended that you set the severity level to 4 to avoid performance issues. You may want to temporally increase it to a higher value when troubleshooting a specific problem.
- **Message number:** A unique 6-digit number that identifies the SYSLOG message.
- **Message text:** The description of the log message. It sometimes includes IP addresses, port numbers, or usernames.

You can filter SYSLOG messages on the Cisco ASA, Cisco PIX, and Cisco FWSM to send only specific events to a particular output destination. In other words, you can configure the device to send all SYSLOG messages to one output destination and also to send a subset of those SYSLOG messages to a different output destination. You can also configure the Cisco ASA, Cisco PIX, and Cisco FWSM to send SYSLOG messages based on specific criteria, such as the following:

- Message ID number (range of 104024 to 105999)
- Severity level
- Message class

For example, you can use the **logging class <message_class>** command to specify the specific class.

TIP All Cisco ASA and Cisco PIX messages are defined in detail at http://www.cisco.com/univercd/cc/td/doc/product/multisec/asa_sw/v_7_2/syslog/logmsgs.htm.

This site also includes the different SYSLOG message classes and associated message ID numbers.

SNMP

SNMP is one of the most basic forms of getting information from your network. It is a Layer 7 protocol designed to obtain information from network devices. This information includes but is not limited to the following:

- Device health statistics (CPU, memory, and so on)
- Device errors
- Network traffic statistics
- Packet rates
- Packet errors

The SNMP solution has three components:

- **An SNMP manager:** The system used to control and monitor the activities of network hosts using SNMP.
- **An SNMP agent:** The software component within the managed device that maintains the data for the device and reports this data, as needed, to managing systems.
- **A Management Information Base (MIB):** An information storage medium that contains a collection of managed objects (MIB modules) within each device. MIB modules are written in the SNMP MIB module language, as defined in STD 58, RFC 2578, RFC 2579, and RFC 2580.

In Chapter 2, you learned about the three versions of SNMP and the security implications of each version. That chapter also showed you how to protect SNMP environments. This section covers the basic commands on how to enable SNMP on Cisco IOS and the Cisco ASA and Cisco PIX security appliances.

Enabling SNMP on Cisco IOS Devices

As a best practice, you should set the system contact, location, and serial number of the SNMP agent so that your management servers can obtain these descriptions. This information is useful when responding to incidents. The following example shows how to enter the contact information on the Cisco IOS device:

```
myrouter#configure terminal
myrouter(config)#snmp-server contact John Route
myrouter(config)#snmp-server location 1st Floor NY Office
myrouter(config)#snmp-server chassis-id ABC12345
```

In the previous example, the name of the administrator is John Route, the device is located on the 1st floor of an office in New York, and the chassis identification number is ABC12345.

The following example shows how you can configure SNMP Version 3 on a Cisco IOS device:

```
myrouter(config)#snmp-server group mygroup v3 auth
```

SNMP Version 3 supports authentication. In the previous example, an SNMP group named mygroup is configured for SNMP Version 3. Authentication is also enabled with the **auth** keyword. When you configure the **snmp-server group** command, there are no default values for authentication. To specify authentication user parameters, use the **snmp-server user** command, as shown in the following example:

```
myrouter(config)#snmp-server user admin1 mygroup v3 auth md5 zxasqw12
*Feb  8 15:45:04.902: Configuring snmpv3 USM user, persisting snmpEngineBoots.
Please Wait...
```

In the previous example, a user (*admin1*) is configured and mapped to the SNMP group *mygroup*. Authentication is done with MD5, and the password is *zxasqw12*. After you invoke this command, the preceding warning message is displayed. You should match all this information in your SNMP management server.

To verify the configuration, you can invoke the **show snmp user** command as follows:

```
myrouter#show snmp user
User name: admin1
Engine ID: 800000090300000013C4EC5528
storage-type: nonvolatile        active
Authentication Protocol: MD5
Privacy Protocol: DES
Group-name: mygroup
```

To view SNMP group information, invoke the **show snmp group** command, as shown in Example 3-4.

Example 3-4 *Output of the* **show snmp group** *Command*

```
myrouter#show snmp group
groupname: ILMI                               security model:v1
readview : *ilmi                              writeview: *ilmi
notifyview: <no notifyview specified>
row status: active
groupname: ILMI                               security model:v2c
readview : *ilmi                              writeview: *ilmi
notifyview: <no notifyview specified>
row status: active
groupname: mygroup                            security model:v3 auth
readview : v1default                          writeview: <no writeview specified>
notifyview: <no notifyview specified>
row status: active
```

The configured group (mygroup) is shown in the highlighted line.

NOTE The following site includes detailed information on how to configure SNMP Version 1 and 2:

http://www.cisco.com/univercd/cc/td/doc/product/software/ios124/124tcg/tnm_c/snmp/confsnmp.htm#wp1032846

This document also includes the following information:

- Configuring the router as an SNMP manager
- Enabling the SNMP Agent Shutdown mechanism
- Defining the maximum SNMP Agent packet size
- Disabling the SNMP Agent
- Limiting the number of Trivial File Transfer Protocol (TFTP) servers used via SNMP

- Configuring SNMP notifications
- Configuring interface index display and interface indexes and configuring long name support
- Configuring SNMP support for VPNs
- Configuring MIB persistence

Enabling SNMP on Cisco ASA and Cisco PIX Security Appliances

The Cisco ASA and the Cisco PIX security appliances support only SNMP Versions 1 and 2c. They both support traps and SNMP read access; however, SNMP write access is not supported. The following example shows how to configure an ASA to receive SNMP Version 2c requests from host 172.18.85.190 on the inside interface:

```
ciscoasa(config)# snmp-server host inside 172.18.85.190 Version 2c
ciscoasa(config)# snmp-server location Raleigh NC Branch
ciscoasa(config)# snmp-server contact Jeff Firewall
ciscoasa(config)# snmp-server community th1s1sacommstrng
```

The ASA in this example is located in a branch office in Raleigh, North Carolina. The point of contact is Jeff Firewall, and the community string is <th1s1sacommstrng>. You can use the **snmp deny version** command to deny SNMP packets from other SNMP versions. The following example shows the available options:

```
ciscoasa(config)# snmp deny version ?
configure mode commands/options:
  1   SNMP version 1
  2   SNMP version 2 (party based)
  2c  SNMP version 2c (community based)
  3   SNMP version 3
```

NOTE You can obtain the MIBs for any Cisco device at http://www.cisco.com/public/sw-center/netmgmt/cmtk/mibs.shtml.

Cisco Security Monitoring, Analysis and Response System (CS-MARS)

CS-MARS enables you to identify, classify, validate, and mitigate security threats. In the previous sections in this chapter, you learned different mechanisms that give you visibility of the network and its devices, such as NetFlow, SYSLOGs, and SNMP. The analysis and manipulation of the data provided by these features can be a time-consuming process and, in some environments, may even be impossible because of the staff requirements.

CS-MARS supports the correlation of events from numerous networking devices from different vendors. The supported devices include:

- Cisco IOS routers and switches
- Cisco ASA
- Cisco PIX
- NetFlow
- Cisco Security Agent
- Cisco Secure ACS
- Cisco IDS/IPS
- Third-party firewalls such as Checkpoint and Netscreen
- Third-party antivirus software
- Third-party IDS/IPS systems such as snort
- Operating system (Windows and UNIX/Linux) events
- Application-specific events

NOTE A complete of list of supported devices can be found at http://www.cisco.com/en/US/products/ps6241/products_device_support_tables_list.html.

For a complete list of available CS-MARS models, go to http://www.cisco.com/go/mars.

CS-MARS provides a powerful and interactive dashboard with several key items. It includes a topology map that comprises real-time hotspots, incidents, attack paths, and detailed investigation with full incident disclosure, allowing immediate verification of valid threats. Figure 3-7 shows the CS-MARS main dashboard.

Note that the system has processed more than 22,000,000 NetFlow events (or flows) over a period of 24 hours, and more than 44,000,000 security and network events. This automated process is accomplished by analyzing device logs such as firewalls and by using intrusion prevention applications, third-party vulnerability assessment data, and Cisco Security MARS endpoint scans to eliminate false positives. Users can quickly fine-tune the system to further reduce false positives. This will be impossible to successfully analyze without the use of a system such as CS-MARS.

Figure 3-8 shows the bottom part of the CS-MARS main dashboard. There you can see a topology map of devices within the network, an attack diagram, and event statistics and graphs.

Figure 3-7 *CS-MARS Main Dashboard*

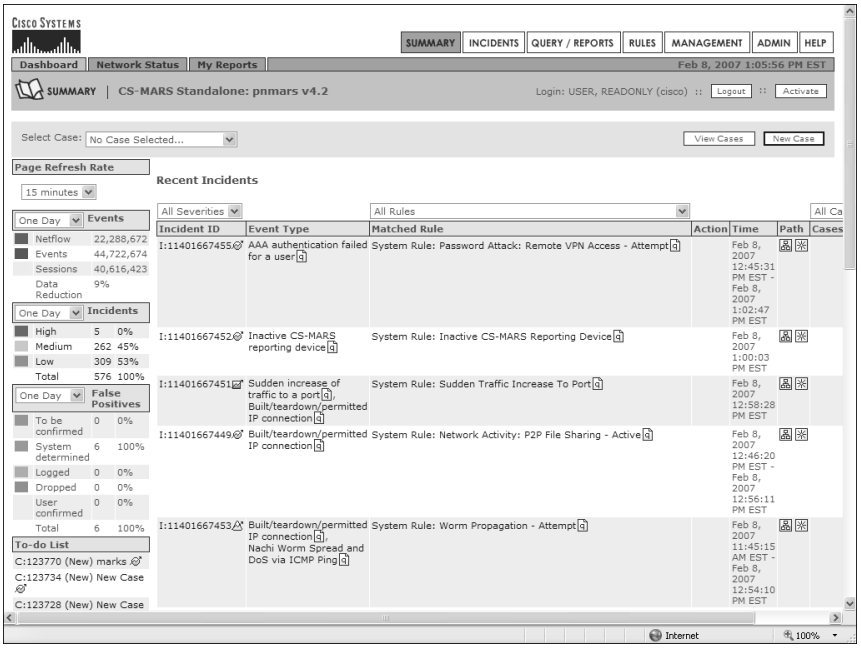

Figure 3-8 *CS-MARS Topology Map, Attack Diagram, and Event Statistics*

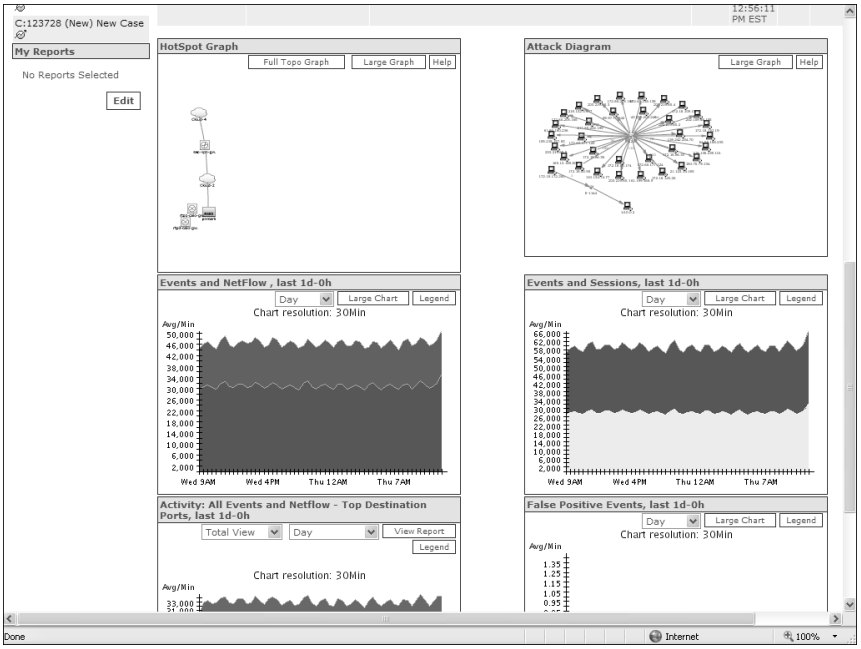

You can view the topology map and attack diagram in full view, as shown in Figure 3-9. Obtaining information about the security incident is simple. If you click on any of the arrows representing the traffic flow, a new window displays with information about the specific incident or session.

Figure 3-9 *CS-MARS Attack Diagram Full View*

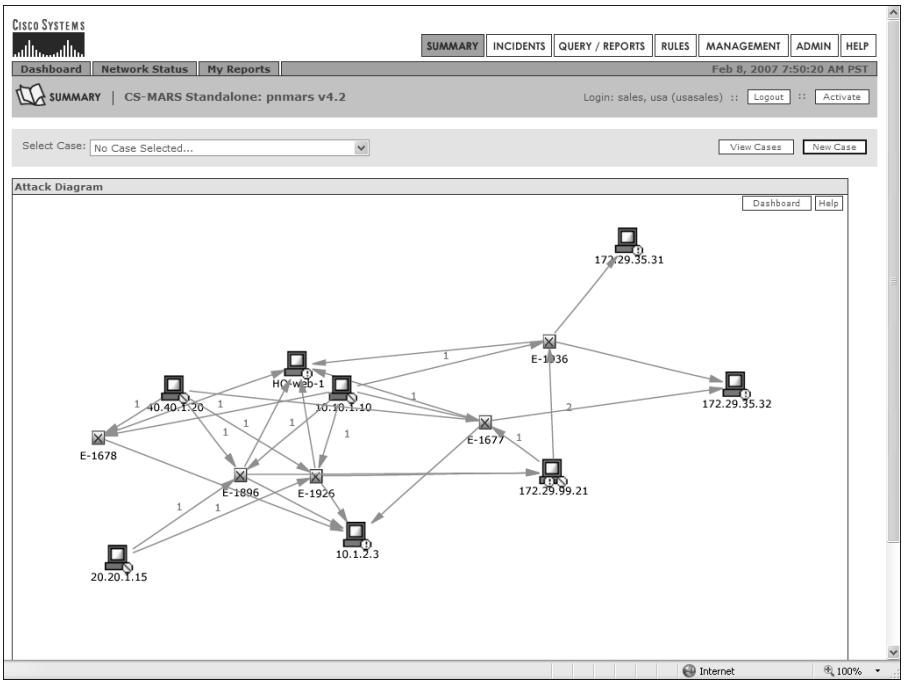

The hosts are color-coded:

- Brown means that the host is the attacker.
- Red means that the host is being attacked.
- Purple means that the host is being attacked and is attacking other hosts in the network.

CS-MARS can do a reverse DNS lookup to give you exact information on the specific hosts and devices. You can run numerous reports in CS-MARS. Figure 3-10 shows an example of reports and graphics you can obtain in CS-MARS.

Figure 3-10 *CS-MARS Detailed Graphics and Reports*

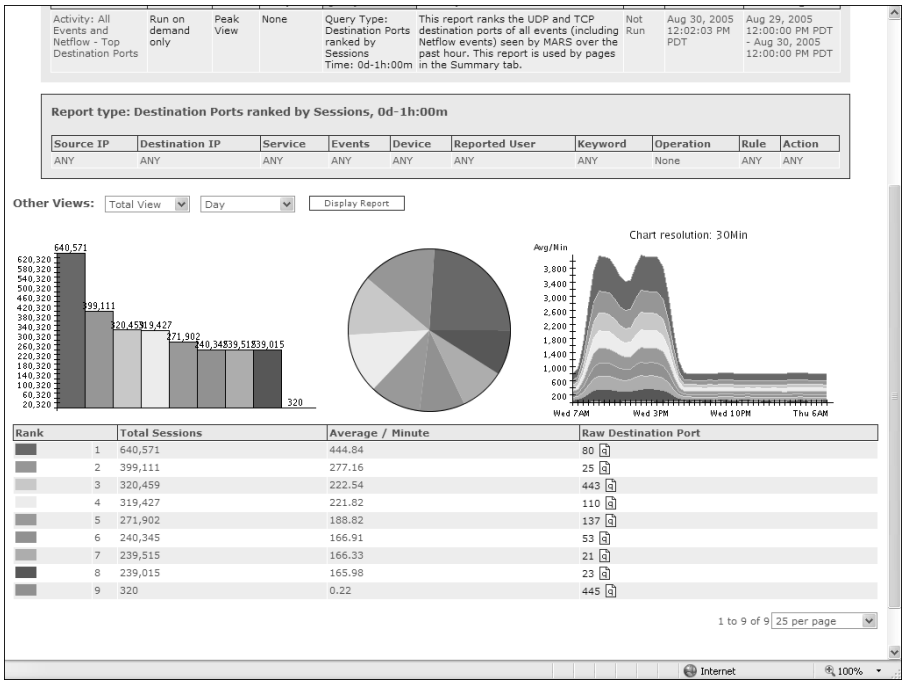

In Figure 3-10, you can see a summary of the most used ports and protocols within a given period. These graphics are based on NetFlow information. The graphic on the right shows the traffic trend. Notice that the traffic starts increasing during normal business hours of 8:00 a.m. to around 5:00 p.m. (0800 to 1700). These types of graphics can help you to create a baseline of what is normal within your network. Then you can identify anomalies and possible security incidents.

NOTE Chapter 12, "Case Studies," includes a case study in which CS-MARS is used to successfully identify, classify, and mitigate an attack. It also includes examples of how to add monitored devices into CS-MARS.

Cisco Network Analysis Module (NAM)

The Cisco Network Analysis Module (NAM) is designed to analyze and monitor traffic in the Catalyst 6500 series switches and Cisco 7600 series Internet routers. It uses remote monitoring (RMON), RMON extensions for switched networks (SMON), and SNMP MIBs to obtain information from the device. The NAM can also collect and analyze NetFlow information on remote devices.

To use the NAM to collect NetFlow data from a remote device, you must configure the remote device to export NDE packets to UDP port 3000 on the NAM. By default, the local supervisor engine of the switch is always available as an NDE device. Optionally, SNMP community strings are used to upload convenient textual strings for interfaces on the remote devices that are monitored in NetFlow records.

NOTE A complete NAM installation and configuration guide is located at http://www.cisco.com/en/US/products/sw/cscowork/ps5401/products_installation_and_configuration_guides_list.html.

Open Source Monitoring Tools

You can use several open source monitoring tools in conjunction with NetFlow. If your organization is small, or if you do not have the budget for more sophisticated monitoring tools, you can take advantage of any of these open source tools that are freely available. Table 3-1 includes the most commonly used open source monitoring tools.

Table 3-1 *Open Source Monitoring Tools*

Tool Name	Website
Caida's Cflowd Analysis Software	http://www.caida.org/tools/measurement/cflowd
My Netflow Reporting System by Dynamic Networks	http://www.dynamicnetworks.us/netflow/index.html
OSU Flow-tools	http://www.splintered.net/sw/flow-tools
Flow Viewer	http://ensight.eos.nasa.gov/FlowViewer
Flowd	http://www.mindrot.org/projects/flowd
NetFlow Monitor (NF)	http://netflow.cesnet.cz
Ntop	http://ntop.ethereal.com/ntop.html
Panoptis	http://panoptis.sourceforge.net
Plixer's Scrutinizer	http://www.plixer.com/products/free-netflow.php
Stager	http://software.uninett.no/stager

Most of these tools are designed to run in common *NIX-type operating systems, including Linux, FreeBSD, Mac OS/X, and Solaris. Some of these tools support the storage of data

in databases such as MySQL and Oracle. Despite the fact that these open source tools are free, they are extremely useful for collecting NetFlow from routers and storing the raw flows for auditing and forensic purposes. The most commonly used tool is the OSU flow-tool, which is typically used in conjunction with other packages that provide detailed graphs, charts, and on-demand queries. Visit each of the websites listed in Table 3-1 to learn more about which tool is most suitable for your environment.

Cisco Traffic Anomaly Detectors and Cisco Guard DDoS Mitigation Appliances

The Cisco traffic anomaly detectors and DDoS mitigation appliances provide a new approach that not only detects increasingly complex and unrepresentative denial of service attacks but also mitigates their effect to ensure business continuity and resource availability. The Cisco DDos solution has two distinct appliances:

- Cisco Traffic Anomaly Detector (TAD) XT
- Cisco Guard XT

This solution is also available in the form of two individual modules for the Catalyst 6500 series switches and the Cisco 7600 Internet routers:

- Catalyst 6500/Cisco 7600 Router Anomaly Guard Module
- Catalyst 6500/Cisco 7600 Router Traffic Anomaly Detector Module

The detectors (whether the appliances or the modules) are designed to promiscuously monitor network traffic while looking for any variation from what is "normal," which may indicate a DDoS attack or a worm outbreak. The Cisco TAD XT alerts the Cisco Guard XT when it detects an anomaly by providing detailed reports and specific alerts.

This solution uses a Multiverification Process (MVP) architecture integrating different verification, analysis, and enforcement techniques. The MVP has five components:

- Static and dynamic DDoS filters
- Active verification (anti-spoofing) implementing source-authentication mechanisms that help ensure proper identification of legitimate traffic
- Anomaly recognition
- Protocol analysis designed to identify Layer 7 attacks, such as HTTP error attacks
- Rate limiting that prevents flows from overwhelming the target while more detailed monitoring is taking place

Figure 3-11 illustrates how the Cisco TAD XT and the Cisco Guard XT work.

Figure 3-11 *Cisco TAD XT Detects an Anomaly and Updates the Guard XT*

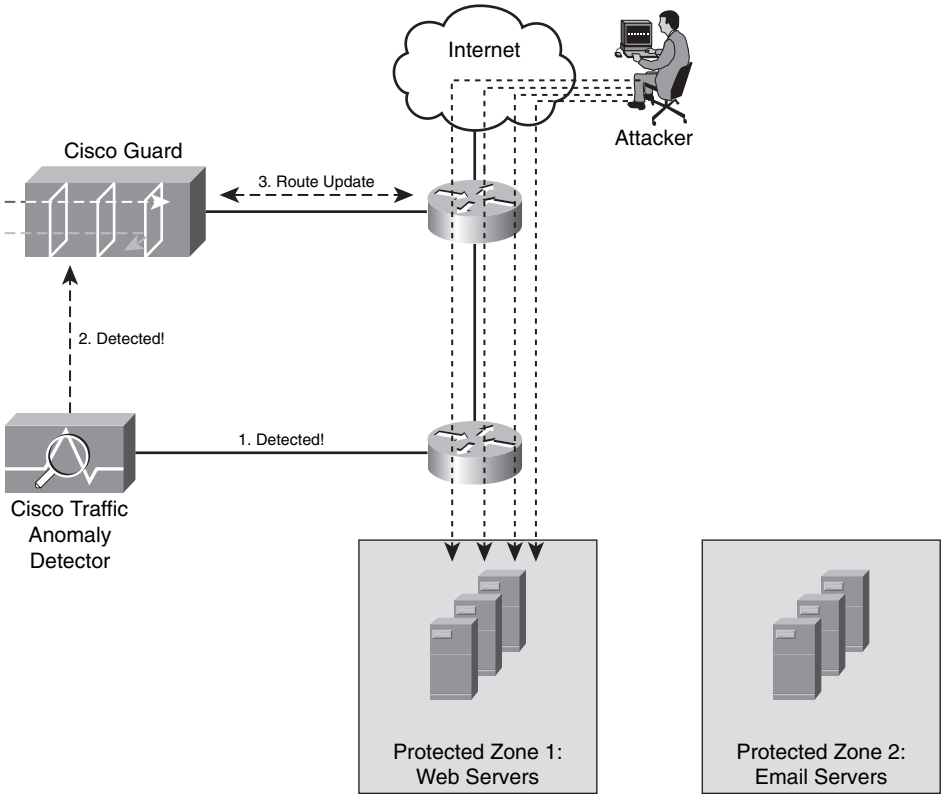

In Figure 3-11, two zones are protected by the Cisco TAD XT: a web server farm and an e-mail server farm. The Cisco Guard is placed at the Internet edge, and the Cisco TAD XT resides a couple of hops in the inside of the corporate network. The following are the steps illustrated in Figure 3-11.

Step 1 An attacker starts a DDoS from the Internet, and the Cisco TAD XT detects the anomaly (spike of traffic).

Step 2 The Cisco TAD XT updates the Cisco Guard XT. The Cisco Guard XT can be triggered in several ways:

- Through direct use of the web-based device manager
- Via the CLI
- Through automatic use of the "protect by packet" feature (illustrated in this example)

Step 3 After the Cisco Guard XT is activated, the Cisco Guard XT performs additional screening, and then the traffic destined to the zone under attack is diverted to the Cisco Guard XT in any of the following ways:

- The Cisco Guard XT can issue a BGP route update telling the router to divert the traffic to the Cisco Guard TX.
- If you are using the Catalyst 6500/7600 modules, the Route Health Injection (RHI) feature can trigger the packet diversion.
- A route is injected externally into the network.

Step 4 The attack traffic is redirected to the Cisco Guard XT, and legitimate traffic is allowed to the protected zone, as illustrated in Figure 3-12.

Figure 3-12 *Attack Traffic Redirected*

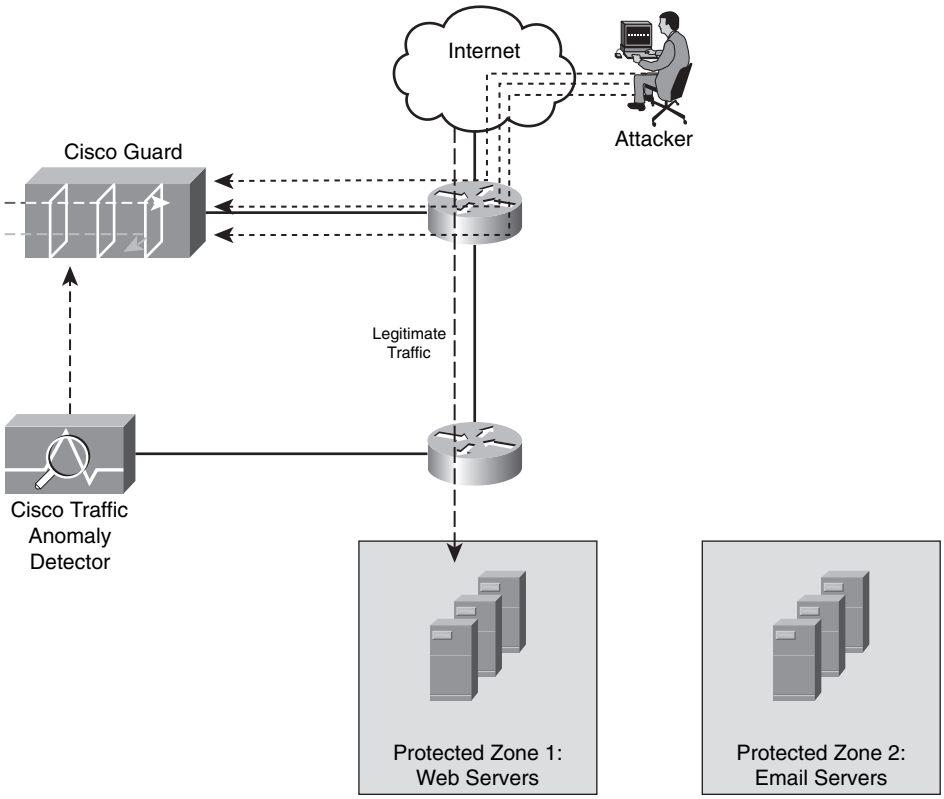

The Cisco Guard can also be deployed with other anomaly detection systems. Examples of this include Arbor's Peakflow SP and Peakflow X. Arbor's Peakflow SP is designed for service providers, and Peakflow X is designed for enterprises. Typically, enterprises deploy the Cisco Guard XT at their Internet edge, or they co-locate it at their Internet service provider network to avoid the unnecessary traffic consuming their bandwidth. Because of this, numerous service providers offer managed network DDoS protection, hosting DDoS protection, peering point DDoS protection, and infrastructure protection services. This is based on a solution that Cisco makes available to service providers called "clean pipes."

NOTE For more information about clean pipes, go to http://www.cisco.com/go/cleanpipes.

Figure 3-13 illustrates the protection cycle that the Cisco Guard XT follows to analyze, filter, and rate-limit the traffic.

Figure 3-13 *Cisco Guard XT Protection Cycle*

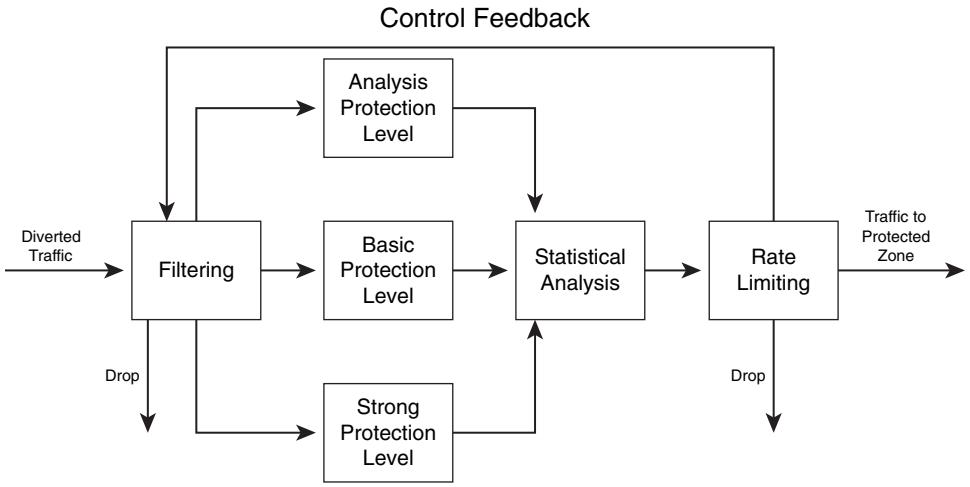

When the traffic is redirected to the Cisco Guard XT, it first filters the traffic using several filtering techniques. If the Cisco Guard XT determines that the packets are malicious, it drops them at this stage. If the packets are not malicious, the packets are sent to different protection levels using several types of authentication methods. Subsequently, the Cisco Guard XT analyzes the traffic flow, drops the traffic that exceeds the defined rate that the zone can handle, and then injects the legitimate traffic back to the zone. A closed-loop feedback cycle dynamically adjusts its protection policies.

NOTE For more detailed information on how to configure the Cisco Guard XT and the Cisco TAD XT, go to http://www.cisco.com/en/US/products/ps5888/products_installation_and_configuration_guides_list.html.

Intrusion Detection and Intrusion Prevention Systems (IDS/IPS)

In Chapter 1, "Overview of Network Security Technologies," you learned the basics about IDS and IPS systems. IDSs are devices that in promiscuous mode detect malicious activity within the network. IPS devices are capable of detecting all these security threats; however, they are also able to drop noncompliant packets inline. Traditionally, IDS systems have provided excellent application layer attack-detection capabilities; however, they were not able to protect against day-zero attacks using valid packets. The problem is that most attacks today use valid packets. On the other hand, now IPS systems such as the Cisco IPS software Version 6.x and later offer anomaly-based capabilities that help you detect such attacks. This is a big advantage, since it makes the IPS devices less dependent on signature updates for protection against DDoS, worms, and any day-zero threats. Just like any other anomaly detection systems, the sensors need to learn what is "normal." In other words, they need to create a baseline of legitimate behavior.

The Importance of Signatures Updates

Traditionally, IPS and IDS systems depend on signatures to operate. Because of this, it is extremely important to tune the IPS/IDS device accordingly and to develop policies and procedures to continuously update the signatures. The Cisco IPS software allows you to automatically download signatures from a management station. Signature updates are posted to Cisco.com almost on a weekly basis. In Chapter 2, you learned about the Cisco Security Center (historically named mySDN or my Self Defending Network). This is an excellent resource to obtain information about the latest IPS signatures and other security intelligence information.

NOTE The Cisco Security Center site is http://www.cisco.com/security.

The Cisco Security Center provides up-to-date security intelligence data, in addition to detailed IDS/IPS signature information.

Although the IPS sensors can work without a license key, you must have a license key to obtain signature updates from Cisco.com. To obtain a license key, you must have a Cisco Service for IPS service contract. For more information, go to http://www.cisco.com/go/license.

The Cisco IPS Device Manager (IDM) is a web-based configuration utility used to manage individual IPS sensors, Catalyst 6500 IPS modules, and the Advanced Inspection and Prevention Security Services Module (AIP-SSM) for the Cisco ASA. You can configure the IPS device via IDM to automatically obtain and install signatures from an FTP or SCP server.

NOTE You cannot automatically download service pack and signature updates from Cisco.com. You need to download service packs and signatures updates from Cisco.com to an FTP or SCP server. Then you can configure your IPS device to access the files on your server. You can also use the Cisco Security Manager IPS Manager Console (IPSMC) to manage your IPS devices. You can configure IPSMC to automatically download the signature updates and service packs from Cisco.com and then install them in your IPS devices. For more information about IPSMC, go to http://www.cisco.com/go/security.

Complete the following steps to configure IDM to automatically download signatures from your FTP or SCP server.

Step 1 Log in to IDM with an administrator account and navigate to **Configuration > Auto Update**.

Step 2 Select the **Enable Auto Update** check box.

Step 3 Enter the IP address of the remote server where the signature update or service packs are saved.

Step 4 Select either **FTP** or **SCP** for your transport mechanism/server type.

Step 5 Enter the path to the directory on the remote server where the updates are located in the **Directory Path**.

Step 6 Enter the username and password of the account in your FTP or SCP server.

Step 7 You can configure the IPS device to check for updates hourly or on a weekly basis. If you want your IPS device to check for updates hourly, check the **Hourly** check box. Then enter the time you want the updates to start and the hour interval at which you want the IPS device to contact your remote server for updates. The IPS sensor checks the directory you specified for new files in your server. Only one update is installed per cycle even if there are multiple available files.

Step 8 Check the **Daily** check box if you want the IPS device to automatically check for updates on a daily basis. Then enter the time you want the updates to start and check the days you want the IPS device to check for updates in your SCP or FTP server.

Step 9 To save and apply your configuration, click **Apply**.

The Importance of Tuning

Chapter 1 showed you the important factors to consider when tuning your IPS/IDS devices. Each IPS/IDS device comes with a preset number of signatures enabled. These signatures are suitable in most cases; however, it is important that you tune your IPS/IDS devices when you first deploy them and then tune them again periodically. You could receive numerous false positive events (false alarms), which could cause you to overlook real security incidents. The initial tuning will probably take more time than any subsequent tuning. The initial tuning process is hard to perform manually, especially in large environments where several IPS/IDS devices are deployed and hundreds of events are generated in short periods. This is why it is important to use event correlation systems to alleviate this process and save numerous hours. CS-MARS is used in the following example to perform initial tuning and event analysis.

In this example, several IPS devices are sending their events to a CS-MARS. The administrator completes the following steps to perform initial tuning:

Step 1 Log in to the CS-MARS via the web interface.

Step 2 Click **Query/Reports** tab.

Step 3 Select the **Activity: All–Top Event Types (Peak View)** option from the second pull-down menu under the **Load Report as On-Demand Query with Filter** section, as shown in Figure 3-14.

Figure 3-14 *CS-MARS Query/Reports*

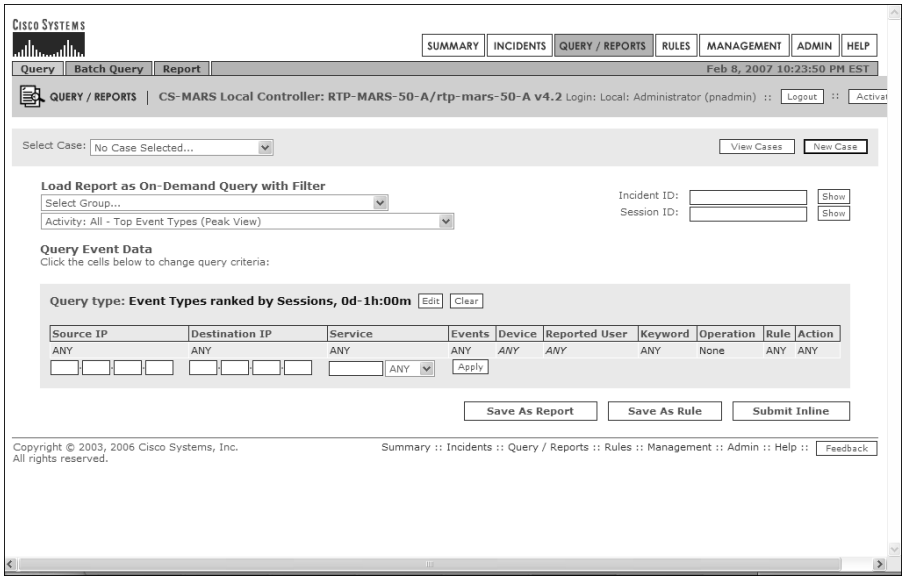

Step 4 Click the **Edit** button to select the time interval for the query and enter **1** day under the **Filter by time** section to trigger the CS-MARS to display the top event types in the past 24 hours, as shown in Figure 3-15.

Figure 3-15 *Selecting the Query Time Interval*

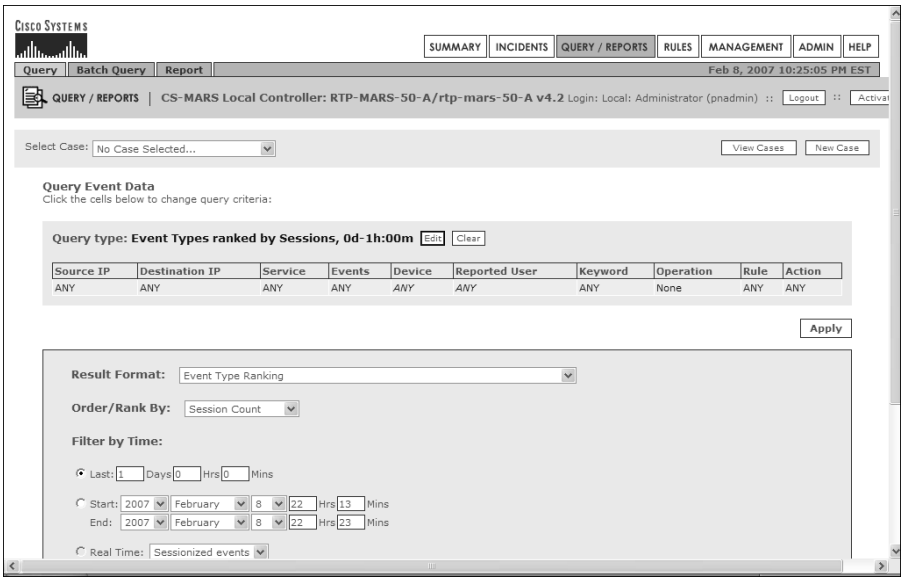

Step 5 Click **Apply** and **Submit Inline** in the next screen to obtain the report. The report in Figure 3-16 is shown. In this report, the administrator notices that there have been more than 480 ARP Reply-to-Broadcast events detected in the past 24 hours.

Step 6 Click the event to obtain more information and read the following from the CS-MARS details screen: "This signature detects an ARP Reply packet where the destination MAC address in the ARP payload is a layer 2 broadcast address. This is not normal traffic and can indicate an ARP poisoning attack."

Step 7 Click **q** by the event and select **Source IP Address Ranking** under the **Result format** section to investigate the source, as shown in Figure 3-17.

Figure 3-16 *Top Event Types*

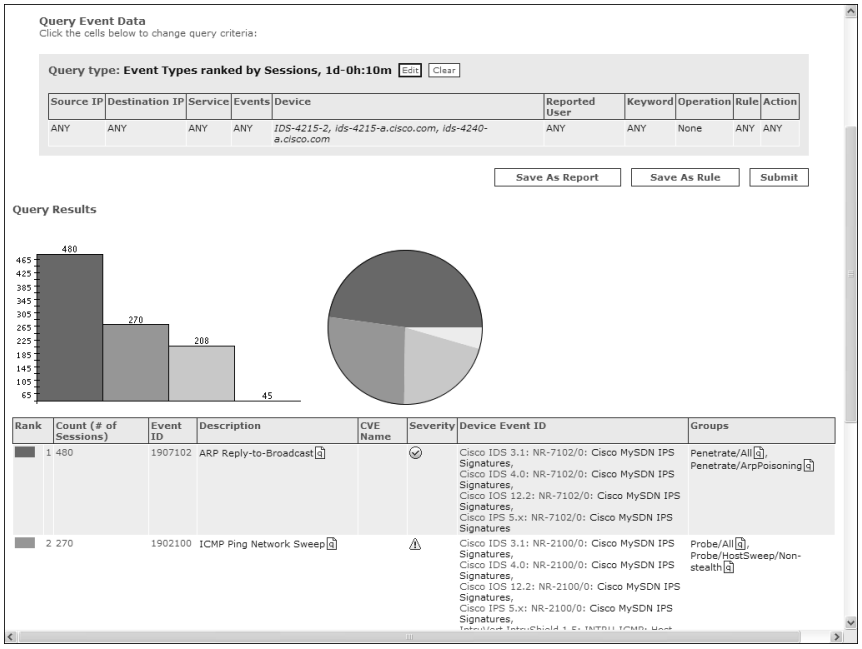

Figure 3-17 *Verifying Sources*

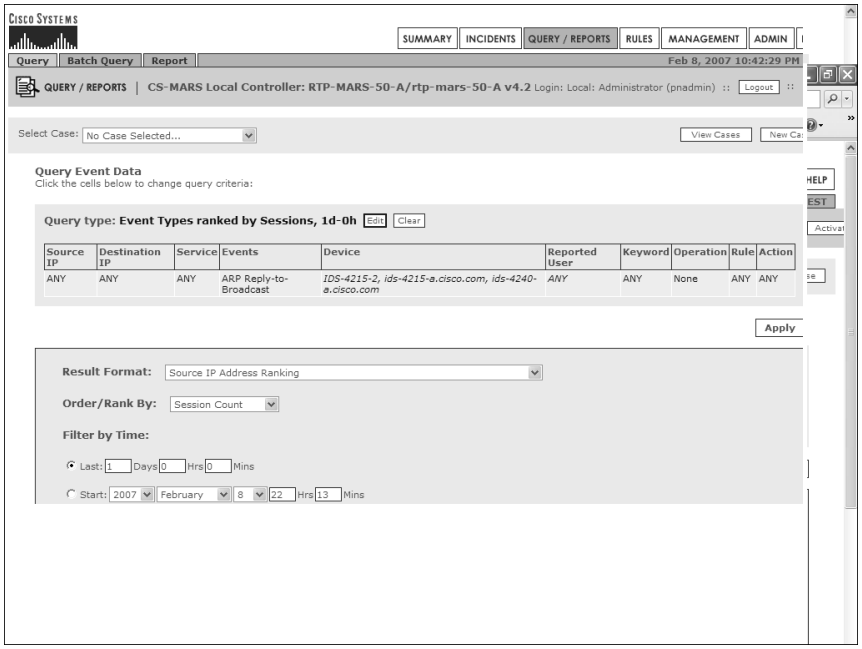

Step 8 Click **Apply** and **Submit Inline** in the following screen to obtain the new report, including the source IP addresses for the *ARP Reply-to-Broadcast* events. The report is shown as illustrated in Figure 3-18.

Figure 3-18 *IP Sources Report*

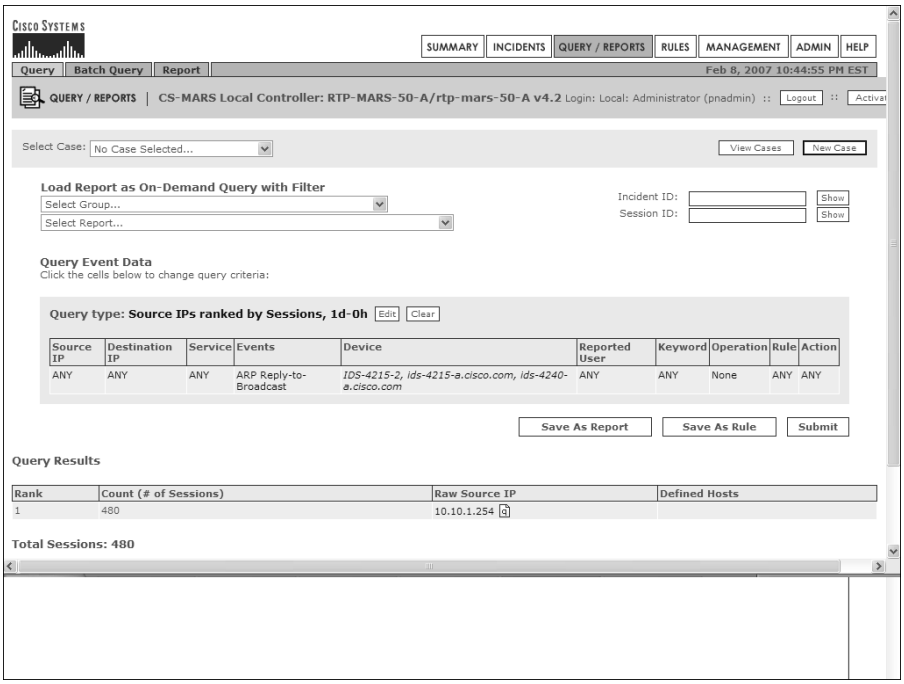

The administrator notices that only one device (10.10.1.254) is triggering these events. After further investigation, he discovers that this is the normal behavior of an application that is running on that machine and marks this incident as a **False Positive** in CS-MARS.

The administrator notices that these events are not shown anymore in CS-MARS; however, they are still shown using the **show events** command in the CLI of the IPS sensors. This is because when you mark an incident/event/session in CS-MARS as a **False Positive**, it does not disable or tune this signature in the actual IPS device. The events are still sent to the CS-MARS from the IPS devices; however, CS-MARS does not process these events. If you do not want the IPS sensor to send or process the events, you must tune or disable the signature on the IPS device. You can tune signatures based on source and destination. For example, in this case, you can tune the IPS signature not to alert you if the host with the

IP address 10.10.1.254 sends this type of packet. However, you can configure the IPS signature to alert you if any other device generates this type of traffic.

Anomaly Detection Within Cisco IPS Devices

When you configure a Cisco IPS device running Versions 6.x and later with anomaly detection services, the IPS device initially goes through a learning process. This is done to configure a set of policy thresholds based on the normal behavior of your network. Three different modes of operation take place when an IPS device is configured with anomaly detection:

- Learning mode
- Detect mode
- Inactive mode

The initial learning mode is performed over a period of 24 hours, by default. The initial baseline is referred to as the knowledge base (KB) of your traffic.

TIP The IPS sensor does not detect attacks during the initial learning phase. If you experience an attack during this period, your results will not reflect a baseline of normal network behavior. This is an important point to take into consideration. Depending on your environment, you may want to have the IPS device in learning mode longer than the default 24 hours because this is a configurable value. Do not initially enable your IPS device with anomaly detection over a weekend if your organization operates mostly during normal business hours and days. This is a huge mistake that many people make.

To configure the IPS sensor using IDM to start the learning mode, go to **Configuration > Policies > Anomaly Detections > ad0 > Learning Accept Mode** and select the **Automatically accept learning knowledge base** check box. In that section, you can also specify the learning period length.

After the learning process, a KB is created that replaces the initial KB. The IPS device then automatically goes into detect mode. Any traffic flows that violate thresholds in the KB trigger the IPS device to generate alerts. The IPS device also keeps track of gradual changes to the KB that do not violate the thresholds and adjusts its configuration.

You can turn off the anomaly detection functionality on your IPS device. This is called being in *inactive mode*. In certain circumstances, this is needed. An example is when you have an asymmetric environment and the IPS device gets traffic from different directions, causing it to operate incorrectly.

NOTE The traffic anomaly engine in Cisco IPS devices uses nine anomaly detection signatures covering TCP, UDP, and other protocols. Each signature has two subsignatures: one for the scanner and the other for the worm-infected host. All of these signatures are enabled by default, and they are in the 13000 range.

Similarly to the Cisco TAD XT, the anomaly detection feature in Cisco IPS devices uses zones. The purpose of configuring zones is to make sure that you do not have false positives and false negatives. A *zone* is a set of destination IP addresses. Three different zones exist:

- **Internal:** You configure this zone with the IP address range of your internal network.
- **Illegal:** You configure this zone with IP address ranges that should never be seen in normal traffic. Here you should use unallocated IP addresses or bogon IP addresses.
- **External:** This is the default zone. By default, it has the Internet range of 0.0.0.0-255.255.255.255.

To configure the Internal zone in your IPS device using IDM, complete the following steps:

Step 1 Navigate to **Configuration > Policies > Anomaly Detections > ad0 > Internal Zone**. The Internal Zone tab appears.

Step 2 Click the **General** tab.

Step 3 Select the **Enable the Internal Zone** check box.

Step 4 Enter your internal subnets/IP address range in the **Service Subnets** field. IDM also allows you to configure protocol and other specific thresholds.

NOTE For more information on how to configure other thresholds and anomaly detection functionality, refer to the Cisco IPS configuration guides located at http://www.cisco.com/univercd/cc/td/doc/product/iaabu/csids/csids13/idmguide/index.htm.

Summary

Identification and classification of security threats mainly concerns visibility. In this chapter, you learned how important it is to have complete network visibility and control to successfully identify and classify security threats in a timely fashion. This chapter also covered different technologies and tools that can be used to obtain information from your network and detect anomalies that can be malicious activity. This chapter provided overviews of Cisco NetFlow, SYSLOG, and SNMP. You also learned about robust event correlation systems, such as CS-MARS and open source monitoring systems that can be used in conjunction with NetFlow to allow you to gain better visibility in your network.

This chapter also provided an overview of anomaly detection solutions, in addition to tips on IPS/IDS tuning and the new anomaly detection features that Cisco IPS software supports.