Introduction to Virtual Appliances

The Cisco FireSIGHT® System combines the security of an industry-leading network intrusion protection system with the power to control access to your network based on detected applications, users, and URLs.

Cisco packages 64-bit virtual Defense Centers® and virtual devices for the VMware vSphere and VMware vCloud Director hosting environments. You can deploy 64-bit virtual Defense Centers and 64-bit virtual managed devices to ESXi hosts using a vCenter, or using vCloud Director. Virtual appliances use e1000 (1 Gbit/s) interfaces, or you can replace the default interfaces with vmxnet3 (10 Gbit/s) interfaces. You can also use VMware Tools to improve the performance and management of your virtual appliances.

The Defense Center provides a centralized management console and database repository for the system. Virtual devices can inspect traffic on virtual or physical networks in either a passive or inline deployment:

- Virtual devices in a passive deployment simply monitor traffic flowing across a network.
- Passive sensing interfaces receive all traffic unconditionally and no traffic received on these interfaces is retransmitted.
- Virtual devices in an inline deployment allow you to protect your network from attacks that might affect the availability, integrity, or confidentiality of hosts on the network. Inline devices can be deployed as a simple intrusion prevention system. You can also configure inline devices to perform access control as well as manage network traffic in other ways.
- Inline interfaces receive all traffic unconditionally, and traffic received on these interfaces is retransmitted unless explicitly dropped by some configuration in your deployment.

Virtual Defense Centers can manage physical devices, Cisco NGIPS for Blue Coat X-Series, and Cisco ASA with FirePOWER Services (ASA FirePOWER), and physical Defense Centers can manage virtual devices. However, virtual appliances do not support any of the system's hardware-based features—virtual Defense Centers do not support high availability and virtual devices do not support clustering, stacking, switching, routing, and so on. For detailed information on physical FireSIGHT System appliances, see the *FireSIGHT System Installation Guide*.

This installation guide provides information about deploying, installing, and setting up virtual FireSIGHT System appliances (devices and Defense Centers). It also assumes familiarity with the features and nomenclature of VMware products, including the vSphere Client, VMware vCloud Director web portal, and, optionally, VMware Tools.

The topics that follow introduce you to FireSIGHT System virtual appliances:

- FireSIGHT System Virtual Appliances, page 1-2
- Understanding Virtual Appliance Capabilities, page 1-3

FireSIGHT System Virtual Appliances

A FireSIGHT System *virtual appliance* is either a traffic-sensing managed *virtual device* or a managing *virtual Defense Center*. For more information, see the following sections:

Guidelines and Limitations

The following limitations exist when deploying virtual Defense Center or devices on VMware:

- vMotion is not supported.
- Cloning a virtual machine is not supported.
- Restoring a virtual machine with a snapshot is not supported.
- Restoring a backup is not supported.

Virtual Defense Centers

A Defense Center provides a centralized management point and event database for your FireSIGHT System deployment. Virtual Defense Centers aggregate and correlate intrusion, file, malware, discovery, connection, and performance data, assessing the impact of events on particular hosts and tagging hosts with indications of compromise. This allows you to monitor the information that your devices report in relation to one another, and to assess and control the overall activity that occurs on your network.

Key features of the virtual Defense Center include:

- device, license, and policy management
- event and contextual information displayed in tables, graphs, and charts
- health and performance monitoring
- external notification and alerting
- correlation, indications of compromise, and remediation features for real-time threat response
- custom and template-based reporting

Virtual Managed Devices

Virtual devices deployed on network segments within your organization monitor traffic for analysis. Virtual devices deployed passively help you gain insight into your network traffic. Deployed inline, you can use virtual devices to affect the flow of traffic based on multiple criteria. Depending on model and license, devices:

- gather detailed information about your organization's hosts, operating systems, applications, users, files, networks, and vulnerabilities
- block or allow network traffic based on various network-based criteria, as well as other criteria including applications, users, URLs, IP address reputations, and the results of intrusion or malware inspections

Virtual devices do **not** have a web interface. You must configure them via console and command line, and you must manage them with a Defense Center.

Understanding Virtual Appliance Capabilities

Virtual appliances have many of the capabilities of physical appliances:

- The virtual Defense Center has the same features as a physical Defense Center, except you cannot create high availability pairs of virtual Defense Centers. With a FireSIGHT license, the virtual Defense Center can monitor 50,000 hosts and users.
- Virtual devices have the traffic and blocking analysis capabilities of physical devices. However, they cannot perform switching, routing, VPN, and other hardware-based, redundancy, and resource-sharing features.

Understanding Virtual Defense Center Capabilities

Table 1-1 Supported Capabilities for Virtual Defense Centers, page 1-3 matches the major capabilities of the system with virtual Defense Centers, assuming you are managing devices that support those features and have the correct licenses installed and applied.

For a brief summary of the features and licenses supported with virtual appliances, see FireSIGHT System Components, page 1-7 and Licensing Virtual Appliances, page 1-11.

Keep in mind that virtual Defense Centers can manage Series 2, Series 3, ASA FirePOWER, and X-Series devices. Similarly, Series 2 and Series 3 Defense Centers can manage virtual devices. The Defense Center column for device-based capabilities (such as stacking, switching, and routing) indicates whether a virtual Defense Center can manage and configure devices to perform those functions. For example, although you cannot configure VPN on a virtual device, you can use a virtual Defense Center to manage Series 3 devices in a VPN deployment.

Table 1-1 Supported Capabilities for Virtual Defense Centers

Feature or Capability	Virtual Defense Center
collect discovery data (host, application, and user) reported by managed devices and build a network map for your organization	yes
view geolocation data for your network traffic	yes
manage an intrusion detection and prevention (IPS) deployment	yes

Table 1-1 Supported Capabilities for Virtual Defense Centers (continued)

Feature or Capability	Virtual Defense Center
manage devices performing Security Intelligence filtering	yes
manage devices performing simple network-based control, including geolocation-based filtering	yes
manage devices performing application control	yes
manage devices performing user control	yes
manage devices that filter network traffic by literal URL	yes
manage devices performing URL filtering by category and reputation	yes
manage devices performing simple file control by file type	yes
manage devices performing network-based advanced malware protection (AMP)	yes
receive endpoint-based malware (FireAMP) events from your FireAMP deployment	yes
manage device-based hardware-based features: • fast-path rules • strict TCP enforcement • configurable bypass interfaces • tap mode • switching and routing • NAT policies • VPN	yes
manage device-based redundancy and resource sharing: • device stacks • device clusters • Cisco NGIPS for Blue Coat X-Series VAP groups • clustered stacks	yes
separate and manage internal and event traffic using traffic channels	yes
isolate and manage traffic on different networks using multiple management interfaces	yes
establish high availability	no
install a malware storage pack	no
connect to an eStreamer, host input, or database client	yes

Understanding Virtual Managed Device Capabilities

Table 1-2 Supported Capabilities for Virtual Managed Devices, page 1-5 matches the major capabilities of the system with virtual managed devices, assuming you have the correct licenses installed and applied from the managing Defense Center.

Keep in mind that although you can use any model of Defense Center running Version 5.4.1 of the system to manage any Version 5.4.1 virtual device, a few system capabilities are limited by the Defense Center model. For example, you cannot use the Series 2 DC500 to manage virtual managed devices performing Security Intelligence filtering, even though virtual managed devices support that capability. For more information, see Understanding Virtual Defense Center Capabilities, page 1-3.

Table 1-2 Supported Capabilities for Virtual Managed Devices

Feature or Capability	Virtual Managed Device
collect discovery data (host, application, and user) reported by managed devices and build a network map for your organization	yes
view geolocation data for your network traffic	yes
network discovery: host, application, and user	yes
intrusion detection and prevention (IPS)	yes
Security Intelligence filtering	yes
access control: basic network control	yes
access control: geolocation-based filtering	yes
access control: application control	yes
access control: user control	yes
access control: literal URLs	yes
access control: URL filtering by category and reputation	yes
file control: by file type	yes
network-based advanced malware protection (AMP)	yes
Automatic Application Bypass	yes
fast-path rules	no
strict TCP enforcement	no
configurable bypass interfaces	no
tap mode	no
switching and routing	no
NAT policies	no
VPN	no
device stacking	no
device clustering	no
clustered stacks	no
traffic channels	no
multiple management interfaces	no
malware storage pack	no
FireSIGHT System-specific interactive CLI	yes
connect to an eStreamer client	no

Operating Environment Prerequisites

You can host 64-bit virtual appliances on the following hosting environments:

- VMware ESXi 5.5 (vSphere 5.5)
- VMware ESXi 5.1 (vSphere 5.1)
- VMware vCloud Director 5.1

You can also enable VMware Tools on all supported ESXi versions. For information on the full functionality of VMware Tools, see the VMware website (http://www.vmware.com/). For help creating a hosting environment, see the VMware ESXi documentation, including VMware vCloud Director and VMware vCenter.

Virtual appliances use Open Virtual Format (OVF) packaging. VMware Workstation, Player, Server, and Fusion do not recognize OVF packaging and are not supported. Additionally, virtual appliances are packaged as virtual machines with Version 7 of the virtual hardware.

The computer that serves as the ESXi host must meet the following requirements:

- It must have a 64-bit CPU that provides virtualization support, either Intel® Virtualization Technology (VT) or AMD Virtualization™ (AMD-V™) technology.
- Virtualization must be enabled in the BIOS settings
- To host virtual devices, the computer must have network interfaces compatible with Intel e1000 drivers (such as PRO 1000MT dual port server adapters or PRO 1000GT desktop adapters).

For more information, see the VMware website: http://www.vmware.com/resources/guides.html.

Each virtual appliance you create requires a certain amount of memory, CPUs, and hard disk space on the ESXi host. Do **not** decrease the default settings, as they are the minimum required to run the system software. However, to improve performance, you can increase a virtual appliance's memory and number of CPUs, depending on your available resources. The following table lists the default appliance settings.

Table 1-3 Default Virtual Appliance Settings

Setting	Default	Adjustable Setting?
memory	4GB	yes, and for a virtual device you **must** allocate: • 4GB minimum • 5GB to use category and reputation based URL filtering • 6GB to perform Security Intelligence filtering using large dynamic feeds • 7GB to perform URL filtering and Security Intelligence
virtual CPUs	4	yes, up to 8
hard disk provisioned size	40GB (device) 250GB (Defense Center)	no

Virtual Appliance Performance

It is not possible to accurately predict throughput and processing capacity for virtual appliances. A number of factors heavily influence performance, such as the:

- amount of memory and CPU capacity of the ESXi host
- number of total virtual machines running on the ESXi host
- number of sensing interfaces, network performance, and interface speed
- amount of resources assigned to each virtual appliance
- level of activity of other virtual appliances sharing the host
- complexity of policies applied to a virtual device

Tip VMware provides a number of performance measurement and resource allocation tools. Use these tools on the ESXi host while you run your virtual appliance to monitor traffic and determine throughput. If the throughput is not satisfactory, adjust the resources assigned to the virtual appliances that share the ESXi host.

You can enable VMware Tools to improve the performance and management of your virtual appliances. Alternatively, you can install tools (such as `esxtop` or VMware/third-party add-ons) on the host or in the virtualization management layer (not the guest layer) on the ESXi host to examine virtual performance. To enable VMware Tools, see the *FireSIGHT System User Guide*.

FireSIGHT System Components

The sections that follow describe some of the key capabilities of virtual Defense Centers and virtual devices that contribute to your organization's security, acceptable use policy, and traffic management strategy. For information on the additional features supported with Series 2 and Series 3 appliances, see the *FireSIGHT System Installation Guide* and the *FireSIGHT System User Guide*.

Tip Many virtual appliance capabilities are license and user role dependent. Where needed, FireSIGHT System documentation outlines the requirements for each feature and task.

The topics that follow describe some of the key capabilities of the FireSIGHT System that contribute to your organization's security, acceptable use policy, and traffic management strategy:

- FireSIGHT, page 1-8
- Access Control, page 1-8
- Intrusion Detection and Prevention, page 1-8
- File Tracking, Control, and Malware Protection, page 1-9
- Application Programming Interfaces, page 1-10

FireSIGHT

FireSIGHT™ is Cisco's discovery and awareness technology that collects information about hosts, operating systems, applications, users, files, networks, geolocation information, and vulnerabilities, in order to provide you with a complete view of your network.

You can use the Defense Center's web interface to view and analyze data collected by FireSIGHT. You can also use this data to help you perform access control and modify intrusion rule states. In addition, you can generate and track indications of compromise on hosts on your network based on correlated event data for the hosts.

Access Control

Access control is a policy-based feature that allows you to specify, inspect, and log the traffic that can traverse your network. An *access control policy* determines how the system handles traffic on your network. You can use a policy that does not include *access control rules* to handle traffic in one of the following ways, using what is called the *default action*:

- block all traffic from entering your network
- trust all traffic to enter your network without further inspection
- allow all traffic to enter your network, and inspect the traffic with a network discovery policy only
- allow all traffic to enter your network, and inspect the traffic with intrusion and network discovery policies

You can include access control rules in an access control policy to further define how traffic is handled by targeted devices, from simple IP address matching to complex scenarios involving different users, applications, ports, and URLs. For each rule, you specify a rule *action*, that is, whether to trust, monitor, block, or inspect matching traffic with an intrusion or file policy.

For each access control policy, you can create a custom HTML page that users see when the system blocks their HTTP requests. Optionally, you can display a page that warns users, but also allows them to click a button to continue to the originally requested site.

As part of access control, the Security Intelligence feature allows you to blacklist—deny traffic to and from—specific IP addresses before the traffic is subjected to analysis by access control rules. If your system supports geolocation, you can also filter traffic based on its detected source and destination countries and continents.

Access control includes intrusion detection and prevention, file control, and advanced malware protection. For more information, see the next sections.

Intrusion Detection and Prevention

Intrusion detection and prevention allows you to monitor your network traffic for security violations and, in inline deployments, to block or alter malicious traffic.

Intrusion prevention is integrated into access control, where you can associate an intrusion policy with specific access control rules. If network traffic meets the conditions in a rule, you can analyze the matching traffic with an intrusion policy. You can also associate an intrusion policy with the default action of an access control policy.

An intrusion policy contains a variety of components, including:

- rules that inspect the protocol header values, payload content, and certain packet size characteristics

- rule state configuration based on FireSIGHT recommendations
- advanced settings, such as preprocessors and other detection and performance features
- preprocessor rules that allow you to generate events for associated preprocessors and preprocessor options

File Tracking, Control, and Malware Protection

To help you identify and mitigate the effects of malware, the FireSIGHT System's file control, network file trajectory, and advanced malware protection components can detect, track, capture, analyze, and optionally block the transmission of files (including malware files) in network traffic.

File Control

File control allows managed devices to detect and block your users from uploading (sending) or downloading (receiving) files of specific types over specific application protocols. You configure file control as part of your overall access control configuration; file policies associated with access control rules inspect network traffic that meets rule conditions.

Network-Based Advanced Malware Protection (AMP)

Network-based *advanced malware protection* (AMP) allows the system to inspect network traffic for malware in several types of files. Virtual devices can store detected files for further analysis to a hard drive.

Regardless of whether you store a detected file, you can submit it to the Collective Security Intelligence Cloud for a simple known-disposition lookup using the file's SHA-256 hash value. You can also submit files for *dynamic analysis*, which produces a threat score. Using this contextual information, you can configure the system to block or allow specific files.

You configure malware protection as part of your overall access control configuration; file policies associated with access control rules inspect network traffic that meets rule conditions.

FireAMP Integration

FireAMP is Cisco's enterprise-class, advanced malware analysis and protection solution that discovers, understands, and blocks advanced malware outbreaks, advanced persistent threats, and targeted attacks.

If your organization has a FireAMP subscription, individual users install *FireAMP Connectors* on their computers and mobile devices (also called *endpoints*). These lightweight agents communicate with the Collective Security Intelligence Cloud, which in turn communicates with the Defense Center.

After you configure the Defense Center to connect to the cloud, you can use the Defense Center web interface to view endpoint-based malware events generated as a result of scans, detections, and quarantines on the endpoints in your organization. The Defense Center also uses FireAMP data to generate and track indications of compromise on hosts, as well as display network file trajectories.

Use the *FireAMP portal* to configure your FireAMP deployment. The portal helps you quickly identify and quarantine malware. You can identify outbreaks when they occur, track their trajectories, understand their effects, and learn how to successfully recover. You can also use FireAMP to create custom protections, block execution of certain applications based on group policy, and create custom whitelists.

See http://amp.sourcefire.com/ for more information.

Network File Trajectory

The network file trajectory feature allows you to track a file's transmission path across a network. The system uses SHA-256 hash values to track files; so, to track a file, the system must either:

- calculate the file's SHA-256 hash value and perform a malware cloud lookup using that value
- receive endpoint-based threat and quarantine data about that file, using the Defense Center's integration with your organization's FireAMP subscription

Each file has an associated trajectory map, which contains a visual display of the file's transfers over time as well as additional information about the file.

Application Programming Interfaces

There are several ways to interact with the system using application programming interfaces (APIs). For detailed information, you can download additional documentation from the Support Site.

eStreamer

The Event Streamer (eStreamer) allows you to stream several kinds of event data from a Cisco appliance to a custom-developed client application. After you create a client application, you can connect it to an eStreamer server (Defense Center or managed device), start the eStreamerservice, and begin exchanging data.

eStreamer integration requires custom programming, but allows you to request specific data from an appliance. If, for example, you display network host data within one of your network management applications, you could write a program to retrieve host criticality or vulnerability data from the Defense Center and add that information to your display.

External Database Access

The database access feature allows you to query several database tables on a Defense Center, using a third-party client that supports JDBC SSL connections.

You can use an industry-standard reporting tool such as Crystal Reports, Actuate BIRT, or JasperSoft iReport to design and submit queries. Or, you can configure your own custom application to query Cisco data. For example, you could build a servlet to report intrusion and discovery event data periodically or refresh an alert dashboard.

Host Input

The host input feature allows you to augment the information in the network map by importing data from third-party sources using scripts or command-line files.

The web interface also provides some host input functionality; you can modify operating system or application protocol identities, validate or invalidate vulnerabilities, and delete various items from the network map, including clients and server ports.

Remediation

The system includes an API that allows you to create remediations that your Defense Center can automatically launch when conditions on your network violate an associated correlation policy or compliance white list. This can not only automatically mitigate attacks when you are not immediately available to address them, but can also ensure that your system remains compliant with your organization's security policy. In addition to remediations that you create, the Defense Center ships with several predefined remediation modules.

Multiple Management Interfaces

You can use *multiple management interfaces* on Series 3 appliances and the virtual Defense Center to improve performance by separating traffic into two traffic channels: *management traffic channel* to carry inter-device communication and *event traffic channel* to carry event traffic such as web access. Both traffic channels can be carried on the same management interface or split between two management interfaces, each interface carrying one traffic channel.

You can create a route from a specific management interface on your Defense Center to a different network, allowing your Defense Center to manage traffic from devices on one network separately from traffic from devices on another network.

Additional management interfaces function the same as the default management interface (such as using high availability between the Defense Centers) with the following exceptions:

- You can configure DHCP on the default (`eth0`) management interface only. Additional (`eth1` and so on) interfaces require unique static IP addresses and hostnames.
- You must configure both traffic channels to use the same management interface when you use a non-default management interface to connect your Defense Center and managed device and those appliances are separated by a NAT device.
- On the 70xx Family, you can separate traffic into two channels and configure those channels to send traffic to one or more management interfaces on the virtual Defense Center. However, because the 70xx Family contains only one management interface, the device receives traffic sent from the Defense Center on only one management interface.

After your appliance is installed, use the web browser to configure multiple management interfaces. To add a management interface to your virtual Defense Center, see Adding and Configuring Interfaces, page 4-9. See Multiple Management Interfaces in the *FireSIGHT System User Guide* for more information.

Licensing Virtual Appliances

You can license a variety of features to create an optimal FireSIGHT System deployment for your organization. You must use the Defense Center to control licenses for itself and the devices it manages.

Cisco recommends you add the licenses your organization has purchased during the initial setup of your Defense Center. Otherwise, any devices you register during initial setup are added to the Defense Center as unlicensed. You must then enable licenses on each device individually after the initial setup process is over. For more information, see Setting Up Virtual Appliances, page 5-1.

A FireSIGHT license is included with each Defense Center purchase, and is required to perform host, application, and user discovery. The FireSIGHT license on a Defense Center also determines how many individual hosts and users you can monitor with the Defense Center and its managed devices, as well as how many users you can allow to perform user control. For a virtual Defense Center, this limit is 50,000 individual hosts and users.

If your Defense Center was previously running Version 4.10.x, you may be able to use legacy RNA Host and RUA User licenses instead of a FireSIGHT license. For more information, see License Settings, page 5-10.

Additional model-specific licenses allow your managed devices to perform a variety of functions, as follows:

Protection

A Protection license allows virtual devices to perform intrusion detection and prevention, file control, and Security Intelligence filtering.

Control

A Control license allows virtual devices to perform user and application control. Although virtual devices do not support any of the hardware-based features granted to Series 2 and Series 3 devices by the Control license (such as switching or routing), virtual Defense Centers can manage those features on physical devices. A Control license requires a Protection license.

URL Filtering

A URL Filtering license allows virtual devices to use regularly updated cloud-based category and reputation data to determine which traffic can traverse your network, based on the URLs requested by monitored hosts. A URL Filtering license requires a Protection license.

Malware

A Malware license allows virtual devices to perform network-based advanced malware protection (AMP), that is, to detect and block malware in files transmitted over your network. It also allows you to view trajectories, which track files transmitted over your network. A Malware license requires a Protection license.

VPN

A VPN license allows you to use a virtual Defense Center to build secure VPN tunnels among the virtual routers on Series 3 devices, or from Series 3 devices to remote devices or other third-party VPN endpoints. A VPN license requires Protection and Control licenses.

Because of architecture and resource limitations, not all licenses can be applied to all managed devices. In general, you cannot license a capability that a device does not support; see Understanding Virtual Appliance Capabilities, page 1-3.

The following table summarizes which licenses you can add to your Defense Center and apply to each device model. The Defense Center rows (for all licenses except FireSIGHT) indicate whether that Defense Center can manage devices using those licenses. For example, you can use a Series 2 DC1000 to create a VPN deployment using Series 3 devices, but you cannot use a DC500 to perform category and reputation-based URL Filtering, regardless of the devices it manages. Note that `n/a` marks Defense Center-based licenses that are not relevant to managed devices.

Table 1-4 Supported Licenses by Model

Models	FireSIGHT	Protection	Control	URL Filtering	Malware	VPN
Series 2 devices: • 3D500, 3D1000, 3D2000 • 3D2100, 3D2500, 3D3500, 3D4500 • 3D6500 • 3D9900	n/a	automatic, no Security Intelligence	no	no	no	no
Series 3 devices: • 7000 Series • 8000 Series	n/a	yes	yes	yes	yes	yes
virtual devices	n/a	yes	yes, but no support for hardware features	yes	yes	no
Cisco ASA with FirePOWER Services	n/a	yes	yes, but no support for hardware features	yes	yes	no
Cisco NGIPS for Blue Coat X-Series	n/a	yes	yes, but no support for hardware features	yes	yes	no
Series 2 Defense Center: • DC500	yes	yes, but no Security Intelligence	yes, but no user control	no	no	yes
Series 2 Defense Centers: • DC1000, DC3000	yes	yes	yes	yes	yes	yes
Series 3 Defense Centers: • DC750, DC1500, DC3500, DC2000, DC4000	yes	yes	yes	yes	yes	yes
virtual Defense Centers	yes	yes	yes	yes	yes	yes

For detailed information on licensing, see the Licensing the FireSIGHT System chapter in the *FireSIGHT System User Guide*.

Security, Internet Access, and Communication Ports

To safeguard the Defense Center, you must install it on a protected internal network. Although the Defense Center is configured to have only the necessary services and ports available, you must make sure that attacks cannot reach it (or any managed devices) from outside the firewall.

If the Defense Center and its managed devices reside on the same network, you can connect the management interfaces on the devices to the same protected internal network as the Defense Center. This allows you to securely control the devices from the Defense Center. You can also configure multiple management interfaces to allow the Defense Center to manage and isolate traffic from devices on other networks.

Regardless of how you deploy your appliances, intra-appliance communication is encrypted. However, you must still take steps to ensure that communications between appliances cannot be interrupted, blocked, or tampered with; for example, with a distributed denial of service (DDoS) or man-in-the-middle attack.

Also note that specific features of the FireSIGHT System require an Internet connection. By default, all appliances are configured to directly connect to the Internet. Additionally, the system requires certain ports remain open for basic intra-appliance communication, for secure appliance access, and so that specific system features can access the local or Internet resources they need to operate correctly.

Tip With the exception of Cisco NGIPS for Blue Coat X-Series and Cisco ASA with FirePOWER Services, FireSIGHT System appliances support the use of a proxy server. For more information, see the *FireSIGHT System User Guide*.

For more information, see:

- Internet Access Requirements, page 1-14
- Communication Ports Requirements, page 1-15

Internet Access Requirements

Virtual Defense Centers are configured to directly connect to the Internet on ports 443/tcp (HTTPS) and 80/tcp (HTTP), which are open by default. On virtual devices, port 443 is open only if you enable a Malware license, so the device can submit files for dynamic analysis. For more information, see Communication Ports Requirements, page 1-15. FireSIGHT virtual appliances support use of a proxy server; for more information see the *FireSIGHT System User Guide*. Note also that a proxy server cannot be used for whois access.

The following table describes the Internet access requirements of specific features of the FireSIGHT System.

Table 1-5 FireSIGHT System Feature Internet Access Requirements

Feature	Internet Access is Required to...	Appliances
dynamic analysis: querying	query the Collective Security Intelligence Cloud for threat scores of files previously submitted for dynamic analysis.	Defense Center
dynamic analysis: submitting	submit files to the Collective Security Intelligence Cloud for dynamic analysis.	Managed devices
FireAMP integration	receive endpoint-based (FireAMP) malware events from the Collective Security Intelligence Cloud.	Defense Center
intrusion rule, VDB, and GeoDB updates	download or schedule the download of a intrusion rule, GeoDB, or VDB update directly to an appliance.	Defense Center

Table 1-5 FireSIGHT System Feature Internet Access Requirements (continued)

Feature	Internet Access is Required to...	Appliances
network-based AMP	perform malware cloud lookups.	Defense Center
RSS feed dashboard widget	download RSS feed data from an external source, including Cisco.	Any except virtual devices, X-Series, and ASA FirePOWER
Security Intelligence filtering	download Security Intelligence feed data from an external source, including the FireSIGHT System Intelligence Feed.	Defense Center
system software updates	download or schedule the download of a system update directly to an appliance.	Any except virtual devices, X-Series, and ASA FirePOWER
URL filtering	download cloud-based URL category and reputation data for access control, and perform lookups for uncategorized URLs.	Defense Center
whois	request whois information for an external host.	Any except virtual devices, X-Series, and ASA FirePOWER

Communication Ports Requirements

FireSIGHT System appliances communicate using a two-way, SSL-encrypted communication channel, which by default uses port 8305/tcp. The system **requires** this port remain open for basic intra-appliance communication. Other open ports allow:

- access to an appliance's web interface
- secure remote connections to an appliance
- certain features of the system to access the local or Internet resources they need to function correctly

In general, feature-related ports remain closed until you enable or configure the associated feature. For example, until you connect the Defense Center to a User Agent, the agent communications port (3306/tcp) remains closed. As another example, port 623/udp remains closed on Series 3 appliances until you enable LOM.

Caution Do not close an open port until you understand how this action will affect your deployment.

For example, closing port 25/tcp (SMTP) outbound on a manage device blocks the device from sending email notifications for individual intrusion events (see the *FireSIGHT System User Guide*). As another example, you can disable access to a physical managed device's web interface by closing port 443/tcp (HTTPS), but this also prevents the device from submitting suspected malware files to the Collective Security Intelligence Cloud for dynamic analysis.

Note that the system allows you to change some of its communication ports:

- You can specify custom ports for LDAP and RADIUS authentication when you configure a connection between the system and the authentication server; see the *FireSIGHT System User Guide*.

- You can change the management port (8305/tcp); see the *FireSIGHT System User Guide*. However, Cisco **strongly** recommends that you keep the default setting. If you change the management port, you must change it for all appliances in your deployment that need to communicate with each other.
- You can use port 32137/tcp to allow upgraded Defense Centers to communicate with the Collective Security Intelligence Cloud. However, Cisco recommends you switch to port 443, which is the default for fresh installations of Version 5.3 and later. For more information, see the *FireSIGHT System User Guide*.

The following table lists the open ports required by each appliance type so that you can take full advantage of FireSIGHT System features.

Table 1-6 Default Communication Ports for FireSIGHT System Features and Operations

Port	Description	Direction	Is Open on...	To...
22/tcp	SSH/SSL	Bidirectional	Any	allow a secure remote connection to the appliance.
25/tcp	SMTP	Outbound	Any	send email notices and alerts from the appliance.
53/tcp	DNS	Outbound	Any	use DNS.
67/udp 68/udp	DHCP	Outbound	Any except X-Series	use DHCP. **Note** These ports are **closed** by default.
80/tcp	HTTP	Outbound	Any except virtual devices, X-Series, and ASA FirePOWER	allow the RSS Feed dashboard widget to connect to a remote web server.
		Bidirectional	Defense Center	update custom and third-party Security Intelligence feeds via HTTP. download URL category and reputation data (port 443 also required).
161/udp	SNMP	Bidirectional	Any except X-Series and ASA FirePOWER	allow access to an appliance's MIBs via SNMP polling.
162/udp	SNMP	Outbound	Any	send SNMP alerts to a remote trap server.
389/tcp 636/tcp	LDAP	Outbound	Any except virtual devices and X-Series	communicate with an LDAP server for external authentication.
389/tcp 636/tcp	LDAP	Outbound	Defense Center	obtain metadata for detected LDAP users.
443/tcp	HTTPS	Inbound	Any except virtual devices, X-Series, and ASA FirePOWER	access the appliance's web interface.

Table 1-6 Default Communication Ports for FireSIGHT System Features and Operations (continued)

Port	Description	Direction	Is Open on...	To...
443/tcp	HTTPS AMQP cloud comms.	Bidirectional	Defense Center	obtain: • software, intrusion rule, VDB, and GeoDB updates • URL category and reputation data (port 80 also required) • the Collective Security Intelligence feed and other secure Security Intelligence feeds • endpoint-based (FireAMP) malware events • malware dispositions for files detected in network traffic • dynamic analysis information on submitted files
			Series 2 and Series 3 devices	download software updates using the device's local web interface.
			Series 3, virtual devices, X-Series, and ASA FirePOWER	submit files to for dynamic analysis.
514/udp	syslog	Outbound	Any	send alerts to a remote syslog server.
623/udp	SOL/LOM	Bidirectional	Series 3	allow you to perform Lights-Out Management using a Serial Over LAN (SOL) connection.
1500/tcp 2000/tcp	Inbound	TCP	Defense Center	allow read-only access to the database by a third-party client.
1812/udp 1813/udp	RADIUS	Bidirectional	Any except virtual devices, X-Series, and ASA FirePOWER	communicate with a RADIUS server for external authentication and accounting.
3306/tcp	User Agent	Inbound	Defense Center	communicate with User Agents.
8302/tcp	eStreamer	Bidirectional	Any except virtual devices and X-Series	communicate with an eStreamer client.
8305/tcp	device management	Bidirectional	Any	securely communicate between appliances in a deployment. **Required**.
8307/tcp	host input client	Bidirectional	Defense Center	communicate with a host input client.
32137/tcp	cloud comms.	Bidirectional	Defense Center	allow upgraded Defense Centers to communicate with the Collective Security Intelligence Cloud cloud.