Introduction to Network Discovery

The FireSIGHT System uses a feature called *network discovery* to monitor traffic on your network and build a comprehensive map of your network assets.

As managed devices passively observe traffic on the network segments you specify, the system compares specific packet header values and other unique data from network traffic against established definitions (called *fingerprints*) to determine the number and types of hosts (including network devices) on your network, as well as the operating systems, active applications, and open ports on those hosts.

You can also configure FireSIGHT System managed devices to monitor user activity on your network, which allows you to identify the source of policy breaches, attacks, or network vulnerabilities.

To supplement the data gathered by the system, you can import records generated by NetFlow-enabled devices, Nmap active scans, the host input feature, and User Agents that reside on a Microsoft Active Directory server and report LDAP authentications. The FireSIGHT System integrates these records with the information it collects via direct network traffic observation by managed devices.

The system can correlate certain types of intrusion, malware, and other events occurring on hosts on your network to determine when hosts are potentially compromised, tagging those hosts with *indications of compromise* (IOC) tags. IOC data can give you a clear, direct picture of the threats to your monitored network as they relate to its hosts.

The system uses all of this information to help you with forensic analysis, behavioral profiling, access control, and mitigating and responding to the vulnerabilities and exploits to which your organization is susceptible.

For more information, see:

- Understanding Discovery Data Collection, page 45-1
- Understanding NetFlow, page 45-16
- Understanding Indications of Compromise, page 45-19
- Creating a Network Discovery Policy, page 45-22

Understanding Discovery Data Collection

License: FireSIGHT

Discovery data includes information on your network’s hosts and the operating systems, active applications, and user activity on those hosts.

To begin collecting discovery data, you must first apply an access control policy. The access control policy defines the traffic that you permit, and therefore the traffic you can monitor with network discovery. Note that this means if you block certain traffic using access control, the system cannot examine that traffic for host, user, or application activity. For example, if you block access to social networking applications, the system does not provide you with any discovery data on social network applications.

After you apply an access control policy, you must configure and apply a network discovery policy, which specifies the network segments and ports you want to monitor with your managed devices, and the kinds of data you want to collect. When you apply the network discovery policy, the system begins generating discovery data, which you can then view and analyze using the Defense Center web interface.

The system stores network discovery data in the Defense Center database; for information on storage limits, see Configuring Database Event Limits, page 63-15. In addition to the database limits, the total number of detected hosts and users the Defense Center can store depends on your FireSIGHT license.

After you reach the licensed user limit, in most cases, the system stops adding new users to the database. To add new users, you must either manually delete old or inactive users from the database, or purge all users from the database. On the other hand, after you reach the licensed host limit, you can configure the system either to stop adding new hosts to the database, or to replace the hosts that have remained inactive for the longest time.

To supplement the data gathered by the system, you can import records generated by NetFlow-enabled devices, Nmap active scans, the host input feature, and User Agents that reside on a Microsoft Active Directory server and report LDAP authentications. The FireSIGHT System integrates these records with the information it collects via direct network traffic observation by managed devices.

For more information, see:

- Understanding Host Data Collection, page 45-2
- Understanding User Data Collection, page 45-3
- Understanding Application Detection, page 45-10
- Understanding Indications of Compromise, page 45-19
- Importing Third-Party Discovery Data, page 45-15
- Uses for Discovery Data, page 45-15

Understanding Host Data Collection

License: FireSIGHT

As the system passively monitors the traffic that travels through your network, it compares specific packet header values and other unique data from network traffic against established definitions (called *fingerprints*) to determine the following information about the hosts on your network, including:

- the number and types of hosts (including network devices such as bridges, routers, load balancers, and NAT devices)
- basic network topology data, including the number of hops from the discovery point on the network to the hosts
- the operating systems running on the hosts
- applications on the hosts and users associated with these applications

If the system cannot identify the operating system of a host, you can use the custom fingerprinting feature to create custom client or server fingerprints. The system uses these fingerprints to identify new hosts. You can map fingerprints to systems in the vulnerability database (VDB) to allow the appropriate vulnerability information to be displayed whenever a host is identified using the custom fingerprint. For more information, see Using Custom Fingerprinting, page 46-7.

You can also add or update host and operating system data through the host input feature. In addition, if you create a NetFlow-enabled discovery rule with host detection enabled, hosts can be added to the network map from NetFlow data.

You can view the hosts detected by the system using the Defense Center web interface:

- For information on viewing and searching for hosts using the event viewer, see Working with Hosts, page 50-18.
- For information on viewing the network map, which is a detailed representation of your network assets and topology, see Using the Network Map, page 48-1.
- For information on viewing host profiles, which are complete views of all the information available for your detected hosts, see Using Host Profiles, page 49-1.

Understanding User Data Collection

License: FireSIGHT

You can use the FireSIGHT System to monitor user activity on your network, which allows you to correlate threat, endpoint, and network intelligence with user identity information. By linking network behavior, traffic, and events directly to individual users, the system can help you to identify the source of policy breaches, attacks, or network vulnerabilities. In other words, the system can tell you the “who” behind the “what.” For example, you could determine:

- who owns the host targeted by an intrusion event that has a Vulnerable (level 1: red) impact level
- who initiated an internal attack or portscan
- who is attempting unauthorized access of a server that has high host criticality
- who is consuming an unreasonable amount of bandwidth
- who has not applied critical operating system updates
- who is using instant messaging software or peer-to-peer file-sharing applications in violation of company IT policy

Armed with this information, you can take a targeted approach to mitigate risk, block users or user activity, and take action to protect others from disruption. These capabilities also significantly improve audit controls and enhance regulatory compliance.

The system downloads the users used in access control policies from the Microsoft Active Directory LDAP server, based on the user awareness settings in the LDAP connection. The User Agent then provides login data for these users and the users are added to the user database. These users are referred to as *access-controlled users*. When you author access control policies that include user conditions, you write those conditions against access-controlled users. For more information, see Adding a User Condition to an Access Control Rule, page 17-2.

When the system detects user data from a user login, either from a User Agent, from application data detected in traffic, or from an email login over POP3, SMTP, or IMAP, the user from the login is checked against the list of users. If the login user matches an existing user reported by an agent, the data from the login is assigned to the user. Logins that do not match existing users cause a new user to be created, unless the login is in SMTP traffic. Non-matching logins in SMTP traffic are discarded.

The following diagram illustrates how the FireSIGHT System collects and stores user data.

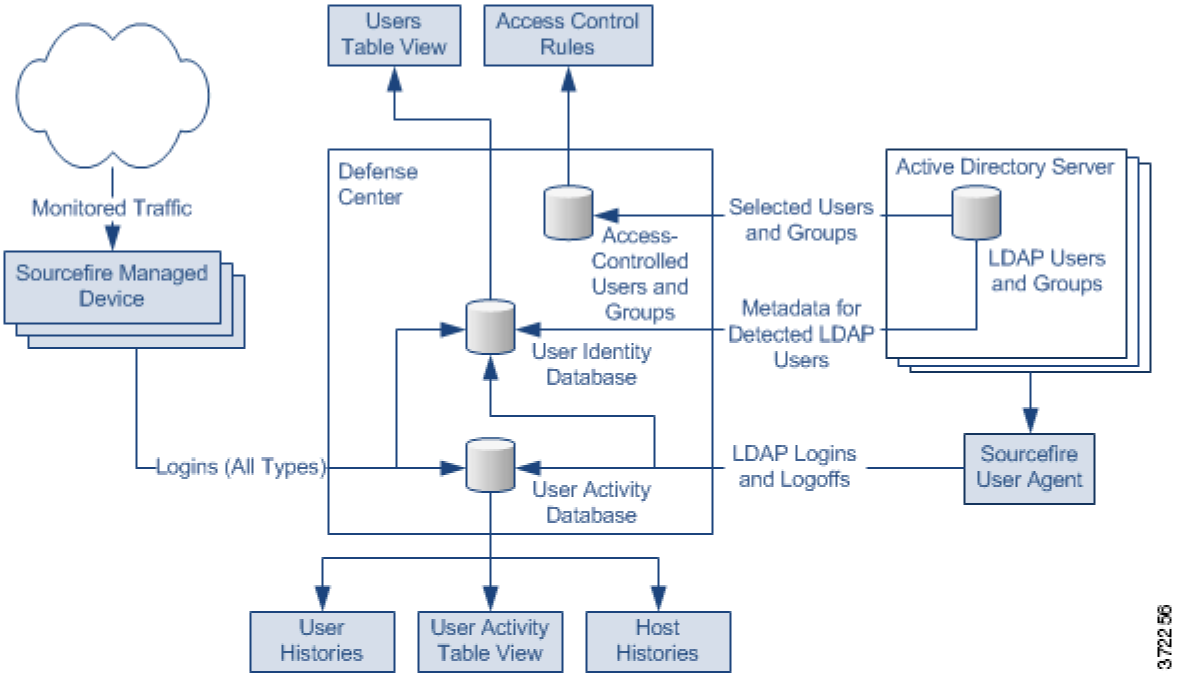

As shown in the diagram, there are three sources for user data, and three places that data is stored. For more information on user data collection, see:

- Managed Devices, page 45-4
- User Agents, page 45-5
- Defense Center-LDAP Server Connections, page 45-7
- Users Database, page 45-7
- User Activity Database, page 45-7
- Access-Controlled Users Database, page 45-8
- User Data Collection Limitations, page 45-9

Managed Devices

License: FireSIGHT

You use the network discovery policy to configure managed devices to passively detect LDAP, AIM, POP3, IMAP, Oracle, SIP (VoIP), FTP, HTTP, MDNS, and SMTP logins on the networks you specify. Note that when you enable discovery of users in a network discovery rule, host discovery is automatically enabled.

Note Managed devices interpret only Kerberos logins for LDAP connections as LDAP authentications. Managed devices cannot detect encrypted LDAP authentications using protocols such as SSL or TLS.

When a device detects a login, it sends the following information to the Defense Center to be logged as user activity:

- the user name identified in the login

- the time of the login
- the IP address involved in the login, which can be the IP address of the user's host (for LDAP, POP3, IMAP, and AIM logins), the server (for HTTP, MDNS, FTP, SMTP and Oracle logins), or the session originator (for SIP logins)
- the user's email address (for POP3, IMAP, and SMTP logins)
- the name of the device that detected the login

If the user was previously detected, the Defense Center updates that user's login history. Note that the Defense Center can use the email addresses in POP3 and IMAP logins to correlate with LDAP users. This means that, for example, if the Defense Center detects a new IMAP login, and the email address in the IMAP login matches that for an existing LDAP user, the IMAP login does not create a new user, rather, it updates the LDAP user's history.

If the user has never been detected before, the Defense Center adds the user to the users database. Unique AIM, SIP, and Oracle logins always create new user records, because there is no data in those login events that the Defense Center can correlate with other login types.

The Defense Center does **not** log user activity or user identities in the following cases:

- if you configured the network discovery policy to ignore that login type, as described in Restricting User Logging, page 45-29
- if a managed device detects an SMTP login, but the users database does not contain a previously detected LDAP, POP3, or IMAP user with a matching email address

User Agents

License: FireSIGHT

If your organization uses Microsoft Active Directory LDAP servers, Cisco recommends that you install User Agents to monitor user activity via your Active Directory servers. If you want to perform user control, you **must** install and use User Agents; the agents associate users with IP addresses, which in turn allows access control rules with user conditions to trigger. You can use one agent to monitor user activity on up to five Active Directory servers.

To use an agent, you must configure a connection between each Defense Center connected to the agent and the monitored LDAP servers. This connection not only allows you to retrieve metadata for the users whose logins and logoffs were detected by User Agents, but also is used to specify the users and groups you want to use in access control rules. For more information on configuring LDAP servers for user discovery, see Retrieving Access-Controlled Users and LDAP User Metadata, page 17-4.

Each agent can monitor logins using encrypted traffic, either through regularly scheduled polling or real-time monitoring. Logins are generated by the Active Directory server when a user logs into a computer, whether at the workstation or through a Remote Desktop login.

Agents can also monitor and report user logoffs. Logoffs are generated by the agent itself when it detects a user logged out of a host IP address. Logoffs are also generated when the agent detects that the user logged into a host has changed, before the Active Directory server reports that the user has changed. Combining logoff data with login data develops a more complete view of the users logged into the network.

Polling an Active Directory server allows an agent to retrieve batches of user activity data at the defined polling interval. Real-time monitoring transmits user activity data to the agent as soon as the Active Directory server receives the data.

You can configure the agent to exclude reporting any logins or logoffs associated with a specific user name or IP address. This can be useful, for example, to exclude repeated logins to shared servers, such as file shares and print servers, as well as exclude users logging into machines for troubleshooting purposes.

The agents send records of all detected logins and logoffs that do not contain an excluded user name or IP address to Defense Centers, which log and report them as user activity. The agents detect the Defense Center version and send the login records in the appropriate data format. This supplements any user activity detected directly by managed devices. The logins reported by User Agents associate users with IP addresses, which in turn allows access control rules with user conditions to trigger.

User Agents monitor users as they log into the network or when accounts authenticate against Active Directory credentials for other reasons. Version 2.1 of the User Agent detects interactive user logins to a host, Remote Desktop logins, file-share authentication, and computer account logins, as well as user logoffs and Remote Desktop sessions where the user has logged off.

The type of login detected determines how the agent reports the login and how the login appears in the host profile. An *authoritative user login* for a host causes the current user mapped to the host IP address to change to the user from the new login. Other logins either do not change the current user or only change the current user for the host if the existing user on the host did not have an authoritative user login to the host. In these cases, if the expected user is no longer logged in, the agent generates a logoff for that user. User logins detected by network discovery only change the current user for the host if the existing user on the host did not have an authoritative user login to the host. Agent-detected logins have the following effect on the network map:

- When the agent detects an interactive login to a host by a user or a Remote Desktop login, the agent reports an authoritative user login for the host and changes the current user for the host to the new user.
- If the agent detects a login for file-share authentication, the agent reports a user login for the host, but does not change the current user on the host.
- If the agent detects a computer account login to a host, the agent generates a NetBIOS Name Change discovery event and the host profile reflects any change to the NetBIOS name.
- If the agent detects a login from an excluded user name, the agent does not report a login to the Defense Center.

When a login or other authentication occurs, the agent sends the following information to the Defense Center:

- the user’s LDAP user name
- the time of the login or other authentication
- the IP address of the user’s host, and the link-local address if the agent reports an IPv6 address for a computer account login

The Defense Center records login and logoff information as user activity. When a User Agent reports user data from a user login or logoff, the reported user is checked against the list of users. If the reported user matches an existing user reported by an agent, the reported data is assigned to the user. Reported users that do not match existing users cause a new user to be created.

Even though the user activity associated with an excluded user name is not reported, related user activity may still be reported. If the agent detects a user login to a machine, then the agent detects a second user login, and you have excluded the user name associated with the second user login from reporting, the agent reports a logoff for the original user. However, no login for the second user is reported. As a result, no user is mapped to the IP address, even though the excluded user is logged into the host.

Note the following limitations on user names detected by the agent:

- User names ending with a dollar sign character ($) reported to the Defense Center update the network map, but do not appear as user logins.
- Defense Center display of user names containing Unicode characters may have limitations.

The total number of detected users the Defense Center can store depends on your FireSIGHT license. After you reach the licensed user limit, in most cases the system stops adding new users to the database. To add new users, you must either manually delete old or inactive users from the database, or purge all users from the database.

Defense Center-LDAP Server Connections

License: FireSIGHT

The Defense Center-LDAP server connection allows you to retrieve metadata for certain detected users. You can retrieve metadata for LDAP users, whether their logins were detected by managed devices or by a User Agent; you can also retrieve metadata for POP3 and IMAP users if those users have the same email address as an LDAP user.

If your organization uses Microsoft Active Directory servers, the connection also allows you to specify the LDAP users and groups you want to use in access control rules. If you want to perform user control, you **must** configure a connection between the Defense Center and an Active Directory server. If your organization does not use Active Directory, you can still detect user logins using managed devices, and you can still obtain metadata for some of those users from an Oracle or OpenLDAP server. However, you cannot perform user control based on those users or their activity.

From the LDAP server, the Defense Center obtains the following information and metadata about each user:

- LDAP user name
- first and last names
- email address
- department
- telephone number

Users Database

License: FireSIGHT

The users database contains a record for each user detected by either managed devices or User Agents. The total number of detected users the Defense Center can store depends on your FireSIGHT license. After you reach the licensed limit, in most cases the system stops adding new users to the database. To add new users, you must either manually delete old or inactive users from the database, or purge all users from the database.

However, the system favors authoritative user logins. If you have reached the limit and the system detects an authoritative user login for a previously undetected user, the system deletes the non-authoritative user who has remained inactive for the longest time, and replaces it with the new user.

You can view the contents of the users database with the Defense Center web interface. For information on viewing, search for, and deleting detected users, see Working with Users, page 50-58.

User Activity Database

License: FireSIGHT

The user activity database contains records of user activity on your network, either from a connection to an Active Directory LDAP server that is also monitored by a User Agent, or though network discovery. The system logs events in the following circumstances:

- when it detects individual logins or logoffs
- when it detects a new user
- when you manually delete a user
- when the system detects a user that is not in the database, but cannot add the user because you have reached your FireSIGHT licensed limit

You can view the user activity detected by the system using the Defense Center web interface. For information on viewing, searching for, and deleting user activity, see Working with User Activity, page 50-64. If you plan to use Version 2.1 of the User Agent to send LDAP login data to your Defense Centers, you must configure a connection for each agent on each Defense Center where you want the agent to connect. That connection allows the agent to establish a secure connection with the Defense Center, over which it can send login data. If the agent is configured to exclude specific user names, login data for those user names are not reported to the Defense Center.

In addition, if you are planning to implement user access control, you must set up a connection to each Microsoft Active Directory server where you plan to collect data, with user awareness parameters configured.

Whenever possible the FireSIGHT System correlates user activity with other types of events. For example, intrusion events can tell you the users who were logged into the source and destination hosts at the time of the event.

The system also uses user activity to generate *host histories*, which track the hosts that each user has logged into, and *user histories*, which track the users that have logged into each individual host. The system provides a graphical representation of the last twenty-four hours of each user's activity and the last twenty-four hours of the logins to each host. For more information, see Understanding User Details and Host History, page 50-62 and Working with User History in the Host Profile, page 49-22.

Access-Controlled Users Database

License: Control

The access-controlled users database contains the users and groups that you can use in access control rules, so that you can perform user control with the FireSIGHT System. These users can be one of two types:

- An *access-controlled user* is a user that you can add to access control rules to perform user control. You specify the groups that access-controlled users must belong to when you configure the Defense Center-LDAP server connection.
- A *non-access-controlled user* is any other detected user.

You specify the groups that access-controlled users must belong to when you configure the Defense Center-LDAP server connection, as described in Retrieving Access-Controlled Users and LDAP User Metadata, page 17-4.

If you plan to use Version 2.1 of the User Agent to send LDAP login and logoff data to your Version 5.x Defense Centers, you must configure a connection for each agent on each Defense Center where you want the agent to connect. That connection allows the agent to establish a secure connection with the Defense Center, over which it can send the user activity data.

If the agent is configured to exclude specific user names, user activity data for those user names are not reported to the Defense Center. These excluded user names remain in the database, but are not associated with IP addresses.

In addition, if you are planning to implement user access control, you must set up a connection to each Microsoft Active Directory server where you plan to collect data, with user awareness parameters configured.

The maximum number of users you can use in access control depends on your FireSIGHT license. When configuring the Defense Center-LDAP server connection, make sure the total number of users you include is less than your FireSIGHT user license. See Understanding FireSIGHT Host and User License Limits, page 65-8 for more information.

User Data Collection Limitations

License: FireSIGHT

The following table describes the limitations of user data collection.

Table 45-1 User Awareness Limitations

Limitation	Description
user control	To perform user control, your organization **must** use Microsoft Active Directory LDAP servers. The system obtains the users and groups you can use in access control rules from Active Directory, and also ties users to IP addresses with the logins and logoffs reported by User Agents installed on Active Directory servers.
non-Kerberos logins for LDAP connections	Managed devices interpret only Kerberos logins for LDAP connections as LDAP authentications. Managed devices cannot detect encrypted LDAP authentications if they use other protocols, such as SSL or TLS. On the other hand, User Agents use the security logs on Active Directory servers to collect user login data and have no such limitations.
login detection	If you want to detect logins to an Active Directory server, you must configure the Active Directory server connection with the server IP address. See the *User Agent Configuration Guide* for more information. If multiple users are logged into a host using remote sessions, the agent may not detect logins from that host properly. See the *User Agent Configuration Guide* for more information on how to prevent this.
logoff detection	Logoffs may not be immediately detected. The timestamp associated with a logoff reflects when the agent detected the user was no longer mapped to the host IP address, which may not correspond to the actual time the user logged off of the host. Logoffs are generated by the agent itself when it detects a user logged out of a host IP address. Logoffs are also generated when the agent detects that the user logged into a host has changed, before the Active Directory server reports that the user has changed.
real-time data retrieval	The Active Directory server must be running Windows Server 2008 or Windows Server 2012.
multiple logins to the same host by different users	The system assumes that only one user is logged into any given host at a time, and that the current user of a host is the last authoritative user login. If only non-authoritative logins have been logged into the host, the last non-authoritative login is considered the current user. If multiple users are logged in through remote sessions, the last user reported by the Active Directory server is the user reported to the Defense Center.

Table 45-1 User Awareness Limitations (continued)

Limitation	Description
multiple logins to the same host by the same user	The system records the first time that a user logs into a specific host and disregards subsequent logins. If an individual user is the only person who logs into a specific host, the only login that the system records is the original login. If another user logs into that host, however, the system records the new login. Then, if the original user logs in again, his or her new login is recorded.
Unicode characters	The user interface may not correctly display user names with Unicode characters.
LDAP user accounts in the users database	If you remove or disable an LDAP user on your LDAP servers, or exclude the user name from being reported to the Defense Center, the Defense Center does not remove that user from the users database, and that user continues to count against your licensed limit for user listed in the database. You must manually purge the user from the database. Note that the user license limit is applied in parallel for access-controlled users; the user count for access-controlled users depends on the number of users retrieved by your LDAP configuration.
AOL Instant Messenger (AIM) login detection	Managed devices can detect AIM logins using the OSCAR protocol only. While most AIM clients use OSCAR, some use TOC2.

Understanding Application Detection

License: FireSIGHT

When the FireSIGHT System analyzes IP traffic, it attempts to identify the commonly used applications on your network. Application awareness is crucial to performing application-based access control.

There are three types of applications that the system detects:

- *application protocols* such as HTTP and SSH, which represent communications between hosts
- *clients* such as web browsers and email clients, which represent software running on the host
- *web applications* such as MPEG video and Facebook, which represent the content or requested URL for HTTP traffic

The system identifies applications in your network traffic either using ASCII or hexadecimal patterns in the packet headers, or the port that the traffic uses. Some application detectors use both port and pattern detection to increase the likelihood of correctly identifying traffic for a particular application. In addition, Secure Socket Layers (SSL) protocol detectors use information from the secured session to identify the application from the session. There are two sources of application detectors in the FireSIGHT System:

- *Cisco-provided detectors*, which detect web applications, clients, and application protocols

  The availability of Cisco-provided detectors for applications (and operating systems, see Understanding Host Data Collection, page 45-2) depend on the version of the FireSIGHT System and the version of the VDB you have installed. Release notes and advisories contain information on new and updated detectors. You can also import individual detectors authored by Professional Services. For a complete list of detected applications, see the Support site.

- *user-defined application protocol detectors*, which you can create to enhance the system's application protocol detection capabilities

You can also detect application protocols through *implied application protocol detection*, which implies the existence of an application protocol based on the detection of a client.

The system characterizes each application that it detects using the criteria described in the following table. The system uses these characteristics to create application filters, or groups of applications. You can use these filters and filters that you create to perform access control, as well as to constrain searches, reports, and dashboard widgets. For more information, see Working with Application Filters, page 3-14.

:

Table 45-2 Application Characteristics

Characteristic	Description	Example
Type	The type of application: • **Application Protocols** represent communications between hosts. • **Clients** represent software running on a host. • **Web Applications** represent the content or requested URL for HTTP traffic.	HTTP and SSH are application protocols. Web browsers and email clients are clients. MPEG video and Facebook are web applications.
Risk	How likely the application is to be used for purposes that might be against your organization's security policy. An application's risk can range from **Very Low** to **Very High**.	Peer-to-peer applications tend to have a very high risk.
Business Relevance	The likelihood that the application is used within the context of your organization's business operations, as opposed to recreationally. An application's business relevance can range from **Very Low** to **Very High**.	Gaming applications tend to have a very low business relevance.
Category	A general classification for the application that describes its most essential function. Each application belongs to at least one category.	Facebook is in the **social networking** category.
Tag	Additional information about the application. Applications can have any number of tags, including none.	Video streaming web applications often are tagged **high bandwidth** and **displays ads**.

To supplement the application data gathered by the system, you can use records generated by NetFlow-enabled devices, Nmap active scans, and the host input feature.

For more information, see:

- Understanding the Application Protocol Detection Process, page 45-11
- Implied Application Protocol Detection from Client Detection, page 45-13
- Special Considerations for Application Protocol Detection: Squid, page 45-14
- Special Considerations: SSL Application Detection, page 45-14
- Special Considerations: Referred Web Applications, page 45-14
- Working with Application Detectors, page 46-17
- Importing Third-Party Discovery Data, page 45-15
- Understanding NetFlow, page 45-16

Understanding the Application Protocol Detection Process

License: FireSIGHT

When the system detects application traffic, it first determines whether the application protocol is running on a port identified by a detector that uses that specific port as its only detection criterion. If the application protocol is running on one of those ports, the system positively identifies the application protocol using the well-known port detector.

Note Because you can create and activate user-defined port-based application protocol detectors on ports used by Cisco-provided detectors, it is possible to override Cisco's detection capabilities. For example, if your user-defined detector identifies all application protocol traffic on port 22 as the `myapplication` application protocol, SSH traffic on port 22 will be misidentified as `myapplication` traffic.

If the application protocol is not running on one of those ports, the system employs a more robust method to identify it based on port and pattern matches. If two detectors both positively identify the traffic, the detector that employs the longer pattern match has precedence. Similarly, detectors with multiple pattern matches have precedence over single pattern matches.

Note that the system identifies only those application protocols running on hosts in your monitored networks, as defined in the network discovery policy. For example, if an internal host accesses an FTP server on a remote site that you are not monitoring, the system does not identify the application protocol as FTP. On the other hand, if a remote or internal host accesses an FTP server on a host you are monitoring, the system can positively identify the application protocol.

An exception occurs if the system can identify the client used in connections between a monitored host accessing a non-monitored server. In that case, the system positively identifies the appropriate application protocol that corresponds with the client in the connection, but does not add the application protocol to the network map. For more information, see Implied Application Protocol Detection from Client Detection, page 45-13. Note that client sessions must include a response from the server for application detection to occur.

The following table outlines how the FireSIGHT System identifies detected application protocols in the Defense Center web interface: the network map, host profiles, event views, and so on.

Table 45-3 FireSIGHT System Identification of Application Protocols

Application	Description
the application protocol name	The Defense Center identifies an application protocol with its name if the application protocol was: • positively identified by the system • identified using NetFlow data and there is a port-application protocol correlation in `/etc/sf/services` • manually identified using the host input feature • identified by Nmap or another active source
`pending`	The Defense Center identifies an application protocol as `pending` if the system can neither positively nor negatively identify the application. Most often, the system needs to collect and analyze more connection data (from which applications are identified) before it can identify a pending application. In the Application Details and Servers tables and in the host profile, the `pending` status appears only for application protocols where specific application protocol traffic was detected (rather than implied by detected client or web application traffic).

Table 45-3 FireSIGHT System Identification of Application Protocols (continued)

Application	Description
`unknown`	The Defense Center identifies an application protocol as `unknown` if the application: • does not match any of the system's detectors • the application protocol was identified using NetFlow data, but there is no port-application protocol correlation in `/etc/sf/services`
blank	All available detected data has been examined and no application protocol was identified. In the Application Details and Servers tables and in the host profile, the application protocol is left blank for non-HTTP generic client traffic with no detected application protocol.

Implied Application Protocol Detection from Client Detection

License: FireSIGHT

If the system can identify the client used in a connection between a monitored host accessing a non-monitored server, the Defense Center infers that the connection is using the application protocol that corresponds with the client. (Because the system tracks applications only on monitored networks, connection logs usually do not include application protocol information for connections where a monitored host is accessing a non-monitored server.)

There are several consequences of the implied detection of an application protocol from the detection of a client:

- Because the system does not generate a New TCP Port or New UDP Port event for these servers, the server does not appear in the Servers table. In addition, you cannot trigger either discovery event alerts or correlation rules using the detection of these application protocol as a criterion.
- Because the application protocol is not associated with a host, you cannot view its details in host profiles, set its server identity, or use its information in host profile qualifications for traffic profiles or correlation rules. In addition, the system does not associate vulnerabilities with hosts based on this type of detection.

You can, however, trigger correlation events on the application protocol information in a connection. You can also use the application protocol information in connection logs to create connection trackers and traffic profiles.

Host Limits and Discovery Event Logging

License: FireSIGHT

When the system detects a client, server, or web application, it generates a discovery event unless the associated host has already reached its maximum number of clients, servers, or web applications.

Host profiles display up to 16 clients, 100 servers, and 100 web applications per host. See Working with Servers in the Host Profile, page 49-15 and Viewing Applications in the Host Profile, page 49-20 for more information.

Note that actions dependent on the detection of clients, servers, or web applications are unaffected by this limit. For example, access control rules configured to trigger on a server will still log connection events.

Special Considerations for Application Protocol Detection: Squid

License: FireSIGHT

The system positively identifies Squid server traffic when either:

- the system detects a connection from a host on your monitored network to a Squid server where proxy authentication is enabled, or
- the system detects a connection from a Squid proxy server on your monitored network to a target system (that is, the destination server where the client is requesting information or another resource)

However, the system cannot identify Squid service traffic if:

- a host on your monitored network connects to a Squid server where proxy authentication is disabled, or
- the Squid proxy server is configured to strip Via: header fields from its HTTP responses

Special Considerations: SSL Application Detection

License: FireSIGHT

The FireSIGHT System provides detectors that can use session information from a Secure Socket Layers (SSL) session to identify the application protocol, client application, or web application in the session.

When the system detects an encrypted connection, it marks that connection as either a generic HTTPS connection or as a more specific secure protocol, such as SMTPS, when applicable. When the system detects an SSL session, it adds `SSL client` to the **Client** field in connection events for the session. If it identifies a web application for the session, the system generates discovery events for the traffic.

For SSL application traffic, managed devices can also detect the common name from the server certificate and match that against a client or web application from an SSL host pattern. When the system identifies a specific client, it replaces `SSL client` with the name of the client.

Because the SSL application traffic is encrypted, the system can use only information in the certificate for identification, not application data within the encrypted stream. For this reason, SSL host patterns can sometimes only identify the company that authored the application, so SSL applications produced by the same company may have the same identification.

In some instances, such as when an HTTPS session is launched from within an HTTP session, managed devices detect the server name from the client certificate in a client-side packet.

To enable SSL application identification, you must create access control rules that monitor responder traffic. Those rules must have either an application condition for the SSL application or URL conditions using the URL from the SSL certificate. For network discovery, the responder IP address does not have to be in the networks to monitor in the network discovery policy; the access control policy configuration determines whether the traffic is identified. You can filter by the `SSL protocol` tag, in the application detectors list or when adding application conditions in access control rules, to identify detectors for SSL applications.

Special Considerations: Referred Web Applications

Web servers sometimes refer traffic to other websites, which are often advertisement servers. To help you better understand the context for referred traffic occurring on your network, the system lists the web application that referred the traffic in the Web Application field in events for the referred session. The VDB contains a list of known referred sites. When the system detects traffic from one of those sites, the referring site is stored with the event for that traffic. For example, if an advertisement accessed via Facebook is actually hosted on Advertising.com, the detected Advertising.com traffic is associated with the Facebook web application. The system can also detect referring URLs in HTTP traffic, such as when a website provides a simple link to another site; in this case, the referring URL appears in the HTTP Referrer event field.

In events, if a referring application exists, it is listed as the web application for the traffic, while the URL is that for the referred site. In the example above, the web application for the connection event for that traffic would be Facebook, but the URL would be Advertising.com. If no referring web application is detected, if the host refers to itself, or if there is a chain of referrals, a referred application may appear as the web application in the event. In the dashboard, connection and byte counts for web applications include sessions where the web application is associated with traffic referred by that application.

Note that if you create a rule to act specifically on referred traffic, you should add a condition for the referred application, rather than the referring application. To block Advertising.com traffic referred from Facebook, for example, add an application condition to your access control rule for the Advertising.com application.

Importing Third-Party Discovery Data

License: FireSIGHT

You can use Nmap active scans to add information about operating systems, applications, and vulnerabilities, supplementing the data gathered by the system. For more information on Nmap scanning and scan results, see Understanding Nmap Scans, page 47-1.

You can also use the host input feature to supplement the information that the system gathers from monitoring network traffic, either by configuring a third-party application to interact with the FireSIGHT System via an API, or by manually adding data. You can create product, vulnerability, and fix mappings to map third-party data to Cisco definitions, enabling impact correlation for operating systems and servers. For more information on the host input feature and mapping third-party data, see the *FireSIGHT System Host Input API Guide* and Importing Host Input Data, page 46-29.

The system reconciles the collected data about operating system and server identities and determines each identity based on fingerprint source priority values, identity conflict resolution settings, and time of collection.

You can also configure your network map to use data from NetFlow-enabled devices to enhance your network map and event tables. For more information, see Understanding NetFlow, page 45-16.

Uses for Discovery Data

License: FireSIGHT

Logging discovery data allows you to take advantage of many features in the FireSIGHT System, including:

- viewing the network map, which is a detailed representation of your network assets and topology that you can view by grouping hosts and network devices, host attributes, application protocols, or vulnerabilities; see Using the Network Map, page 48-1
- viewing host profiles, which are complete views of all the information available for your detected hosts; see Using Host Profiles, page 49-1
- viewing dashboards, which (among other capabilities) can provide you with an at-a-glance view of your network assets and user activity; see Using Dashboards, page 55-1
- viewing detailed information on the discovery events and user activity logged by the system; see Working with Discovery Events, page 50-1

- creating reports based on discovery data; see Working with Reports, page 57-1
- performing application and user control, that is, writing access control rules using application and user conditions; see Controlling Application Traffic, page 16-2 and Adding a User Condition to an Access Control Rule, page 17-2
- associating hosts and any servers or clients they are running with the exploits to which they are susceptible, which allows you to identify and mitigate vulnerabilities, evaluate the impact that intrusion events have on your network, and tune intrusion rule states so that they provide maximum protection for your network assets; see Working with Vulnerabilities in the Host Profile, page 49-25, Using Impact Levels to Evaluate Events, page 41-37, Understanding Indications of Compromise, page 45-19, and Tailoring Intrusion Protection to Your Network Assets, page 33-1
- alerting you via email, SNMP trap, or syslog when the system generates either an intrusion event with a specific impact flag, or a specific type of discovery event; see Configuring External Alerting, page 43-1
- monitor your organization’s compliance with a white list of allowed operating systems, clients, application protocols, and protocols; see Using the FireSIGHT System as a Compliance Tool, page 52-1
- creating correlation policies with rules that trigger and generate correlation events when the system generates discovery events or detects user activity; see Configuring Correlation Policies and Rules, page 51-1
- if you log NetFlow connections, using that connection data; see Logging Connections to the Defense Center or External Server, page 38-5

Understanding NetFlow

License: FireSIGHT

NetFlow is an embedded instrumentation within Cisco IOS Software that characterizes network operation. Standardized through the RFC process, NetFlow is available not only on Cisco networking devices, but can also be embedded in Juniper, FreeBSD, and OpenBSD devices.

NetFlow-enabled devices are widely used to capture and export data about the traffic that passes through those devices. NetFlow-enabled devices have a database called the NetFlow cache that stores records of the flows that pass through the devices. A flow, called a *connection* in the FireSIGHT System, is a sequence of packets that represents a session between a source and destination host, using specific ports, protocol, and application protocol.

For the networks you specify, FireSIGHT System managed devices detect the records exported by NetFlow-enabled devices, generate connection events based on the data in those records, and finally send those events to the Defense Center to be logged in the database. You can also configure the system to add host and application protocol information to the database, based on the information in NetFlow connections.

You can use this discovery and connection data to supplement the data gathered directly by your managed devices. This is especially useful if you have NetFlow-enabled devices deployed on networks that your managed devices cannot monitor.

You configure NetFlow data collection, including connection logging, using rules in the network discovery policy. Contrast this with connection logging for connections detected by FireSIGHT System managed devices, which you configure per access control rule, as described in Logging Connections Based on Access Control Handling, page 38-15. Because NetFlow data collection is linked to networks rather than access control rules, you do not have as much granular control over which connections you want to log, Also, the system automatically saves all NetFlow-based connection events to the Defense Center connection event database; you cannot send them to the system log or an SNMP trap server.

For more information, see:

- Differences Between NetFlow and FireSIGHT Data, page 45-17
- Preparing to Analyze NetFlow Data, page 45-18
- Uses for Discovery Data, page 45-15
- Logging Connections to the Defense Center or External Server, page 38-5

Differences Between NetFlow and FireSIGHT Data

License: FireSIGHT

With one exception (TCP flags), the information available in NetFlow records is more limited than the information generated by monitoring network traffic using managed devices. Because the system cannot directly analyze the traffic represented by NetFlow data, when the system processes NetFlow records it uses various methods to convert that data into connection logs as well as into host and application protocol records.

There are several differences between converted NetFlow data and the discovery and connection data gathered directly by your managed devices. You should keep the differences in mind when performing analysis that requires:

- statistics on the number of detected connections
- operating system and other host-related information (including vulnerabilities)
- application data, including client information, web application information, and vendor and version server information
- knowing which host in a connection is the initiator and which is the responder

Tip For each field in a connection event, Table 39-1 on page 39-12 indicates the available data depending on whether the connection was detected directly by FireSIGHT System managed devices, or if the connection event is based on NetFlow data.

Number of Connection Events Generated Per Monitored Session

For connections detected directly by managed devices, depending on the access control rule action, you can log a bidirectional connection event at the beginning or end of a connection, or both.

However, because NetFlow-enabled devices export unidirectional connection data, the system always generates at least two connection events for each connection detected by NetFlow-enabled devices, depending on how you configured the devices. This also means that a summary’s connection count is incremented by two for every connection based on NetFlow data, providing an inflated count of the number of connections that are actually occurring on your network.

Note that if you configure your NetFlow-enabled devices to output records only when the connection ends, the system generates two connection events for that session. On the other hand, if you configure your NetFlow-enabled devices to output records at a fixed interval even if a connection is still ongoing, the system generates a connection event for each record exported by the device. For example, if you configure your NetFlow-enabled devices to output records for long-running connections every five minutes, and a particular connection lasts twelve minutes, the system generates six connection events for that session:

- one pair of events for the first five minutes
- one pair for the second five minutes
- a final pair when the connection is terminated

For this reason. Cisco **strongly** recommends that you configure your NetFlow-enabled devices to output records only when monitored sessions close.

Host and Operating System Data

Although you can configure the network discovery policy to add hosts to the network map based on NetFlow records, the host profile does not include any operating system or NetBIOS data for the hosts involved in the connection, nor can the system identify if the hosts are network devices (bridges, routers, NAT devices, or load balancers). You can, however, manually set a host's operating system identity using the host input feature.

Application Data

For connections detected directly by managed devices, the system can identify application protocols, clients, and web applications by examining the packets in the connection.

When the system processes NetFlow records, the system uses a port correlation in `/etc/sf/services` to extrapolate application protocol identity. However, there is no vendor or version information for those application protocols, nor do connection logs contain information on client or web applications used in the session. You can, however, manually provide this information using the host input feature.

Note that a simple port correlation means that application protocols running on non-standard ports may be unidentified or misidentified. Additionally, if no correlation exists, the system marks the application protocol as `unknown` in connection logs.

Vulnerability Mappings

The FireSIGHT System cannot determine which vulnerabilities might affect hosts added to the network map based on NetFlow records, unless you use the host input feature to manually set either a host's operating system identity or an application protocol identity. Note that because there is no client information in NetFlow connections, you cannot associate client vulnerabilities with NetFlow hosts.

Initiator and Responder Information in Connections

For connections detected directly by managed devices, the system can identify which host is the initiator, or source, and which is the responder, or destination. However, NetFlow data does not contain initiator or responder information.

When the system processes NetFlow records, it uses an algorithm to determine this information based on the ports each host is using, and whether those ports are well-known:

- If both or neither port being used is a well-known port, the system considers the host using the lower-number port to be the responder.
- If only one of the hosts is using a well-known port, the system considers that host to be the responder.

For this purpose, a well-known port is any port that is either numbered from 1 to 1023, or that contains application protocol information in `/etc/sf/services` on the managed device.

Preparing to Analyze NetFlow Data

License: FireSIGHT

Before you configure the FireSIGHT System to analyze NetFlow data, you must enable the NetFlow feature on the routers or other NetFlow-enabled devices you plan to use, and configure the devices to export NetFlow version 5 data to a destination network where the sensing interface of a managed device is connected.

Note that the system can parse both NetFlow version 5 and NetFlow version 9 records. Your NetFlow-enabled devices **must** use one of those versions if you want to use them with your FireSIGHT System deployment. In addition, the system requires that specific fields be in the templates and records that your NetFlow-enabled devices broadcast. If your NetFlow-enabled devices are using version 9, which you can customize, you **must** make sure that the templates and records that the devices broadcast contain the following fields, in any order:

- IN_BYTES (1)
- IN_PKTS (2)
- PROTOCOL (4)
- TCP_FLAGS (6)
- L4_SRC_PORT (7)
- IPV4_SRC_ADDR (8)
- L4_DST_PORT (11)
- IPV4_DST_ADDR (12)
- LAST_SWITCHED (21)
- FIRST_SWITCHED (22)
- IPV6_SRC_ADDR (27)
- IPV6_DST_ADDR (28)

Because the FireSIGHT System uses managed devices to analyze NetFlow data, your deployment must include at least one managed device that can monitor your NetFlow-enabled devices. At least one sensing interface on that managed device must be connected to a network where it can collect the data that your NetFlow-enabled devices export. Because the sensing interfaces on managed devices do not usually have IP addresses, the system does not support the direct collection of NetFlow records.

In addition, Cisco **strongly** recommends that you configure your NetFlow-enabled devices to output records only when monitored sessions close. If you configure your NetFlow-enabled devices to output records at fixed intervals, analyzing the connection data derived from the NetFlow records may be more complicated; see Number of Connection Events Generated Per Monitored Session, page 45-17.

Finally, note that the Sampled NetFlow feature available on some NetFlow-enabled devices collects NetFlow statistics on only a subset of packets that pass through the devices. Although enabling this feature can improve CPU utilization on the NetFlow-enabled device, it may affect the data you are collecting for analysis by the system.

Understanding Indications of Compromise

License: FireSIGHT

As a part of network discovery, the FireSIGHT System's Data Correlator can correlate various types of data (intrusion events, Security Intelligence, connection events, and malware events) associated with hosts to determine whether a host on your monitored network is likely to be compromised by malicious means. These correlations are known as indications of compromise (IOC). You activate this feature by enabling it and any of many Cisco-predefined *IOC rules* in the discovery policy editor. When the feature is enabled, you can also edit rule states for individual hosts from that host's host profile. Each IOC rule corresponds to one specific *IOC tag*, which is associated with a host.

In addition to the Data Correlator, Cisco's endpoint-based Collective Security Intelligence Cloud data can also generate IOC tags from IOC rules. Because this data examines activity on a host itself — such as actions taken by or on individual programs — it can provide insights into possible threats that network-only data cannot. FireAMP IOC data from endpoints is transmitted via the Cisco cloud connection.

Hosts with active IOC tags appear in the IP Address columns of event views with a compromised host icon (🖥) instead of the normal host icon (🖥). Event views for events that can trigger IOC tags indicate whether an event triggered an IOC.

Understanding Indications of Compromise Types

License: FireSIGHT

There are many indications of compromise (IOC) rule and tag types. All are Cisco-predefined, and one IOC rule corresponds to one IOC tag. Because IOC rules trigger based on data provided by other features of the FireSIGHT System (and, for some events, the Cisco cloud), those features must be available and active for IOC rules to set IOC tags. When Cisco develops new endpoint-based malware event IOC types, the system automatically downloads them via the cloud and begins to use them. The lists below detail IOC rule types, the features with which they are associated, and any additional licensing requirements (beyond the FireSIGHT license required for network discovery):

- Endpoint-Based Malware Event IOC Types, page 45-20
- Intrusion Event IOC Types, page 45-21
- Security Intelligence Event IOC Types, page 45-21

Endpoint-Based Malware Event IOC Types

License: FireSIGHT

The following list contains examples of IOC types that are associated with endpoint-based malware events, which require a subscription to the Cisco cloud. In addition to the IOC types listed below, Cisco periodically develops new types, which your system downloads and implements automatically via your connection to the cloud.

For more information on configuring endpoint-based malware protection, see Working with Cloud Connections for FireAMP, page 37-24 and Network-Based AMP vs Endpoint-Based FireAMP, page 37-7.

- Adobe Reader Compromise — Adobe Reader launched shell
- Adobe Reader Compromise — PDF Compromise Detected by FireAMP
- CnC Connected — Suspected Botnet Detected by FireAMP
- Dropper Infection — Dropper Infection Detected by FireAMP
- Excel Compromise — Excel Compromise Detected by FireAMP
- Excel Compromise — Excel launched shell
- Generic IOC Detected by FireAMP
- Java Compromise — Java Compromise Detected by FireAMP

- Java Compromise — Java launched shell
- Malware Detected — Threat Detected by FireAMP - Not Executed
- Malware Detected — Threat Detected in File Transfer
- Malware Executed — Threat Detected by FireAMP - Executed
- Microsoft Calculator Compromise — Microsoft Calculator Compromise Detected by FireAMP
- Microsoft Notepad Compromise — Microsoft Calculator Compromise Detected by FireAMP
- PowerPoint Compromise — PowerPoint Compromise Detected by FireAMP
- PowerPoint Compromise — PowerPoint launched shell
- QuickTime Compromise — QuickTime Compromise Detected by FireAMP
- QuickTime Compromise — QuickTime launched shell
- Word Compromise — Word Compromise Detected by FireAMP
- Word Compromise — Word launched shell

Intrusion Event IOC Types

License: FireSIGHT+Protection

The following IOC types are associated with intrusion events, which require a Protection license. For more information on viewing intrusion events and configuring intrusion detection and protection, see Controlling Traffic Using Intrusion and File Policies, page 18-1 and Viewing Intrusion Events, page 41-9.

- CnC Connected — Intrusion Event - malware-backdoor
- CnC Connected — Intrusion Event - malware-cnc
- Exploit Kit — Intrusion Event - exploit-kit
- Impact 1 Attack — Impact 1 Intrusion Event - attempted-admin
- Impact 1 Attack — Impact 1 Intrusion Event - attempted-user
- Impact 1 Attack — Impact 1 Intrusion Event - successful-admin
- Impact 1 Attack — Impact 1 Intrusion Event - successful-user
- Impact 1 Attack — Impact 1 Intrusion Event - web-application-attack
- Impact 2 Attack — Impact 2 Intrusion Event - attempted-admin
- Impact 2 Attack — Impact 2 Intrusion Event - attempted-user
- Impact 2 Attack — Impact 2 Intrusion Event - successful-admin
- Impact 2 Attack — Impact 2 Intrusion Event - successful-user
- Impact 2 Attack — Impact 2 Intrusion Event - web-application-attack

Security Intelligence Event IOC Types

License: FireSIGHT+Protection

Supported Devices: Any except Series 2

Supported Defense Centers: Any except DC500

The CnC Connected — Security Intelligence Event - CnC type is associated with Security Intelligence events, a type of connection event. The Security Intelligence feature requires a Protection license. For more information on configuring Security Intelligence and viewing Security Intelligence events, see Blacklisting Using Security Intelligence IP Address Reputation, page 13-1 and Viewing Connection and Security Intelligence Data, page 39-14.

Viewing and Editing Indications of Compromise Data

License: FireSIGHT

Outside the network discovery policy itself, you can view and edit indications of compromise (IOC) data in several other parts of the FireSIGHT System web interface:

- In the dashboard, the Threats tab of the Summary Dashboard displays, by default, IOC tags by host and new IOC rules triggered over time. The Custom Analysis widget offers presets based on IOC data. For information, see Using Dashboards, page 55-1 and Configuring the Custom Analysis Widget, page 55-15.
- The Indications of Compromise section of the Context Explorer displays graphs of hosts by IOC category and IOC categories by host. For information, see Understanding the Indications of Compromise Section, page 56-4.
- Event views for discovery (IOC), connection, Security Intelligence, intrusion, and malware events display (in the IOC column) whether an event triggered an IOC rule. Endpoint-based malware events that trigger IOC rules have the event type FireAMP IOC and appear with an event subtype that specifies the compromise. You can write compliance rules against all IOC data that appears in the event viewer. For more information, see the following sections:
- Viewing Connection and Security Intelligence Data, page 39-14
- Viewing Intrusion Events, page 41-9
- Working with Malware Events, page 40-16
- Working with Indications of Compromise, page 50-31
- Configuring Correlation Policies and Rules, page 51-1
- The Indications of Compromise tab of the network map lists hosts on your monitored network, grouped by IOC tag. For information, see Working with the Indications of Compromise Network Map, page 48-4.
- In the host profile view for a potentially compromised host, you can view all IOC tags associated with that host, resolve any or all of its IOC tags, and configure IOC rule states. For information, see Working with Indications of Compromise in the Host Profile, page 49-8.

Creating a Network Discovery Policy

License: FireSIGHT

The network discovery policy on the Defense Center controls how the system collects data on your organization’s network assets and which network segments and ports are monitored.

Discovery rules within the policy specify what networks and ports the FireSIGHT System monitors to generate discovery data based on network data in traffic, and what zones the policy is applied to. Within a rule, you can configure whether hosts, applications, and users are discovered. You can create rules to exclude networks and zones from discovery. When you create a rule for discovery from a NetFlow device, you can choose to just log connections.

The network discovery policy has a a single default rule in place, configured to discover applications in any IPv4 traffic on the 0.0.0.0/0 network. Note that you must have applied an access control policy to the targeted device before you can apply a network discovery policy. The rule does not exclude any networks, zones, or ports, host and user discovery is not configured, and a NetFlow device is not configured. Note that the policy is applied to any managed devices by default when they are registered to the Defense Center. To begin collecting host or data, you must add or modify discovery rules and reapply the policy to a device.

Remember that the access control policy defines the traffic that you permit, and therefore the traffic you can monitor with network discovery. Note that this means if you block certain traffic using access control, the system cannot examine that traffic for host, user, or application activity. For example, if you block access to social networking applications in the access control policy, the system will not provide you with any discovery data on those applications.

If you want to adjust the scope of network discovery, you can create additional discovery rules and modify or remove the default rule. You can configure discovery of data from NetFlow devices and can restrict the protocols for traffic where user data is discovered on your network.

If you want to use the FireSIGHT System to perform intrusion detection and prevention but do not need to take advantage of discovery data, you can optimize performance by disabling new discovery. First, make sure that your applied access control policies do not contain rules with user, application, or URL conditions. Then, remove all rules from your network discovery policy and apply it to your managed devices. For more information on configuring access control rules, see Tuning Traffic Flow Using Access Control Rules, page 14-1.

If you enable user discovery in your discovery rules, you can detect users through user login activity in traffic over a set of application protocols. You can disable discovery in particular protocols across all rules if needed. Disabling some protocols can help avoid reaching the user limit associated with your FireSIGHT license, reserving available user count for users from the other protocols.

Advanced network discovery settings allow you to manage what data is logged, how discovery data is stored, what indications of compromise (IOC) rules are active, what vulnerability mappings are used for impact assessment, and what happens when sources offer conflicting discovery data. You can also add NetFlow devices and sources for host input.

For more information, see:

- Working with Discovery Rules, page 45-23
- Restricting User Logging, page 45-29
- Configuring Advanced Network Discovery Options, page 45-30
- Applying the Network Discovery Policy, page 45-37

Working with Discovery Rules

License: FireSIGHT

Discovery rules allow you to tailor the information discovered for your network map to include only the specific data you want. Rules in your network discovery policy are evaluated sequentially. Note that while you can create rules with overlapping monitoring criteria, doing so may affect your system performance.

When you exclude a host or a network from monitoring, the host or network does not appear in the network map and no events are reported for it. Cisco recommends that you exclude load balancers (or specific ports on load balancers) and NAT devices from monitoring. These devices may create excessive and misleading events, filling the database and overloading the Defense Center. For example, a monitored NAT device might exhibit multiple updates of its operating system in a short period of time. If you know the IP addresses of your load balancers and NAT devices, you can exclude them from monitoring.

Tip The system can identify many load balancers and NAT devices by examining your network traffic. To determine which hosts on your network are load balancers and NAT devices, apply your network discovery policy, wait for the system to populate the network map, then perform a search of hosts constraining on host type.

In addition, if you need to create a custom server fingerprint, you should temporarily exclude from monitoring the IP address that you are using to communicate with the host you are fingerprinting. Otherwise, the network map and discovery event views will be cluttered with inaccurate information about the host represented by that IP address. After you create the fingerprint, you can configure your policy to monitor that IP address again. For more information, see Fingerprinting Servers, page 46-11.

Cisco also recommends that you **not** monitor the same network segment with NetFlow-enabled devices and FireSIGHT System managed devices. Although ideally you should configure your network discovery policy with non-overlapping rules, the system does drop duplicate connection logs generated by managed devices. Note that you **cannot** drop duplicate connection logs for connections detected by both a managed device and a NetFlow-enabled device.

For more information, see the following sections:

- Understanding Device Selection, page 45-24
- Understanding Actions and Discovered Assets, page 45-24
- Understanding Monitored Networks, page 45-25
- Understanding Zones in Network Discovery Policies, page 45-25
- Understanding Port Exclusions, page 45-26
- Adding a Discovery Rule, page 45-26
- Creating Network Objects, page 45-28
- Creating Port Objects, page 45-29

Understanding Device Selection

License: FireSIGHT

If you select a NetFlow device in a discovery rule, the rule is limited to discovery of NetFlow data for the specified networks. Select the NetFlow device before you configure other aspects of rule behavior, as the available rule actions change when you select a NetFlow device. In addition, you cannot configure port exclusions for NetFlow traffic.

Before you can select a NetFlow device in a network discovery rule, you must configure a connection to the NetFlow device in the network discovery advanced settings. For more information, see Adding NetFlow-Enabled Devices, page 45-34.

Understanding Actions and Discovered Assets

License: FireSIGHT

When you configure a discovery rule, you must select an action for the rule. The action determines what assets are discovered or excluded when the system processes the rule. However, note that the affect of a rule action depends on whether you are using the rule to discover data from a managed device or from a NetFlow-enabled device.

Note that if you create a network discovery policy without any rules that discover hosts or users, applying the policy disables new discovery for the appliance. To optimize performance when using managed devices only for intrusion prevention, remove all discovery rules from your policy and apply it to the active devices.

The following table describes what assets are discovered by rules with the specified action settings in those two scenarios.

Table 45-4 Discovery Rule Actions

Action	Managed Device	NetFlow
Exclude	Excludes the specified network from monitoring. If the source or destination host for a connection is excluded from discovery, the connection is recorded but discovery events are not created for excluded hosts.	
Discover: Hosts	Adds hosts to the network map based on discovery events. (Optional, unless user discovery is enabled, then required.)	Adds hosts to the network map based on NetFlow records. (Required)
Discover: Applications	Adds applications to the network map based on application detectors. Note that you cannot discover hosts or users in a rule without also discovering applications. (Required)	Adds application protocols to the network map based on NetFlow records and the port-application protocol correlation in `/etc/sf/services`. (Optional)
Discover: Users	Adds users to the users table and logs user activity based on activity detected in traffic matching the user protocols configured in the network discovery policy. (Optional)	n/a
Log NetFlow Connections	n/a	Logs NetFlow connections only. Does not discover hosts or applications.

Understanding Monitored Networks

License: FireSIGHT

A discovery rule causes discovery of monitored assets only in traffic to and from hosts in the specified networks. For a discovery rule, discovery occurs for connections that have at least one IP address within the networks specified, with events generated only for IP addresses within the networks to monitor. The default discovery rule discovers applications only on the `0.0.0.0/0` and `::/0` networks.

For rules with a specified NetFlow device and the **Log Network Connections** option enabled, connections to and from IP addresses in the specified networks are also logged. Note that network discovery rules provide the only way to log NetFlow network connections.

You can also use network object or object groups to specify the networks to monitor. If you modify a network object used in the network discovery policy, you must reapply the policy for those changes to take effect for discovery.

Understanding Zones in Network Discovery Policies

License: FireSIGHT

For performance reasons, you should configure each discovery rule so that the zones in the rule include the sensing interfaces on your managed devices that are physically connected to the networks-to-monitor in the rule.

Unfortunately, you may not always be kept informed of network configuration changes. A network administrator may modify a network configuration through routing or host changes without informing you, which may make it challenging to stay on top of proper network discovery policy configurations. If you do not know how the sensing interfaces on your managed devices are physically connected to your network, leave the zone configuration as the default, which is to apply the discovery rule to all zones in your deployment. (If no zones are excluded, the discovery policy is applied to all zones.)

Understanding Port Exclusions

License: FireSIGHT

Just as you can exclude hosts from monitoring (see Understanding Actions and Discovered Assets, page 45-24), you can exclude specific ports from monitoring.

For example, load balancers can report multiple applications on the same port in a short period of time. You can configure your network discovery policy so that it excludes that port from monitoring, such as excluding port 80 on a load balancer that handles a web farm.

As another scenario, your organization may use a custom client that uses a specific range of ports. If the traffic from this client generates excessive and misleading events, you can exclude those ports from monitoring. Similarly, you may decide that you do not want to monitor DNS traffic. In that case, you could configure your policy so that it does not monitor port 53.

When adding ports to exclude, you can decide whether to use a reusable port object from the Available Ports list, add ports directly to the source or destination exclusion lists, or create a new reusable port and then move it into the exclusion lists.

Note that you cannot configure NetFlow-enabled devices to exclude ports from monitoring.

Adding a Discovery Rule

License: FireSIGHT

You can configure discovery rules to tailor the discovery of host and application data to your needs. Note that when you modify an object referenced in a rule, you must reapply the network discovery policy for those changes to take effect.

To add a discovery rule:

Access: Admin/Discovery Admin

Step 1 Check your access control policies to ensure that you are logging connections as needed for the traffic where you want to discover network data.

For more information, see Logging Connections Based on Access Control Handling, page 38-15. To discover the most data, log at the end of the connection for traffic you want to discover.

Step 2 Select **Policies > Network Discovery**.

The Network Discovery Policy page appears.

Step 3 Click **Add Rule**.

The Add Rule pop-up window appears.

Step 4 You have two options:

- If you plan to use the rule to monitor NetFlow traffic, within the Add Rule pop-up window, click **NetFlow Device**.

  The NetFlow Device page appears.

  Note that the NetFlow page is available only if you have added a NetFlow device to the discovery policy. For more information, see Adding NetFlow-Enabled Devices, page 45-34.

- If you plan to use the rule to monitor managed devices, skip to step 6.

For more information, see Differences Between NetFlow and FireSIGHT Data, page 45-17 and Understanding Device Selection, page 45-24

Step 5 Select the IP address for the NetFlow device you want to use from the drop-down list.

Step 6 Set the action for the rule:

- To exclude all traffic that matches the rule from network discovery, select **Exclude**. Note that the Port Exclusions tab is disabled when you select this rule action.
- To discover the selected types of data in traffic that matches the rule, select **Discovery** and select or clear the appropriate data type check boxes.

  If monitoring managed device traffic, application logging is required. If monitoring users, host logging is required. If monitoring NetFlow traffic, note that you cannot log users and that logging applications is optional.

- If monitoring NetFlow traffic, to use the rule to log connections in NetFlow traffic, select **Log NetFlow Connections**. Note that this option only appears after you have selected a NetFlow device in the rule.

Note The system detects connections in NetFlow traffic based on network discovery policy settings. Connection logging in managed device traffic is configured in the access control policy. For more information, see Logging Connections in Network Traffic, page 38-1.

For more information on rule actions and discovery of assets, see Understanding Actions and Discovered Assets, page 45-24

Step 7 Every discovery rule must include at least one network. Optionally, to restrict the rule action to specific networks, click the **Networks** tab, select a network from the **Available Networks** list, and click **Add**, or type the network below the Networks list and click **Add**.

For information on network monitoring, see Understanding Monitored Networks, page 45-25. For information on adding network objects to the Available Networks list, see Creating Network Objects, page 45-28. Note that If you modify a network object used in the network discovery policy, you must reapply the policy for those changes to take effect for discovery.

Step 8 Optionally, to restrict the rule actions to traffic in specific zones, click **Zones**, select a zone or zones from the **Available Zones** list, and click **Add**.

For information on selecting zones for monitoring, see Understanding Zones in Network Discovery Policies, page 45-25.

Step 9 To exclude ports from monitoring, click **Port Exclusions**.

The Port Exclusions page appears.

Step 10 To exclude specific source ports from monitoring, you have two options:

- Select a port or ports from the **Available Ports** list and click **Add to Source**.
- To exclude traffic from a specific source port without adding a port object, under the **Selected Source Ports** list, select the appropriate protocol from the **Protocol** drop-down list, type a port number from 1 to 65535 into the **Port** field, and click **Add**.

For information on excluding ports from monitoring, see Understanding Port Exclusions, page 45-26. For information on adding port objects to the Available Ports list, see Creating Port Objects, page 45-29. Note that if you modify a port object used in the network discovery policy, you must reapply the policy for those changes to take effect for discovery.

Step 11 To exclude specific destination ports from monitoring, you have two options:

- Select a port or ports from the **Available Ports** list and click **Add to Destination**.
- To exclude traffic from a specific destination port without adding a port object, under the **Selected Destination Ports** list, select the appropriate protocol from the **Protocol** drop-down list, type a port number from 1 to 65535 into the **Port** field, and click **Add**.

Step 12 If you are finished editing the rule, click **Save** to return to the discovery policy rule list.

You must apply the network discovery policy for your changes to take effect. For more information, see Applying the Network Discovery Policy, page 45-37.

Creating Network Objects

License: FireSIGHT

The list of available networks that appears in a discovery rule contains reusable network object and groups that can be used anywhere in the FireSIGHT System. You can add new network objects to the list. Note that when you modify an object referenced in a rule, you must reapply the network discovery policy for those changes to take effect.

To create a new network object:

Admin/Discovery Admin

Step 1 Select **Policies > Network Discovery**.

The Network Discovery Policy page appears.

Step 2 Click **Add Rule**.

The Add Rule pop-up window appears.

Step 3 On the Networks page, click the add icon (⊕).

The Network Objects pop-up window appears.

Step 4 Type a **Name** for the network object. You can use any printable standard ASCII characters except a pipe (|) or curly braces ({}).

Step 5 For each IP address, CIDR block, and prefix length you want to add to the network object, type its value and click **Add**.

Step 6 Click **Save** to add the network object to the Available Networks list.

Tip If the network does not immediately appear on the list, click the refresh icon (⟳).

Creating Port Objects

License: FireSIGHT

The list of available ports that appears in a discovery rule contains reusable port objects and groups that can be used anywhere in the FireSIGHT System. You can add new port objects to the list. Note that when you modify an object referenced in a rule, you must reapply the network discovery policy for those changes to take effect.

To create a new port object:

Admin/Discovery Admin

Step 1 Click **Port Exclusions**.

The Port Exclusions page appears.

Step 2 To add a port to the Available Ports list, click the add object icon (⊕).

The Port Objects pop-up window appears.

Step 3 Supply a **Name** for the port object. You can use any printable standard ASCII characters except a pipe (`|`) or curly braces (`{}`).

Step 4 In the **Protocol** field, specify the protocol of the traffic you want to exclude.

Select **TCP**, **UDP**, or **Other** and choose an option from the drop-down list to select a protocol or **All**.

Step 5 In the **Port(s)** field, enter the ports you want to exclude from monitoring.

You can specify a single port, a range of ports using the dash (-), or a comma-separated list of ports and port ranges. Allowed port values are from 1 to 65535.

Step 6 Click **Save** to add the port to the Available Ports list.

Tip If the port does not immediately appear on the list, click the refresh icon (🔄).

Restricting User Logging

License: FireSIGHT

When you apply a network discovery policy with rules that discover users, users are discovered in traffic that uses the AIM, IMAP, LDAP, Oracle, POP3, SMTP, FTP, HTTP, MDNS, and SIP protocols. These users are added to the users table, accessible through the Analysis menu. You can restrict the protocols where user activity is discovered to reduce the total number of detected users so you can focus on users likely to provide the most complete user information.

The total number of detected users the Defense Center can store depends on your FireSIGHT license. After you reach the licensed limit, in most cases the system stops adding new users to the database. To add new users, you must either manually delete old or inactive users from the database, or purge all users from the database. Restricting protocol detection helps minimize user name clutter and preserve FireSIGHT user licenses.

For example, obtaining user names through protocols such as AIM, POP3, and IMAP may introduce user names not relevant to your organization due to network access from contractors, visitors, and other guests.

As another example, AIM, Oracle, and SIP logins may create extraneous user records. This occurs because these login types are not associated with any of the user metadata that the system obtains from an LDAP server, nor are they associated with any of the information contained in the other types of login that your managed devices detect. Therefore, the Defense Center cannot correlate these users with other types of users.

Keep in mind that only managed devices can detect non-LDAP user logins. If you are using only User Agents installed on Microsoft Active Directory servers to detect user activity, restricting non-LDAP logins has no effect. Also, you cannot restrict SMTP logging. This is because users are not added to the database based on SMTP logins; although the system detects SMTP logins, the logins are not recorded unless there is already a user with a matching email address in the database.

You can choose to record failed login attempts for failed user logins detected in LDAP, POP3, FTP, or IMAP traffic. A failed login attempt does not add a new user to the list of users in the database. Note that the User Agent does not report failed login activity. The user activity type for detected failed login activity is Failed User Login.

Note that the system cannot distinguish between failed and successful HTTP logins. To see HTTP user information, you must enable **Capture Failed Login Attempts**.

To restrict the protocols where user logins are detected:

Admin/Discovery Admin

Step 1 Select **Policies > Network Discovery**.

The Network Discovery Policy page appears.

Step 2 Click **User**.

The User page appears.

Step 3 Select check boxes for protocols where you want to detect logins or clear check boxes for protocols where you do not want to detect logins.

Step 4 Optionally, to record failed login attempts detected in LDAP, POP3, FTP, or IMAP traffic, or to capture user information for HTTP logins, enable **Capture Failed Login Attempts**.

Step 5 Click **Save** to save the network policy.

You must apply the network discovery policy for your changes to take effect. For more information, see Applying the Network Discovery Policy, page 45-37.

Configuring Advanced Network Discovery Options

License: FireSIGHT

The Advanced tab of the network discovery policy allows you to configure policy-wide settings for what events are detected, how long discovery data is retained and how often it is updated, what vulnerability mappings are used for impact correlation, and how operating system and server identity conflicts are resolved. In addition, you can add host input sources and NetFlow-enabled devices to allow import of data from other sources.

Note that the database event limits for discovery and user activity events are set in the system policy. For more information, see Configuring Database Event Limits, page 63-15.

To configure advanced settings:

Admin/Discovery Admin

Step 1 Select **Policies > Network Discovery**.

The Network Discovery Policy page appears.

Step 2 Click **Advanced**.

The Advanced page appears.

Step 3 Edit advanced settings as needed:

- Configuring General Settings, page 45-31
- Configuring Identity Conflict Resolution, page 45-32
- Enabling Vulnerability Impact Assessment Mappings, page 45-33
- Setting Indications of Compromise Rules, page 45-33
- Adding NetFlow-Enabled Devices, page 45-34
- Configuring Data Storage, page 45-35
- Configuring Discovery Event Logging, page 45-36
- Adding Identity Sources, page 45-37

Step 4 When you finish configuring settings, click **Save** to save the policy.

Step 5 When the policy is complete and saved, apply the policy to put the updated settings into effect. For more information, see Applying the Network Discovery Policy, page 45-37.

Configuring General Settings

License: FireSIGHT

The general settings control how often the system updates information in the network map and whether server banners are captured during discovery.

Capture Banners

Select this check box if you want the system to store header information from network traffic that advertises server vendors and versions (“banners”). This information can provide additional context to the information gathered. You can access server banners collected for hosts by accessing server details.

Update Interval

The interval at which the system updates information (such as when any of a host’s IP addresses was last seen, when an application was used, or the number of hits for an application). The default setting is 3600 seconds (1 hour).

Note that setting a lower interval for update timeouts provides more accurate information in the host display, but generates more network events.

To update general settings:

Admin/Discovery Admin

Step 1 Click the edit icon (✎) next to **General Settings**.

The General Settings pop-up window appears.

Step 2 Update the settings as needed.

Step 3 Click **Save** to save the general settings and return to the Advanced tab of the network discovery policy.

You must apply the network discovery policy for your changes to take effect. For more information, see Applying the Network Discovery Policy, page 45-37.

Configuring Identity Conflict Resolution

License: FireSIGHT

The system matches fingerprints for operating systems and servers against patterns in traffic to determine what operating system and which applications are running on a particular host. To provide the most reliable operating system and server identity information, the system collates fingerprint information from several sources.

The system uses all passive data to derive operating system identities and assign a confidence value. For more information on current identities and how the system selects the current identity, see Enhancing Your Network Map, page 46-4.

By default, unless there is an identity conflict, identity data added by a scanner or third-party application overrides identity data detected by the FireSIGHT System. You can use the Identity Sources settings to rank scanner and third-party application fingerprint sources by priority. The system retains one identity for each source, but only data from the highest priority third-party application or scanner source is used as the current identity. Note, however, that user input data overrides scanner and third-party application data regardless of priority.

An identity conflict occurs when the system detects an identity that conflicts with an existing identity that came from either the active scanner or third-party application sources listed in the Identity Sources settings or from a FireSIGHT System user. By default, identity conflicts are not automatically resolved and you must resolve them through the host profile or by rescanning the host or re-adding new identity data to override the passive identity. However, you can set your system to always automatically resolve the conflict by keeping the passive identity or to always resolve it by keeping the active identity.

Generate Identity Conflict Event

Enable this option to generate an event when an identity conflict occurs on a host in the network map.

Automatically Resolve Conflicts

You have the following options:

- To force manual conflict resolution of identity conflicts, select **Disabled** from the **Automatically Resolve Conflicts** drop-down list.
- To use the passive fingerprint when an identity conflict occurs, select **Identity** from the **Automatically Resolve Conflicts** drop-down list.
- To use the current identity from the highest priority active source when an identity conflict occurs, select **Keep Active** from the **Automatically Resolve Conflicts** drop-down list.

To update identity conflict resolution settings:

Admin/Discovery Admin

Step 1 Click the edit icon (✎) next to **Identity Conflict Settings**.

The Edit Identity Conflict Settings pop-up window appears.

Step 2 Update the settings as needed.

Step 3 Click **Save** to save the identity conflict settings and return to the **Advanced** tab of the network discovery policy.

You must apply the network discovery policy for your changes to take effect. For more information, see Applying the Network Discovery Policy, page 45-37.

Enabling Vulnerability Impact Assessment Mappings

License: FireSIGHT

You can configure how the FireSIGHT System performs impact correlation with intrusion events. Your options are as follows:

- Select **Use Network Discovery Vulnerability Mappings** if you want to use system-based vulnerability information to perform impact correlation.
- Select **Use Third-Party Vulnerability Mappings** if you want to use third-party vulnerability references to perform impact correlation. For more information, see Mapping Third-Party Vulnerabilities, page 46-33 or the *FireSIGHT System Host Input API Guide*.

You can select either or both of the check boxes. If the system generates an intrusion event and the host involved in the event has servers or an operating system with vulnerabilities in the selected vulnerability mapping sets, the intrusion event is marked with the Vulnerable (level 1: red) impact icon. For any servers which do not have vendor or version information, note that you need to configure vulnerability mapping in the system policy. For more information, see Mapping Vulnerabilities for Servers, page 63-30.

If you clear both check boxes, intrusion events will **never** be marked with the Vulnerable (level 1: red) impact icon. For more information, see Using Impact Levels to Evaluate Events, page 41-37.

To update vulnerability settings:

Admin/Discovery Admin

Step 1 Click the edit icon (✎) next to **Vulnerabilities to use for Impact Assessment**.

The Edit Vulnerability Settings pop-up window appears.

Step 2 Update the settings as needed.

Step 3 Click **Save** to save the vulnerability settings and return to the Advanced tab of the network discovery policy.

You must apply the network discovery policy for your changes to take effect. For more information, see Applying the Network Discovery Policy, page 45-37.

Setting Indications of Compromise Rules

License: FireSIGHT

For your system to detect and tag indications of compromise (IOC), you must first activate at least one IOC rule in your discovery policy. Each IOC rule corresponds to one type of IOC tag, and all IOC rules are predefined by Cisco; you cannot create original rules. You can enable any or all rules, depending on the needs of your network and organization. For example, if hosts using software such as Microsoft Excel never appear on your monitored network, you may decide not to enable the IOC tags that pertain to Excel-based threats. For more information on the IOC feature, see Understanding Indications of Compromise, page 45-19.

You must also enable the FireSIGHT System features associated with the IOC rules you enable, such as intrusion and malware protection; if a rule's associated feature is not enabled, no relevant data is collected and the rule cannot trigger. For more information on the types of IOC rules and their associated features, see Understanding Indications of Compromise Types, page 45-20.

To set indications of compromise rules in the discovery policy:

Admin/Discovery Admin

Step 1 Click the edit icon (✎) next to **Indications of Compromise Settings**.

The Edit Indications of Compromise Settings pop-up window appears.

Step 2 To toggle the entire IOC feature off or on, click the slider next to **Enable IOC**.

Step 3 To enable or disable individual IOC rules, click the slider in the rule's **Enabled** column.

Step 4 Click **Save** to save your IOC rule settings and return to the Advanced tab of the discovery policy.

Your changes are saved.

You must apply the network discovery policy for your changes to take effect. For more information, see Applying the Network Discovery Policy, page 45-37.

Adding NetFlow-Enabled Devices

License: FireSIGHT

If you have enabled the NetFlow feature on your NetFlow-enabled devices, you can use the connection data exported by these devices to supplement the connection data collected by Cisco devices.

Before you can use them in discovery rules, you must configure the NetFlow-enabled devices you plan to use (see Preparing to Analyze NetFlow Data, page 45-18), then add them to the network discovery policy.

For more information on using NetFlow data with the FireSIGHT System, including information on additional prerequisites, see Understanding NetFlow, page 45-16.

To add NetFlow-enabled devices for connection data collection:

Admin/Discovery Admin

Step 1 Select **Policies > Network Discovery**.

The Network Discovery Policy page appears.

Step 2 Click **Advanced**.

The Advanced page appears.

Step 3 Click the add icon (⊕) next to NetFlow Devices.

The Add NetFlow Device pop-up window appears.

Step 4 In the **IP Address** field, enter the IP address of the NetFlow-enabled device you want to use to collect connection data.

Step 5 To add additional NetFlow-enabled devices, repeat steps 3 and 4.

Tip To remove a NetFlow-enabled device, click the delete icon (🗑) next to the device you want to remove. Keep in mind that if you use a NetFlow-enabled device in a discovery rule, you must delete the rule before you can delete the device from the Advanced page. For more information, see Working with Discovery Rules, page 45-23.

Step 6 Click **Save**.

The device appears on the list of NetFlow-enabled devices.

You must apply the network discovery policy for your changes to take effect. For more information, see Applying the Network Discovery Policy, page 45-37.

Configuring Data Storage

License: FireSIGHT

Data storage settings control the kinds of data stored in the database, and therefore determine the data that the FireSIGHT System can use. These settings also control how long data is retained in the network map.

The following options comprise the network discovery data storage settings.

When Host Limit Reached

You can control how hosts are handled when the Defense Center reaches its host limit (as determined by the FireSIGHT license) and the network map is full. This option is especially valuable if you want to prevent spoofed hosts from taking the place of valid hosts in the network map. To drop old hosts, select **Drop hosts** from the **When Host Limit Reached** drop-down list.To drop new hosts, select **Don't insert new hosts** from the **When Host Limit Reached** drop-down list. For more information, see Understanding FireSIGHT Host and User License Limits, page 65-8.

Host Timeout

The amount of time that passes, in minutes, before the system drops a host from the network map due to inactivity. The default setting is 10080 minutes (7 days). Individual host IP and MAC addresses can time out individually, but a host does not disappear from the network map unless all of its associated addresses have timed out.

To avoid premature timeout of hosts, make sure that the host timeout value is longer than the update interval in the network discovery policy. For more information on the update interval, see Configuring General Settings, page 45-31.

Server Timeout

The amount of time that passes, in minutes, before the system drops a server from the network map due to inactivity. The default setting is 10080 minutes (7 days).

To avoid premature timeout of servers, make sure that the service timeout value is longer than the update interval in the network discovery policy. For more information, see Configuring General Settings, page 45-31.

Client Application Timeout

The amount of time that passes, in minutes, before the system drops a client from the network map due to inactivity. The default setting is 10080 minutes (7 days).

You should make sure that the client timeout value is longer than the update interval in the network discovery policy. For more information, see Configuring General Settings, page 45-31.

To update data storage settings:

Admin/Discovery Admin

Step 1 Click the edit icon (✎) next to **Data Storage Settings**.

The Data Storage Settings pop-up window appears.

Step 2 Update the settings as needed.

Step 3 Click **Save** to save the data storage settings and return to the Advanced tab of the network discovery policy.

You must apply the network discovery policy for your changes to take effect. For more information, see Applying the Network Discovery Policy, page 45-37.

Configuring Discovery Event Logging

License: FireSIGHT

The Event Logging Settings control whether discovery and host input events are logged. If you do not log an event, you cannot retrieve it in event views or use it to trigger correlation rules.

To set event logging settings:

Admin/Discovery Admin

Step 1 Click the edit icon (✎) next to **Event Logging Settings**.

The Event Logging Settings pop-up window appears.

Step 2 Select or clear the check boxes next to the discovery and host input event types you want to log in the database. See Understanding Discovery Event Types, page 50-9 and Understanding Host Input Event Types, page 50-13 for information about each event type.

Step 3 Click **Save** to save the event logging settings and return to the Advanced tab of the network discovery policy.

You must apply the network discovery policy for your changes to take effect. For more information, see Applying the Network Discovery Policy, page 45-37.

Adding Identity Sources

License: FireSIGHT

You can add new active sources through this page, or change the priority or timeout settings for existing sources. Note that adding a scanner to this page does not add the full integration capabilities that exist for the Nmap scanners, but does allow integration of imported third-party application or scan results. If you import data from a third-party application or scanner, remember to make sure that you map vulnerabilities from the source to the vulnerabilities in the network map. For more information, see Mapping Third-Party Vulnerabilities, page 46-33.

To add identity sources:

Admin/Discovery Admin

Step 1 Click the edit icon (✎) next to **OS and Server Identity Sources**.

The Edit OS and Server Identity Sources pop-up window appears.

Step 2 To add a new source, click **Add Source**.

The Add Identity Source pop-up window appears.

Step 3 Type a **Name** for the source.

Step 4 Select the input source type from the **Type** drop-down list:

- Select **Scanner** if you plan to import scan results using the AddScanResult function.
- Select **Application** if you do not plan to import scan results.

Step 5 To indicate the duration of time that should elapse between the addition of an identity to the network map by this source and the deletion of that identity, select **Hours**, **Days**, or **Weeks** from the **Timeout** drop-down list and type the appropriate duration.

Tip To delete a source that you added, click the delete icon (🗑) next to the source.

Step 6 Optionally, to promote a source and cause the operating system and application identities to be used in favor of sources below it in the list, select the source and click the up arrow.

Step 7 Optionally, to demote a source and cause the operating system and application identities to be used only if there are no identities provided by sources above it in the list, select the source and click the down arrow.

Step 8 Click **Save** to save the identity source settings and return to the Advanced tab of the network discovery policy.

You must apply the network discovery policy for your changes to take effect. For more information, see Applying the Network Discovery Policy, page 45-37.

Applying the Network Discovery Policy

License: FireSIGHT

By default, the network discovery policy is applied to any targeted zones on managed devices when they are registered with the Defense Center. Applying the network discovery policy allows the system to begin monitoring your network according to your specifications. If you change the network discovery policy, you must reapply it before your changes take effect.

When you reapply the network discovery policy:

- the system deletes and then rediscovers MAC address, TTL, and hops information from the network map for the hosts in your monitored networks
- the affected managed devices discard any discovery data that has not yet been sent to the Defense Center

When you apply a network discovery policy, make sure that you have already applied an access control policy to all devices managed by the Defense Center. If an access control policy has not been applied to each device, the network discovery policy apply fails. Note that you cannot apply a network discovery policy on a Defense Center where no FireSIGHT license is installed.

If you modify a network or port object used in the network discovery policy, you must reapply the policy for those changes to take effect for discovery.

Note that you cannot apply a network discovery policy to stacked devices running different versions of the FireSIGHT System (for example, if an upgrade on one of the devices fails).

To apply the network discovery policy:

Admin/Security Approver

Step 1 Select **Policies > Network Discovery**.

The Network Discovery Policy page appears.

Step 2 Click **Apply**.

A message appears, confirming that you want to apply the policy to all zones targeted by access control policies on the Defense Center.

Step 3 Click **Yes** to apply the policy.