CA Network Flow Analysis

Administrator Guide

Release 9.3.0

CA Technologies Product References

This document references the following CA Technologies products:

- CA Infrastructure Management
- CA Infrastructure Management Data Aggregator
- CA Network Flow Analysis
- CA Performance Center
- CA ReporterAnalyzer
- CA Single Sign-On

Related Documentation

CA provides a full set of technical documentation in the CA Network Flow Analysis Documentation Bookshelf. Access the bookshelf by clicking the Help link in the CA Network Flow Analysis user interface. You can open the guides in PDF and HTML format from the Documentation Bookshelf.

The documentation may have been updated since its release. To get the latest CA Network Flow Analysis documentation updates and localized documentation, download the Bookshelf from CA Support.

The documentation set for CA Network Flow Analysis 9.3.0 includes the following guides:

- *Online help*: Assistance for Administrators and operators, available through the Help link in the user interface.
- *Administrator Guide*: How to set up and maintain CA Network Flow Analysis.
- *Operator Guide*: How to use the NFA console to create, view, and manage reports.
- *Installation Guide*: How to install the software and perform one-time configuration tasks.
- *Upgrade Guide*: How to upgrade the software and perform initial configuration tasks.
- *Release Notes*: Summary of CA Network Flow Analysis enhancements, fixes, and open issues.
- *CA Anomaly Detector Guide:* How to install, upgrade, configure, and use CA Anomaly Detector.
- *CA Anomaly Detector Release Notes:* Overview of the product, system requirements/recommendations, and features.

The product PDFs are in the following directory:
<install_path>\Reporter\NetQoS.ReporterAnalyzer.WebSite\help\en_US\NFA_Bookshelf\Bookshelf_Files\PDF

To view the documentation PDF files, make sure that Adobe Reader is installed.

Contact CA Technologies

Contact CA Support

For your convenience, CA Technologies provides one site where you can access the information that you need for your Home Office, Small Business, and Enterprise CA Technologies products. At http://ca.com/support, you can access the following resources:

- Online and telephone contact information for technical assistance and customer services
- Information about user communities and forums
- Product and documentation downloads
- CA Support policies and guidelines
- Other helpful resources appropriate for your product

Providing Feedback About Product Documentation

If you have comments or questions about CA Technologies product documentation, you can send a message to techpubs@ca.com.

To provide feedback about CA Technologies product documentation, complete our short customer survey which is available on the CA Support website at http://ca.com/docs.

Contents

Chapter 8: Making Additional Customizations 103

Chapter 9: Maintenance and Data Collection 137

Chapter 1: Introduction

This section contains the following topics:

CA Network Flow Analysis

CA Network Flow Analysis provides network traffic analysis with real-time visibility into the traffic throughout your enterprise. You can access as much as one year of flow data for your entire network.

CA Network Flow Analysis is designed to be integrated as a data source for either CA Performance Center or CA NetQoS Performance Center--whichever program your enterprise uses. The Performance Center Console displays report data from CA Network Flow Analysis and any other programs that are integrated as data sources.

This guide describes administration tasks that you perform in the NFA console. Other administration tasks are performed in the Performance Center Console--the management of users, roles, permissions, SNMP profiles, and some types of groups.

Install Performance Center and register CA Network Flow Analysis as a data source so you can use all the features. You can access Performance Center from the NFA console as soon as you register the product.

Note: This guide uses the term *Performance Center* to refer to CA Performance Center and CA NetQoS Performance Center collectively. Program-specific page names or functions may be identified by the full program name or acronym, which is *CA PC* for CA Performance Center and *NPC* for CA NetQoS Performance Center.

Third Party Acknowledgment and License Agreements

Third-party software was used in the creation of CA Network Flow Analysis. All third-party software has been used in accordance with the terms and conditions for use, reproduction, and distribution as defined by the applicable license agreements.

Information about third-party license agreements is provided in the following document, which is installed automatically with the software:

<install_path>\ThirdPartyContent\ThirdPartyLicenseInfo.pdf

Chapter 2: Initial Configuration

After the CA Network Flow Analysis software is installed, some configuration is needed. The topics that follow provide information to help you configure CA Network Flow Analysis initially.

This section contains the following topics:

Single Sign-On

The Single Sign-On tool authenticates you at start-up so you can log in once and can use several related products without logging in again. If you log in to CA Performance Center, for example, you can jump to information in CA Network Flow Analysis without logging in again.

Single Sign-On is installed automatically when you install CA Network Flow Analysis.

Notes:

- The Single Sign-On tool uses Lightweight Directory Access Protocol (LDAP) if LDAP is set up. The Single Sign-On tool does not use LDAP by default.
- For more information about configuring the Single Sign-On tool for LDAP authentication and other options, see the *Single Sign-On Guide*, which is available in the CA Network Flow Analysis Bookshelf and from CA Support Online.

Configure the Product to Work with Performance Center

To complete some essential administration tasks, you must register CA Network Flow Analysis with Performance Center (either CA Performance Center or CA NetQoS Performance Center). Once you register the product and complete the administration tasks, operators can see meaningful reports in the NFA console and in the Performance Center Console.

To configure CA Network Flow Analysis properly, you must perform the following tasks:

- Register the product as a data source for Performance Center.
- Set up the following elements in the Performance Center Console:
  - SNMP profiles
  - (Optional) IP domains
  - User accounts, permissions, and roles
  - (Recommended) Groups
- Configure flow collection
  - Add Harvesters.
  - (Three-tier distributed architecture only) Add DSAs.
  - Enable the routers and interfaces to export flows.
- (Optional) Configure traps.
- (Optional) Verify that CA Network Flow Analysis is receiving data.

These tasks are described in the following topics.

This section contains the following topics:

Register CA Network Flow Analysis

Register the product in the Performance Center Console--on the Manage Data Sources page (CA PC) or the Data Source List page (NPC).

Note: For information about the number of data sources that you can use, see the *Release Notes* for your Performance Center version.

Follow these steps:

1. Verify that no one else is running a session of CA Network Flow Analysis. If multiple users write to the database simultaneously, problems can result.
2. Log in to the Performance Center Console as a user who has the Administrator role.
3. Click Admin, Data Sources.

   The current list of registered data sources are shown on the Manage Data Sources page (CA PC) or the Data Source List page (NPC).
4. Click Add (CA PC) or New (NPC).

   The Add Data Source dialog opens.
5. Select the type of data source you want to add from the Source Type list.

   Note: All CA products that can be registered as data sources are shown in the Source Type list. The list is not filtered to hide products that are installed already.
6. Enter the Host Name of the data source.

   The hostname is the IP address or DNS hostname of the server that hosts the database for the data source. For a distributed deployment of the product, enter the hostname of the NFA console or stand-alone server.
7. Enter the port to use for contacting CA Network Flow Analysis.

   For more information about the port setting, see the *CA Single Sign-On User Guide*.
8. Select the protocol to use for contacting the data source. Select **https** if your network uses SSL for communications. Verify that you have configured the system correctly before you select the **https** option.

   Note: For more information about using SSL for communication between the product and Performance Center, see the *CA Single Sign-On User Guide.*
9. (Optional) Enter a Display Name for the CA Network Flow Analysis data source.

   If you do not enter a name, a default name is created by combining the data source type and hostname. For example, you can use a default name like NetworkFlowAnalysis@xxx.x.x.xx or you can use a name like NetworkFlowAnalysis_NewYork.

10. Confirm whether the web console address is the same as the Host Name. If it is not, take the following steps:
    - Clear the 'Same as Data Source' check box.
    - Provide the web console hostname, port, and protocol.
11. Click Save.

    The updated list shows the data sources that are registered.

Test Data Source Connections

In most cases, the status indicates that data source registration has completed successfully. If the status indicates an error, use the test feature on the Manage Data Sources page (CA PC) or Data Source List page (NPC).

The Test button initiates a test to confirm that a new data source is registered and connected correctly. The test checks for version compatibility and verifies that the data source is not registered with a different instance of the software.

If the test fails, verify that the DNS hostname or IP address of the data source server is correct.

Verify SNMP Profiles

CA Network Flow Analysis sends secure SNMP information to Performance Center at registration. Subsequently, this information is transformed into SNMP profiles. *SNMP profiles* are definitions that contain the information necessary to enable secure queries of device MIBs using SNMP. Profile information is updated at each synchronization (every five minutes).

During initial product setup, verify that the available SNMP profiles are adequate to monitor your environment. The available SNMP profiles are listed on the Manage SNMP Profiles page (CA PC) or SNMP Profile List page (NPC).

View the SNMP Profiles List

You can view a list of SNMP profiles that have already been defined. The list includes high-level information about the contents of each profile.

Tenant-Specific Profiles in CA Performance Center:

- If no tenant definitions have been created, the definitions in the SNMP Profile List are shared among all registered data sources. The global administrator sees a list of SNMP profiles that are not explicitly associated with a tenant.
- Tenant administrators only see the items that are associated with their tenant.

Follow these steps:

1. Log in as a user with the Administrator role.
2. Display the list of SNMP profiles:
   - CA PC: Select Admin, System Settings: SNMP Profiles.
   - NPC: Select Admin, NetQoS Settings: SNMP Profiles.

   The page displays the current list of SNMP profiles.

   The following information is listed for each profile:

   Order

   Determines the order in which the secure information contained in an SNMP profile is used to try to query a selected device. If the query fails, the next profile is used, in priority order.

   Profile Name

   Defines a name for the SNMP profile. Profile names must be unique, cannot be duplicated across SNMP versions, and are not case-sensitive.

   Port (CA PC Only)

   Identifies the port that is used to make SNMP connections to devices.

   Default: UDP 161.

   SNMP Version

   Specifies the version of SNMP that the profile uses. Because SNMPv1 and SNMPv2C are similar from a security standpoint, they share a single option. SNMPv3 is a separate option.

Authentication Protocol

(SNMPv3 Only) Specifies the authentication protocol to use when contacting devices associated with this profile. The following algorithms for authenticating SNMPv3 packets are supported:

- None (do not attempt authentication)
- MD5 (Message Digest 5)
- SHA (Secure Hash Algorithm)

Privacy Protocol

Identifies the encryption protocol that is used to contact associated devices, if any. Always 'None' if no authorization protocol is in use.

Use by Default

Indicates whether the information in this profile is used when not explicitly assigned to a device. If disabled, this profile is excluded from discovery in data sources that support the exclusion of profiles.

To perform any action on this page, select a profile, then click a button.

Add SNMP Profiles

Administrators can create SNMP profiles in the Performance Center Console. You can create SNMPv1/v2C or SNMPv3 profiles.

Follow these steps:

1. Log in to the Performance Center Console as a user with the Administrator role.
2. Display the list of SNMP profiles:
   - CA PC: Select Admin, System Settings: SNMP Profiles.
   - NPC: Select Admin, NetQoS Settings: SNMP Profiles.

   The page displays the current list of SNMP profiles.
3. Click New.

   The Add SNMP Profile dialog opens.
4. Complete the fields and change default settings as needed. Some fields apply only to SNMPv3 or SNMPv.1/v2C

   Profile Name

   Defines a name for the SNMP profile. Profile names must be unique, cannot be duplicated across SNMP versions, and are not case-sensitive.

   SNMP Version

   Specifies the version of SNMP that the profile uses. Because SNMPv1 and SNMPv2C are similar from a security standpoint, they share a single option. SNMPv3 is a separate option.

Port

(Optional for SNMPv1/v2C) Identifies the port that is used to make SNMP connections to devices associated with this profile.

Default: UDP 161.

User Name

(SNMPv3 Only) Identifies the user for the profile, whose secret keys were used potentially to authenticate and encrypt the SNMPv3 packets. The User Name is a character string.

Context Name

(SNMPv3 Only) Specifies a collection of management information that is accessible by an SNMP entity. The Context Name is necessary for providing end-to-end identification and for retrieving data from an SNMPv3 agent. The Context Name is an octet string.

Community Name

(SNMPv1/v2C Only) Defines a secure string that lets the data source query the MIB of the associated device. The community that you supply must provide read-only access to the device MIB.

Note: In the default SNMP profile, the community is 'public'.

Verify Community Name

Confirms the secure community string (name).

Authentication Protocol

(SNMPv3 Only) Specifies the authentication protocol to use when contacting devices associated with this profile. The following algorithms for authenticating SNMPv3 packets are supported:

- None (do not attempt authentication)
- MD5 (Message Digest 5)
- SHA (Secure Hash Algorithm)

Authentication Password

(SNMPv3 Only) Specifies the password for authentication using SNMPv3 and the selected authentication protocol.

Note: For SNMP profiles that are used with CA Network Flow Analysis make sure you specify a password that is at least eight characters long. If the password does not meet this requirement, the SNMP profile will be invalid and no SNMP data will be returned when the SNMP profile is used for polling. In this case the corresponding interfaces and devices in views, reports, and dashboards are missing device names, interface names, interface speeds, and interface utilization data.

Verify Authentication Password

Confirms the authentication password.

Privacy Protocol

(Optional) Specifies the encryption protocol to use for data flows sent to any devices or servers associated with this profile. Select one of the following protocols to create a valid SNMP profile for CA Network Flow Analysis polling:

- None (do not encrypt communications)
- DES
- AES (128-bit encryption)
- Triple DES

Note: If AES 192 and AES 256 options are listed, do not select either of those options: They are not supported for CA Network Flow Analysis. If the SNMP profile you create is not valid, no SNMP data is returned for devices and interfaces that use it. In this case the corresponding interfaces and devices in views, reports, and dashboards are missing device names, interface names, interface speeds, and interface utilization data.

The privacy protocol option is not enabled until authentication is enabled for this profile.

Privacy Password

Defines the password that is used when exchanging encryption keys. See the Note for a possible length requirement.

Verify Privacy Password

Defines the password used when exchanging encryption keys.

Use by default for new devices (CA PC)

Enable this profile for auto-discovery (NPC)

Specifies whether the profile is used by default to contact any new devices that are discovered from monitored traffic. If it fails, the next profile in priority order is used. Disable this parameter to exclude a profile from discovery.

5. Click Save.
6. You return to the list of SNMP profiles. The new profile appears in the list.

   Performance Center performs a global synchronization to send the profile information to CA Network Flow Analysis.

Verify IP Domains

The IP domains feature helps to address potential IP address conflicts. Domain identifiers indicate that two managed items that otherwise appear as duplicate IP addresses are actually two *different* managed items.

IP domains also let a global administrator in Performance Center control which managed items are visible to and accessible by a particular administrator or user.

Information about custom IP domains is sent down to the data sources during synchronization. Domains are available for use during configuration. You can use the NFA console Administration functions to add interfaces, custom virtual interfaces (CVIs), routers, and Harvesters to the custom domains that you create.

During initial setup, verify that the existing IP domains are adequate to monitor your environment. To see the domains, complete one of the following actions in the Performance Center Console:

- (CA PC) Select Admin, IP Domains, and review the domains on the Manage IP Domains page.
- (NPC) Select Admin, Groups, and expand the All Groups tree on the Manage Groups page.

View the IP Domains List

IP domains are required for monitoring multiple environments with overlapping IP addresses. Set up all the domains that you need before you begin to export flow data.

Follow these steps:

1. Log in to the Performance Center Console as a user with the required administrative role rights.
2. Display the current domains by completing one of the following actions:
   - (CA PC) Select Admin, Custom Settings: IP Domains.

     The Manage IP Domains page opens and displays the current domains.
   - (NPC) Select Admin, NetQoS Settings: Groups, and expand the All Groups tree on the Manage Groups page.

     The current domains are shown under All Domains. To display the parameters for a domain in CA NetQoS Performance Center, select the domain and click the Properties tab.

If you have not created any custom IP domains, only the Default Domain appears in the list. This predefined domain has a 'null' setting for all parameters.

Any custom domains that you have created include values for the following parameters:

Name

Identifies the domain.

Description

(Optional) Describes this domain namespace, such as naming the enterprise that owns it.

Primary DNS Address

Is the IP address of the primary name server for this domain.

Primary DNS Port

Is the port number that the primary name server uses.

Secondary DNS Address

Is the IP address of the secondary name server for this domain. Can be the same as the primary address.

Secondary DNS Port

Is the port number that the secondary name server uses.

Add Custom IP Domains

Administrators can set the domain assignment for Harvesters, routers, interfaces, and custom virtual interfaces (CVIs). We recommend that you set up any custom tenants and domains you need before you add the Harvesters.

Having the appropriate IP domains set up helps to achieve the following goals:

- Assign the correct tenant-domain when you add Harvesters so that their routers and interfaces inherit the correct associations. The routers have the appropriate SNMP profiles available to poll their interfaces.
- Make specific content accessible only to the operators who need to monitor it.
- Enable Administrators to create domain-specific ToS labels, protocol groups, and Autonomous System (AS) names in CA Network Flow Analysis.
- Avoid IP address conflicts.

  For example, suppose a router with a single IP address has interfaces that belong to different enterprises. The domain identifiers clarify that the interfaces are different managed items, even though they have the same IP address.

The Default Domain is created automatically. The Default Domain includes any items that are not assigned to a custom domain.

Follow these steps:

1. Log in to the Performance Center Console as a user with the required administrative role rights.
2. Display the current domains by completing one of the following actions:
   - (CA PC) Select Custom Settings: IP Domains.

     The Manage IP Domains page opens and displays the current domains.
   - (NPC) Select Admin, NetQoS Settings: Groups, then expand the All Groups tree on the Manage Groups page.

     The current domains are shown under All Domains.

   If you have not created any custom IP domains, only the Default Domain appears in the list.
3. Create a new domain:
   - (CA PC) Click New.

     The IP Domains Administration dialog opens.
   - (NPC) Right-click All Domains and select Add New Domain.

     The Add Domain dialog opens.

4. Supply information for the following parameters:

Domain Name

Identifies the domain.

Description

(Optional) Describes this domain namespace, such as naming the enterprise that owns it.

Device Name Alias

(CA PC only) Indicates the alias to use for a managed device. A device alias is a user-configured name that is applied to the associated managed item in CA Performance Center. Click Browse to navigate to and import a CSV file of aliases. The CSV file contains a comma-separated list of IP address-to-device alias mappings.

Aliases that are associated with the primary IP address of a device take precedence over aliases that are associated with any secondary IP addresses. Look for the primary IP address in the Address column of the Inventory Devices list. We recommend always using the primary IP address of the device in the CSV file.

For example:
`172.24.36.107,Austin Router`

Browse to select the file and click Open.

If you include aliases for devices you are managing already, it can take up to 5 minutes to begin synchronizing these aliases with CA Performance Center.

Note: To remove an alias, import a CSV file that includes the IP address for the device and a blank alias column. To change an alias, modify the alias entry in the CSV file and reimport the file.

Interface Description Override

(CA PC only) Indicates the alternate description to use for an interface. Interface descriptions appear in CA Performance Center already, but you can provide an alternate description. Click Browse to navigate to and import a CSV or TXT file of alternate descriptions. The file contains a comma-separated list of values that include the device IP address, interface name, interface description, and alternate interface description (alias) mappings.

For example:
`172.24.36.107,ethernet_7,vmxnet3 Ethernet Adapter,Connection to Dallas`

Note: Use the primary IP address of the associated device in the CSV or TXT file. Secondary IP addresses are not supported. Look for the primary IP address in the Address column of the Inventory Devices list.

Browse to select the file and click Open.

If you include alternate descriptions for interfaces you are managing already, it can take up to 5 minutes to begin synchronizing these descriptions with CA Performance Center.

Note: You can use the same alternate interface descriptions for more than one interface.

To remove an alternate description, import a CSV or TXT file that includes the IP address for the device, the interface name, the interface description, and a blank alias column. When you remove an alternate description, the original interface description reappears in CA Performance Center views.

Important! If you use a spreadsheet program to remove all of the alternate descriptions from a CSV file, include a column heading for the interface description override column in the imported file. If you do not include this column heading, the original interface descriptions will not reappear in CA Performance Center views.

To change a description, modify the alias entry in the CSV or TXT file and re-import the file.

DNS Settings check box

(CA PC only) If selected, displays the Primary DNS/Port and Secondary DNS/Port options.

Primary DNS Address

Is the IP address of the primary name server for this domain.

Primary DNS Port

Is the port number that the primary name server uses.

Secondary DNS Address

Is the IP address of the secondary name server for this domain. Can be the same as the primary address.

Secondary DNS Port

Is the port number that the secondary name server uses.

Enable DNS Proxy Address

(NPC only) Indicates whether the proxy address is enabled for this IP domain.

DNS Proxy Address

(NPC only) Is the IP address of the DNS proxy server.

This setting is required only if your network is located behind a DNS proxy server.

5. Click Save.

   The new IP domain appears on the page.

6. Repeat the steps as required to add more IP domains.

Configure Flow Collection

The next step is to configure the routers in CA Network Flow Analysis and verify that they are sending data to the Harvesters.

CA Network Flow Analysis can begin to collect flows as soon as you complete the following tasks:

- Recommended: Begin by adding the domains that you need (see page 23).
- Add the Harvesters (see page 26).
- (Three-tier distributed architecture Only) Add DSAs (see page 30)
- Configure the routers and interfaces to export flow data (see page 28).

Add One or More Harvesters

Add one or more Harvesters to enable data to be processed and displayed.

Prerequisite:

Recommended: If you have not already done so, register CA Network Flow Analysis and set up the domains before you add any Harvesters.

Follow these steps:

1. Open the NFA console, logged in with Administrator rights. For example, enter the following address in a browser:

   http://<ipaddress>/ra/

   User name: *admin*

   Password: *admin*

2. Open the Harvester page:

   a. Select Administration from the NFA console menu.

      The Administration page opens.

   b. Select System: Harvester from the Administration page menu.

      The Harvester page opens and displays the current list of Harvesters.

3. Click Add.

   The Add Harvester dialog opens.

4. Enter the following information:

   IP Address

   Address of the Harvester server.

   Description

   Identifying text about the Harvester, which appears in the Harvester page table.

   Domain

   Parent tenant and domain combination for the Harvester and for any routers and interfaces that begin to supply data to the Harvester.

   Changing this setting affects the tenant-domain association for new routers that begin exporting flow data and any new interfaces that begin generating flow.

   In a multi-tenant environment, the Harvester tenant affects which SNMP profiles are available for routers to poll interfaces.

   The domain affects which operators and reports have access to the data from routers and interfaces.

   This option is visible only in an environment that contains multiple domains.

5. Click Save.

   The new Harvester is added and appears in the Harvester list, provided that the IP address passes the connection test. If the test connection to the web service fails, an error message opens.

   The usual process is to add one or more Harvesters, then configure the router interfaces to export flow to the Harvesters. If you configure the routers to export flow to the Harvesters first, the NFA console immediately begins to collect data from the new Harvester. In this case, the domain for the routers is set at the time you add the parent Harvester.

Note: Make sure the Harvesters you add have not been deleted from the Harvester page previously. To add a Harvester instance successfully in CA Network Flow Analysis 9.3.0 after deleting it, the Harvester installation server must be re-imaged and the Harvester software must be re-installed.

Verify the Harvester Domain

Verify that each Harvester is associated with the appropriate domain before you set up routers to export flow data. If you have not already done so, set up any needed custom tenants and domains before you proceed.

Note: The tenant feature is applicable only to deployments that include CA Performance Center. If your deployment uses CA NetQoS Performance Center, the tenant setting is always Default Tenant.

Follow these steps:

1. Open the Harvester page:
   a. Select Administration from the NFA console menu.

      The Administration page opens.

   b. Select System: Harvester from the Administration page menu.

      The Harvester page opens and displays the current list of Harvesters.
2. Click Edit on the row for the Harvester that you want to edit.

   The Edit Harvester dialog opens.
3. (Optional) Change the Domain setting (tenant-domain combination) as needed.

   Default: Default Tenant \ Default Domain.

   You can also change the IP Address and Description.

   If no custom IP domains have been created, the Harvester table includes only the IP Address and Description columns.
4. Click Save when your changes are complete.

   Your changes are saved immediately.

Set Up the Routers

Enable NetFlow on each CA Network Flow Analysis router by completing the steps in this topic. You can configure routers to export any of the following flow protocols:

- NetFlow v5, v7, v9, and Sampled NetFlow
- sFlow version 5
- IPFIX, J-Flow, cFlow, and NetStream flow that complies with the standards for NetFlow v5, v7, or v9

Notes:

- Configure flow from each source to be exported to a single Harvester. If flow from one source is exported to multiple Harvesters, a number of problems result. If this occurs, contact CA Support for help.

- NetFlow provides a broad view of your network packet streams by creating flow records for all packets. The data from these flow records represents all packets. Sampled NetFlow/IPFIX and sFlow take samples from your packet streams, producing fewer flow records and lessening the impact to a collector. The lower your sampling rate, the less precise the data is likely to be.

In order for data from non-sampled flows to appear in reports of 15-minute (historical) data, the following minimum fields are required:

- One of the following: 1 - IN_BYTES, 85 - IN_PERMANENT_BYTES, 231 - FW_INITIATOR_OCTETS, or 232 - FW_RESPONDER_OCTETS
- 4 - PROTOCOL
- 7 - L4_SRC_PORT
- 8 - IPV4_SRC_ADDR
- 10 - INPUT_SNMP
- 11 - L4_DST_PORT
- 12 - IPV4_DST_ADDR
- 14 - OUTPUT_SNMP

Complete these tasks:

1. Back up the current router configuration.
2. Configure NetFlow export for each interface individually:
   a. Set the flow export version.
   b. Set the flow source IP address. Cisco recommends that you configure a loopback source interface. The IP addresses of non-loopbacked interfaces can change.
   c. Set the flow destination IP address and set the destination port to 9995. If you are using a custom value for the harvester listening port, use that value as the destination port. The port values must match or the Harvester does not receive flow data.
   d. Set the flow expiration timeout to 1 minute.
3. Enable flow for each interface.
   - NetFlow v5 or v5-compatible flow:
     - Monitoring multiple interfaces on a router: Use either all ingress or all egress. Use the same option for all of the interfaces. Ingress and egress values may vary slightly due to routers dropping packets and changing ToS values as traffic travels between interfaces.
     - Monitoring a single known interface on a router: Use ingress and egress. This option results in fewer total flows from the router to the Harvester and puts less load on the network and the Harvester.

- NetFlow v9 or v9-compatible flow:

  The Harvester identifies and deduplicates multiple flows on a single router, so you can use ingress and egress on multiple interfaces. You may find it most efficient to use this option for two or three interfaces. You have the option to enable ingress and egress across all interfaces, but this configuration may put an unnecessary burden on the Harvester.

4. Configure SNMP index persistence on each router that supports this feature.

Add DSAs (Three-Tier Deployment)

Applies to: Three-tier architecture in a distributed deployment: NFA console, Harvester, and DSA components installed on separate servers

If you have a three-tier architecture deployment, add one or more DSAs (Data Storage Appliances) to store the 15-minute (historical) data. If you have a stand-alone system or two-tier architecture, you do not add any DSAs.

We recommend that you add at least one DSA within 30 minutes of starting flow collection. Until you add a DSA to your three-tier deployment, 15-minute data is not available for reports. Reports that show a time range of more than 2 hours do not show any data.

Note: Do not add a DSA instance instead of editing the IP address of a retired DSA. In this case, routers continue to send data to the retired DSA--and that data is not available in reports. If you have to delete a DSA, contact CA Support for assistance.

Follow these steps:

1. Open the NFA console, logged in with Administrator rights.
2. Display the DSA page in the NFA console user interface:
   a. Select Administration from the NFA console menu.

      The Administration page opens.

   b. Select System: DSA from the Administration menu.

      The DSA page opens and shows a list of the current DSAs.
3. Click Add.

   The Add DSA dialog opens.
4. Enter the IP address for the DSA server.
5. Click Test Connection.

   A connection test is performed to determine whether the NFA console server can contact the DSA server and can verify that MySQL is installed. If the test succeeds, a "Test success" message opens.

6. Respond to the test results:
   a. If the Test success message opens, click OK to close it.
   b. If an error message opens, close the error message and respond as described in Troubleshoot DSA Addition (see page 32).
7. Once the test completes successfully, click Save in the Add DSA dialog.

   The connection test is performed, followed by a test to locate DSA settings on the target server.
8. Note the test results:
   - If no error message opens, the tests have succeeded. The following events result:
     - The dialog closes and the DSA is added to the DSA list.
     - The Harvesters begin to include the new DSA in the destinations for new enabled interfaces that report 15-minute data.
     - The NFA console pushes settings down to the new DSA.
     - The DSA is configured to retrieve 15-minute data files from the NFA console.
     - Data from the DSA begins to be available for reports in approximately 30 minutes.
   - If an error message opens, close the error message and respond as described in Troubleshoot DSA Addition (see page 32).

Notes:

- If a DSA does not begin to collect the 15-minute data within 30 minutes after flow collection begins, problems can result. The processed data accumulates on the NFA console server and processing slows. If the problem continues, the NFA console stops collecting the 15-minute data and unprocessed data accumulates on the Harvesters. If Watchdog traps are configured, the Watchdog sends out alerts that the Harvester or Reaper is falling behind. If the problem is left unchecked, the CA Network Flow Analysis services on the Harvesters may stop running.
- Each DSA is enabled to collect data for a maximum of 5,000 enabled interfaces that have reported data.
- For information about the data types, storage lifespan, minimum thresholds, and report types for the 15-minute data that is stored on DSA servers, see the topic 15-Minute Data (see page 170) in the *CA Network Flow Analysis Administrator Guide.*

Troubleshoot DSA Addition

Applies to: Three-tier architecture in a distributed deployment: NFA console, Harvester, and DSA components installed on separate servers

If one of the following error messages opens when you click Test Connection or Save in the Add DSA dialog, close the error message and respond as described.

Test Connection Errors

Use the following tips for troubleshooting error messages that may open when you click Test Connection in the Add DSA dialog

- "An invalid server IP address was entered":

  You entered an IP address in an invalid format. Make sure that you enter the IP address correctly.

- "System.Web.Services.Protocols.SoapException...'":

  Verify that the NetQoS MySql service is running on the NFA console server. If the service is not running, start it.

- "Unable to connect to any of the specified MySQL hosts":

  Start the NetQoS MySql service on the DSA server. Verify that the NetQoS MySql service is running on the DSA server. If the service is not running, start it.

- "Unknown database 'nqrptr'":

  The DSA database nqrptr was not found on the target server. Verify that the DSA software installation was successful.

Save Errors

Use the following tips for troubleshooting error messages that may open when you click Save in the Add DSA dialog

- "An existing record is already in use":

  Enter the IP address of a DSA that is not already in the DSA list.

- "Connection must be valid and open'":

  Verify that the NetQoS MySql service is running on the DSA server. If the service is not running, start it.

- "System.Web.Services.Protocols.SoapException...'":

  Verify that the target server is running and can be reached by the NFA console server. Verify that the NetQoS MySql service is running on the NFA console server. If the service is not running, start it.

- "Table 'nqrptr.settings' doesn't exist"

  The DSA settings table was not found. The DSA software was not installed successfully on the target server.

  - Verify that you entered the correct IP address for the DSA--not the IP address for a Harvester or for the NFA console.
  - Verify that the DSA software installation was successful.

Verify That Data Is Received

To verify that data is received, complete the following tasks:

- Verify that the interfaces are enabled (see page 33)
- Verify that the interfaces are visible in the NFA console (see page 34)

Verify That the Interfaces Are Enabled

Once you configure CA Network Flow Analysis to receive flow data, verify that the expected interfaces are monitored.

Follow these steps:

1. Open the NFA console, logged in with Administrator rights.
2. Open the Available Interfaces page:
   a. Select Administration from the NFA console menu.

      The Administration page opens.

   b. Select System: Enable Interfaces in the Administration menu.

      The Available Interfaces page opens.
3. Expand the router details to display the interface list: Click the arrow to the left of the router name.
4. Review the list to review whether the interfaces are enabled.
5. Change the interface status: Select the check box next to one or more interfaces, then click Enable or Disable.

   The selected interfaces are immediately enabled or disabled.
6. Repeat these steps for each router.

Verify That the Interfaces Are Visible in the NFA Console

Make sure that the configured interfaces are visible in the NFA console.

1. Verify that interfaces are visible on the Enterprise Overview page:
   a. Log in to the NFA console.
   b. Click Enterprise Overview in the main menu.
   c. Make sure the report views show interfaces.

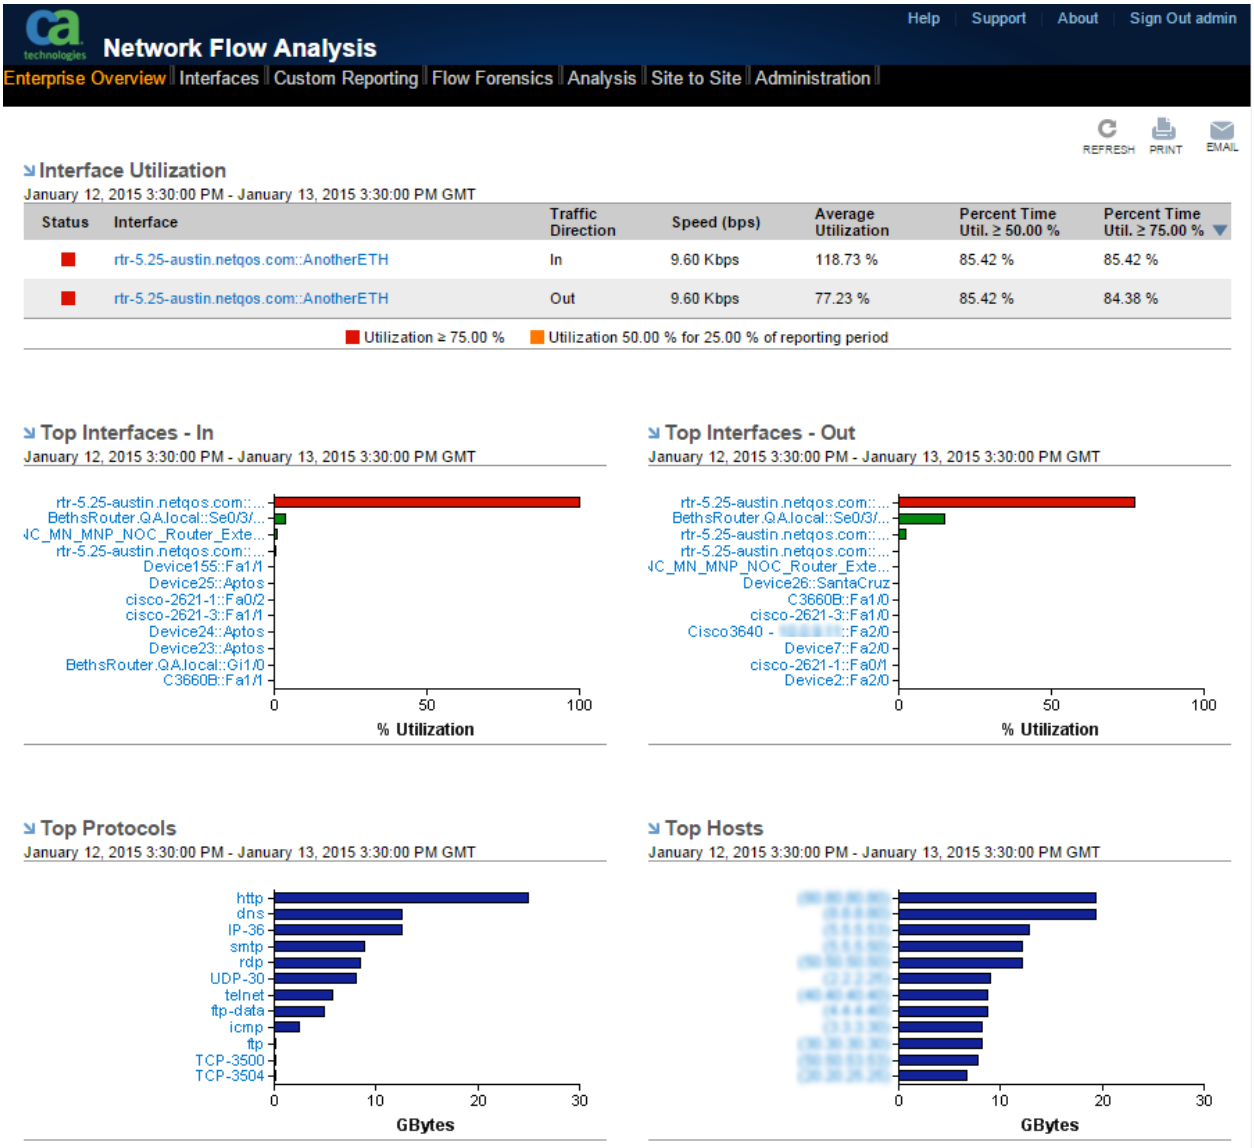

2. Verify that interfaces are visible in the Interface Index:
   a. Click Interfaces in the NFA console menu.

      The Interface Index opens.

   b. Make sure that the Interface Index includes the interfaces that you expect to see.

      You can locate interfaces by expanding router details to view interfaces. Alternatively, you can use the Search box to find routers or interfaces by entering all or part of a name or description.
3. If interfaces are not visible, perform the following preliminary troubleshooting tasks:
   - Verify that the CA Network Flow Analysis services are running on the NFA console server.
   - Review the logs. View the logs in the NFA console or open the logs from the <install_path>\reporter\Logs directory.

Change the Domain of Interfaces and CVIs

Interfaces and CVIs inherit their initial tenant-domain setting from the parent router and Harvester when the Harvester is added and the program begins to collect data from the router and interfaces. If the Harvester is not associated with a custom domain, the routers and interfaces are assigned to the Default Domain as their data begins to be collected.

You can edit the settings for interfaces and CVIs at any time. The domain setting does not have to match the parent router or Harvester. Changing the domain can affect which operators and reports have access to the interfaces' data.

Note: The tenant feature is applicable only to deployments that include CA Performance Center. If your deployment uses CA NetQoS Performance Center, the tenant setting is always Default Tenant.

Follow these steps:

1. Open the Active Interfaces page:
   a. Select Administration from the NFA console menu.

      The Administration page opens.

   b. Select Interfaces: Physical & Virtual from the Administration menu.

      The Active Interfaces page opens.

2. Locate and select the check box next to one or more interfaces that you want to associate with a tenant-domain.
   - To search for parent routers, interfaces, or CVIs, enter all or part of a router IP address, a router or interface name, or an interface description in the Search field, and then click Search. Expand the router details.
   - To navigate to an interface or CVI manually, go to the page that contains the parent router and click the arrow next to the router name. The router details expand to show the interfaces and CVIs.
3. Click Edit.

   The editing dialog opens. The Domain selection list is included in the dialog only if multiple domains exist.
4. Select a tenant-domain option from the Domain list.
5. Click Save.

   The dialog closes. The changes are shown on the Active Interfaces page.

Note: You can also change the tenant-domain setting for Harvesters (see page 86) and routers (see page 60).

Configure Traps

Trap configuration is complete when you have finished the following tasks:

- Create the traps that you need, as described in the *CA Network Flow Analysis Administrator Guide* "Create Traps" topic.
- Enable traps to be displayed as events in the Performance Center Console:
  1. Open the Application Settings page in the NFA console.
  2. Set the Trap Destination value to match the IP address of one of the following servers:
     - (CA PC) NFA console or stand-alone server that is registered as a data source
     - (NPC) Event Manager server
- (Optional) Enable Watchdog trap notifications to be sent to your trap receiver: Open the Watchdog Settings page in the NFA console. Configure values for the Trap Destination, Email Address, and other Watchdog settings.

- (Optional) Verify that the events are displayed on the Performance Center Console--on the Events page (CA PC) or the Event List page (NPC). If the events are not shown as expected, verify that the following conditions are met:
  - The logs show that events have been generated and have been forwarded to the Event Manager.
  - The Event Manager host name is resolvable by the DNS server for CA Network Flow Analysis.
  - The Trap Destination value on the Application Settings page in the NFA console matches the IP address of one of the following servers:
    - (CA PC) NFA console or stand-alone server that is registered as a data source
    - (NPC) Event Manager server
  - (NPC Only): The Event Manager is installed.

Set Up User Accounts

One predefined user account, admin, is included with the installation of CA Network Flow Analysis. The admin account has full administrative privileges.

The administrator must create a user account for each person who will use the product--administrators and operators. Custom user accounts enhance security and take advantage of the narrowly defined role rights that determine access to product features and data.

Custom user accounts are best deployed in a well-planned system that includes custom groups. Custom groups are assigned as permissions to let product operators view only the data, menus, and dashboards that they need to perform their daily tasks.

View a List of User Accounts

You can see a high-level overview of user account settings in the Performance Center Console. Until you create custom user accounts only the two factory user accounts are available.

Follow these steps:

1. Log in to the Performance Center Console as a user with the required administrative role rights.
2. Click Admin, User Settings: Users.

   The current list of user accounts is displayed on the Manage Users page (CA PC) or the User List page (NPC).

The table includes the following information about each user account:

Name

Is a login name for the user account.

Role

Is the role assigned to the user account.

CAPC Privilege / NPC Privilege

Identifies the level of access to registered data sources, such as CA Network Flow Analysis.

Permission

Lists the permission groups that are assigned to this account. You can view permission groups as nested locations in the Groups tree.

Default: '/All Groups'.

Status

Indicates whether the user account is enabled or disabled.

Note: Additional status values may be listed in the Console for CA NetQoS Performance Center: Built-in (configured automatically) and Online (currently logged in).

To perform any action on this page, click one of the buttons along the bottom.

Add User Accounts

Add a user account for each person who will operate the products. For security purposes, operators should not share user accounts.

Note: This topic describes the steps to perform this task in CA Performance Center. If you register CA Network Flow Analysis as a data source for CA NetQoS Performance Center, the steps are slightly different.

Follow these steps:

1. Log in to the Performance Center Console as a user with the required administrative role rights.
2. Confirm that the required roles and groups exist.
3. Select Admin, User Settings: Users.

   The current list of user accounts is displayed.
4. Click New.

   The Create New User wizard (CA PC) or Add User page (NPC) opens.

5. Enter information for the following account parameters:

   Name

   Is a login name for the user account. Limited to 50 characters.

   Description

   (Optional) Describes the user account to help you identify it.

   Email Address

   (Optional) Associates an email address with the user account.

   Preferred Language (CA PC only)

   Specifies the language spoken by the operator associated with the user account.

   Authentication Type

   Identifies the authentication method that applies to this user account. The method must match Single Sign-On configuration. Select one of the following:

   - Performance Center (CA PC) or Product (NPC) —The default authentication scheme deployed by Performance Center.
   - External—A third-party authentication scheme, such as LDAP or SAML.

   Password

   Defines a password for the user account. The password is limited to 32 characters.

   Time Zone

   Corresponds to the time zone in which the user will view data.

   Default: UTC (Coordinated Universal Time).

   Role

   Is the role assigned to the user account.

   Account Status (CA PC only)

   Determines whether the account is activated for use (Enabled).

   User Options (NPC only)

   Determines whether the account is activated for use (Enabled) and whether the user is allowed to share views with other products (Allow user to generate view URLs).

6. Assign access permissions to the user, as described in Assign Permission Groups to User Accounts (see page 42).

7. Add permission groups to the user account, described in Assign Product Privileges (see page 43).
8. Click Save.

   The new user appears in the list of user accounts.

Role Rights

The rights that are assigned to each role determine user access to dashboards and menus. For example, role rights control the types of views that users can see and control whether users can export data, customize settings, and set up schedules for sending email reports.

Administrators can grant additional rights to users by editing their role. The Edit Role dialog lists role rights that are assigned to roles. The Manage Users or User List page shows the role that is assigned to each user.

Note: Do not remove the administrative role rights from your primary administrator account. Administrative access to the console is required.

Add Role Rights for Users

If the predefined user roles do not fit your requirements, you can add custom user roles. Ideally, you create the roles that each unique product operator needs to be able to perform his or her job responsibilities.

Custom roles work best within a system of custom groups. Custom groups let you precisely grant access to dashboards and product features while restricting access to sensitive data. The same groups that you create to organize data can serve as "permission groups" when you set up user account permissions.

A new role has no rights until you add them. The following graphic shows the Add Role dialog in the CA Performance Center Console with a role that is beginning to be defined.

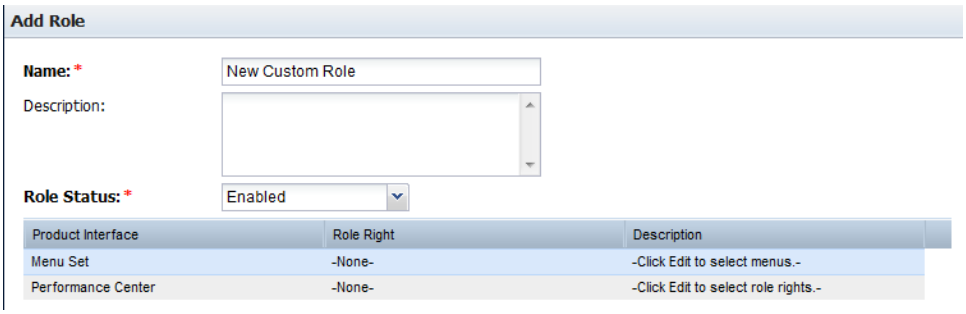

Follow these steps:

1. Log in to the Performance Center Console as a user with the required administrative role rights.
2. Navigate to the Manage Roles or Roles List page.

   The page displays the current list of roles.
3. Click New.

   The Add Role dialog opens.
4. Supply the required information and make selections in the fields provided:

   Name

   (Optional) Identifies the role. Limited to 45 characters.

   Description

   (Optional) Describes the role. For example, identifies the job-related duties that the associated user performs.

   Enable Role

   Enables the role to make it active. Required to give users with this role the access granted by role rights.
5. Specify the menus that will be visible to users with the new role:
   a. Select Menu Set (CA PC) or select a menu or product from the list at the bottom of the dialog (NPC).
   b. Click Edit.

      The Edit Menu Set dialog opens. Menus in the 'Available Menus' list can be added to the role.
   c. Click an item on the left that you want to add to the role, then click the right arrow.

      Use Shift + Click or Ctrl + Click to select multiple items.

      Each selected item moves to the Selected Menus list.
   d. (Optional) Use the Up and Down arrows to move items around in the list. The order of menus in the list determines their order on the Dashboards tab.
   e. Click OK.

      You return to the Add Role page.
6. Set the Performance Center rights for the role:
   a. Select Performance Center (CA PC) or NetQoS Performance Center (NPC).
   b. Click Edit.

      A dialog opens, which you use to select Performance Center access rights.

c. Click an item on the left that you want to add to the role, then click the right arrow.

   The access right moves to the Selected Rights list.

d. (Optional) Use the Up and Down arrows to move items around in the list. The order of role rights determines their priority in cases where rights overlap.

e. Click OK.

   You return to the Add Role page.

7. Set the CA Network Flow Analysis rights for the role:

   a. Select the name of the registered CA Network Flow Analysis instance.

   b. Click Edit.

      A dialog opens, which you use to select access rights for CA Network Flow Analysis in the same way you selected access rights for Performance Center.

   c. When the access rights are set up correctly, click OK.

      The new role is created and appears in the Role List.

8. Repeat the previous step to set the rights for any additional data source that you want to include.

9. Click Save on the Add Role page.

   You return to the Manage Roles page (CA PC) or Roles List page (NPC).

Note: When you finish creating a role, assign it to a user account as a separate step. Roles are inoperative until they are assigned to user accounts. Only users with the 'Administer Users' and 'Administer Roles' role rights can assign roles to user accounts.

Assign Permission Groups to User Accounts

Individual operators require data access permissions to monitor data in the products. Access permissions are based on groups. You can assign access permissions according to your plan for custom groups. Your goal as the administrator is to make sure that all operators see only the data they require to do their job.

For example, suppose you create custom groups and assign them as permissions to IT staff. When staff members log in to Performance Center, they can view data from the systems that are assigned to them.

Follow these steps:

1. Log in to the Performance Center Console as a user with administrative privileges.

2. Click Admin, User Settings: Users.

   The Manage Users page (CA PC) or User List page (NPC) opens.

3. Select a user account that you want to change, and click Edit.

   The Edit User wizard or dialog opens.

4. Display the permission groups:

   - (CA PC) Click the Access Permissions button.
   - (NPC) Locate the Permission Groups pane in the middle of the page.

   The group settings are displayed.

5. Add permission groups to the user account

   - Expand the groups in the Available Groups tree on the left so that subgroups are shown.
   - Select a group or subgroup.
   - Click the right arrow or Add button to add the group.
   - Repeat as necessary.

   The selected permission groups appear in the Selected Groups pane.

6. Select the default group for the user--the data that appears by default in the dashboards for the user:

   - (CA PC) Right-click the target group and select Make Default.
   - (NPC) Select the target group and click Make Default.

7. Click Save.

   The changes are saved to the user account, and you return to the Manage Users page.

   When the user logs in, data from the default group appears in dashboards by default.

Assign Product Privileges

Each registered data source has its own product privilege setting, which grants unique privileges within that product interface. Administrators give users product privileges for each data source. For example, the product privilege determines whether a user can log in to CA Network Flow Analysis or drill down from a Performance Center view to details in the NFA console. Privileges are specific to the data source instance.

The default administrator account, admin, is locked to prevent changes to product privileges. This account must have Administrator privileges for all registered data sources. If you select a group of accounts that includes the admin account, you cannot edit the product privileges for any of the selected accounts.

CA Network Flow Analysis Product Privileges

A user must have product privileges for the CA Network Flow Analysis data source to log in to the NFA console. Product privileges also determine whether a user can access the Administration page, and can perform certain functions:

Administrator

Gives access to the Administration page in the NFA console and to all functions. Functions include creating and managing user accounts, roles, groups, SNMP profiles, and scheduling for reports.

Power User

Gives user-level access and any additional abilities that the Role setting grants. For CA Network Flow Analysis, the Power User privilege is equivalent to the Administrator privilege.

User

Gives access to Top Interfaces reports and Interface Utilization reports on the Enterprise Overview page.

A User with the appropriate Permission Group settings also has access to the following reports:

- Top Hosts and Top Protocols reports on the Enterprise Overview page, if the user also has access to All Groups
- Interfaces page reports for the interfaces that are accessible to the user
- Existing reports on the Custom Reporting, Flow Forensics, and Analysis pages
- Menus that an administrator has assigned to the User role

The Role and Permission Group settings determine whether the User also can run existing reports, create reports, and manage reports. To create reports, a User must have access to All Groups.

None

Has no access to a data source. The user who has this product privilege cannot log in to the NFA console or drill down from a Performance Center view to the NFA console. By default, all users have this product privilege setting for all data sources.

Note: The same user account can have different privileges for different data sources.

Set Up Groups

We recommend that you create custom groups to help manage items in the Performance Center Console. Custom groups are required to let operators see performance data from the routers they manage.

Properly configured, groups can prevent operators from viewing particular types of data for security reasons. The administrator can selectively grant users access to data in their area of responsibility, such as a physical location or subnet.

Create Custom Groups

Before you start creating groups, plan a strategy and a structure. Consider the types of access permissions that operators require to perform their monitoring duties. If necessary, you can discuss your organizational and monitoring goals with a CA technical representative.

Create groups under the All Groups node in the Groups tree, or within an existing custom or site group. You cannot add groups to system groups, which appear "locked" in the Groups tree.

You can add a maximum of 2000 child groups to a parent group.

Follow these steps:

1. Log in to the Performance Center Console as a user with the required administrative role rights.
2. Navigate to the Manage Groups page.

   The page displays current groups in a tree structure.
3. Expand nodes in the Groups tree to find a location for the new group.
4. Right-click the node, and select Add Group (CA PC) or Add New Group (NPC).

   The Add Group window opens with the New tab selected by default.
5. Supply values for the following parameters:

   Group Name

   Specifies a name for the group. Do not use the following special characters in group names: /&\,%.

   Description

   (Optional) Helps you identify the group.

6. Confirm the setting for the following parameter:

   Include the children of managed items

   Adds the children of managed items automatically when the items are added to this group. If this option is disabled and you add a router, the router interfaces are not included. As a result, the data from those interfaces is not visible in drilldown views.

   Default: Selected (CA PC) or Not Selected (NPC).

   Note: Clear this option for a custom group that contains routers or the group will not be usable in in the NFA console.

7. Select Custom or Site from the Group Type list.

   If you selected Site as the type, specify values for the additional parameters that appear, including Location.

8. Click Save.

   The new group appears in the Groups tree.

   The group contains no items until you add them. You have two options for adding items to a custom group:

   - Manually populate the group by adding items in the Manage Groups interface.
   - Create rules to manage group membership.

Results of Unregistering

On rare occasions, you may want to unregister CA Network Flow Analysis. For example, you would unregister a CA Network Flow Analysis instance before registering it with a different Performance Center instance. Unregister only if it is really necessary.

If you unregister CA Network Flow Analysis, the following rules apply:

- Users: Users who are not associated with CA Network Flow Analysis are deleted. Existing User IDs remain unchanged. You cannot add new users or edit user account settings while unregistered.
- Roles: Roles are not deleted. Users continue to have their previous roles. Existing Role IDs remain unchanged. You cannot change the roles or permissions for existing users while unregistered.

- Groups:
  - Groups that do not exist in CA Network Flow Analysis are deleted. You cannot add or change groups while unregistered.
  - Nested groups that are associated with an interface are displayed as interface groups in the NFA console.
  - Groups that are not associated with an interface are displayed as permissions.
- Single Sign-On and LDAP: Single Sign-On and LDAP values remain unchanged.

Note: For a more detailed description of the results of unregistering during an upgrade of CA Network Flow Analysis, see the *CA Network Flow Analysis Upgrade Guide*.

Complete Post-Installation Tasks

You may want to complete a number of additional administrative tasks, depending on your environment and on the number of users. You should complete some tasks as soon as CA Network Flow Analysis is running. Consider the following tasks:

- **Verify that required settings are complete**

  Make sure that you have completed all the installation tasks. Certain settings are required to enable functionality, such as emailing reports and triggering SNMP traps.

- **Adjust settings to improve performance**

  Certain report settings can affect performance, such as the frequency of DNS host name resolution.

- **Check interface speeds**

  If interfaces are at or above 100 percent utilization in the NFA console (for example, as seen on the Enterprise Overview page), consider adjusting the interface speeds. Additional information about editing interface settings is in the topic Expire Stale Addresses (see page 151).

- **Control router display**

  You can exclude domains from the display to control which routers are displayed in the NFA console.

  Note: This action is applicable only in a deployment that includes multiple domains.

- **Disable monitoring of router-generated traffic**

  If you do not want reports to show broadcast traffic for each router, change the Pump Broadcast/Multicast application setting to False.

- **Monitor product components**

  The Watchdog services help you monitor components. Make sure that the appropriate settings are configured to notify you of component issues as soon as possible. Specify thresholds, an email address for receiving messages, and other settings. Additional information is in the topic How to Monitor Network Flow Analysis Components (see page 142).

- **Adjust security settings**

  To export reports to *comma-separated value* (CSV) files successfully, the security settings in your environment may make it necessary to add the IP address of the NFA console to your list of trusted sites.

Note: For installation and upgrade information, see:

- *CA Network Flow Analysis Installation Guide*
- *CA Network Flow Analysis Upgrade Guide*

Chapter 3: Administration Page Options

The CA Network Flow Analysis Administration page lets you review, administer, and customize the network data view. When you open this page, the System Status is displayed. The System Status uses the following status icons:

- *Green* check mark: Component is running without any errors.
- *Red* exclamation mark: Component is running but has errors.

Click a red exclamation mark to view the associated error report. The error report includes the IP address, the type of component, and more information about the error. No error messages are displayed for a green check mark.

Notes:

- To perform administration tasks, log in as a user who has Administrator rights.
- A few options take you to a page in the Console for the Performance Center version that you use with CA Network Flow Analysis. If the product is currently registered as a data source for CA Performance Center or CA NetQoS Performance Center, the following options are enabled:
  - CA PC Groups or NPC Groups
  - Users
  - Roles
  - SNMP Profiles

A menu of administrative options is on the left side of the page. The following list describes the pages and functions that correspond to the options.

Interfaces Menu

Physical & Virtual

View the status of the interfaces on the Active Interfaces page. You can also edit, delete, or merge interfaces or create and edit custom virtual interfaces (CVIs).

Aggregations

Aggregate interfaces on the Interface Aggregations page. Interface aggregations let you view and report on interfaces as a unified group.

Alerts Menu

Alerts

Review, add, edit, and delete traps on the Trap Configuration page.

Define an Application: Groups Submenu

Application Definitions

Create and edit rules for application mapping, port priorities, or Reserved Seating on the Application Definitions page.

Protocol Names

Use the Protocol Configuration page to edit protocol names and descriptions. In a deployment that has multiple domains, protocol names are domain-specific.

ToS Names

Use the ToS Configuration page to edit a ToS label (description) in the context of a particular domain. You also can add, remove, or edit groups of ToS values. In a deployment that has multiple domains, ToS names are domain-specific.

AS Names

Use the AS Names page to search and edit Autonomous System names in the context of a particular domain. In a deployment that has multiple domains, AS names are domain-specific.

ToS Groups

Review, add, delete, or edit ToS groups on the ToS Group Configuration page.

Protocol Groups

Review, add, delete, or edit protocol groups on the Protocol Group Configuration page.

Administration: Reporting Submenu

CA PC Groups or NPC Groups

Open the Manage Groups page in the Performance Center Console. Review system groups and add, remove, or edit custom groups and site groups. This option is enabled if CA Network Flow Analysis is registered as a data source for Performance Center.

Reporting Periods

Use the Reporting Periods Configuration page to add, edit, and delete reporting periods. Operators can use the reporting periods to limit the time frames for the data in Interface reports.

Time Filters

View, add, edit, and delete time filters on the Time Filter Configuration page.

Scheduled Emails

Use the Scheduled Emails page to review the reports that are scheduled for email delivery. You can change the destination address, Subject line, accompanying message text, and schedule options.

Addresses

Use the Address-Hostname Configuration page to specify a mask and options for resolving IP addresses to DNS names. You can list, edit, delete, and expire IP addresses. Expiring IP addresses schedules them to be refreshed. In a deployment that has multiple domains, address configuration is domain-specific.

Administration: Authentication Submenu

Users

Open the page for managing user accounts in the Performance Center Console. Review the user accounts and their settings. Add, remove, and edit user accounts. This option is enabled if the product is registered as a data source for Performance Center.

Roles

Open the page to manage roles in the Performance Center Console. Review the existing role names and the capabilities and menus that are available for each role. Add, delete, and edit roles. This option is enabled if the product is registered as a data source for Performance Center.

Administration: System Submenu

Enable Interfaces

View the router list and status of routers on the Available Interfaces page. Enable, disable, or delete interfaces on this page.

Harvester

Use the Harvester page to view, add, edit, and delete the Harvesters that generate report data.

 **DSA**

Three-tier architecture deployments only: Use the DSA page to add a DSA to your configuration or edit the DSA IP address. In a three-tier architecture deployment, one or more DSAs store the 15-minute (historical) data for reports.

Application Settings

Edit a wide range of application settings on the Application Settings page.

SNMP Profiles

Open the page to manage SNMP profiles in the Performance Center Console. View, add, edit, delete, and re-order the SNMP profiles that Harvesters use for polling. In a multi-tenant environment, a Harvester uses the SNMP profiles that are assigned to its tenant. This option is enabled if the product is registered as a data source for Performance Center.

Templates

Use the Interface Templates page to view, add, edit, and delete the templates that determine the way interface names and descriptions are displayed.

Administration: Health Menu

System Status

View the overall status of the CA Network Flow Analysis components on the System Status page. Click a warning icon to display a list of the problem reports for the component.

 Data Storage Appliances (DSAs) status is useful only for three-tier architecture deployments. DSAs are not included in two-tier deployments. In a two-tier deployment, the DSA status always shows a green system status icon.

Watchdog Settings

View and edit the Watchdog configuration settings on the Watchdog Settings page.

Anomaly Detector Menu

Anomaly Detector

Display the Anomaly Detector window to perform administrative functions for CA Anomaly Detector. These functions have been added to the NFA console to accommodate users in a deployment that includes CA Performance Center.

Notes:

- If your deployment includes CA NetQoS Performance Center, the same functions are available in that Console. The functions in both locations save your settings to the same database.

  For a CA NetQoS Performance Center deployment, we recommend that you use the functions in the CA NetQoS Performance Center Console. Working in a single location helps to reduce the possibility of multiple users writing to the database simultaneously. If that happens, unexpected results can occur.

The Anomaly Detector window gives you access to the following tabs:

- View Monitored Products: Add products for CA Anomaly Detector to monitor and review the list of monitored products.

  The other page functions are enabled as soon as you add an instance of CA Network Flow Analysis to be monitored.

- View Collection Sources: View, enable, and disable the collection sources that the CA Anomaly Detector monitors.
- View Sensors: View and edit the default configuration settings for the sensors.
- View Alert Targets: Configure alert targets for snmp_traps and syslogging.

For more information about the functions, see the *CA Anomaly Detector Guide.*

About Menu

About

Display the Version Information page to view the product version number and installation date. This page also contains links to the product version history, the Administrator Guide, and the Operator Guide.

Chapter 4: Working with Interfaces and Routers

You use a number of different pages to perform the tasks for interfaces and routers:

Active Interfaces Page

The Active Interfaces page shows interfaces, routers, custom virtual interfaces (CVIs), and interface aggregations. The page shows only the routers and interfaces that have contributed to the collected data at some time. You can create and delete CVIs, merge interfaces, and edit properties. You also can delete routers and interfaces without deleting them from the system entirely--for example, to address capacity issues.

Select Interfaces: Physical & Virtual from the Administration page, then use the Active Interfaces page to perform the following main tasks:

- View details (see page 56).
- Edit details (see page 60).
- Delete interfaces (see page 63) or routers (see page 61) from the page.
- Create custom virtual interfaces (CVIs) (see page 66).
- Merge interfaces (see page 69).

Available Interfaces Page

The Available Interfaces page shows information about all of the interfaces and routers, including the ones that have not contributed to the collected data at any time. You can enable or disable interfaces, delete routers and their interfaces from the system, and perform some polling troubleshooting.

Select System: Enable Interfaces from the Administration page, then use the Available Interfaces page to perform the following tasks:

- Check polling and view details about routers (see page 72) and interfaces (see page 74).
- Enable or disable interfaces (see page 75).
- Delete end-of-life routers from the system (see page 76).

Interface Templates Page

You can create a custom interface template to change the way interface names and descriptions appear in a number of reports.

Select System: Templates from the Administration page, then use the Interface Templates page to perform the following tasks:

- Create and apply a template (see page 77)
- Edit a template (see page 79)

Application Settings Page

You can use a setting on the Application Settings page to control whether device names are included in interface names.

Select Application Settings, then use the Application Settings page to perform the following task:

- Add or remove the device name from interface names (see page 80)

Active Interfaces Page

Use the Active Interfaces page to view the routers, interfaces, custom virtual interfaces (CVIs), and interface aggregations.The page shows only the routers and interfaces that have contributed to the collected data at some time. You can perform the following actions on the Active Interfaces page:

- Review the routers, interfaces, and other elements without seeing the routers and interfaces that have never contributed to the collected data.
- Create and delete CVIs.
- Merge interfaces.
- Delete interface aggregations.
- Delete routers and interfaces to reclaim capacity without deleting them from the Available Interfaces page.
- Edit properties, such as:
  - Router name, domain, assigned SNMP profile/version, port, and interface naming template
  - Interface name, description, speed, type, and domain
  - CVI name, description, speed, type, domain, and subnets

Follow these steps:

1. Open the Active Interfaces page:
   a. Select Administration from the NFA console menu.

      The Administration page opens.

   b. Select Interfaces: Physical & Virtual from the Administration menu.

      The Active Interfaces page opens.

The Active Interfaces page includes the following options:

Search

Use a text string to search for routers, interfaces, or other elements with a matching address or name (see page 59).

Edit

Modify properties of the selected item or items (see page 60).

Delete

Remove routers (see page 61) and interfaces (see page 63) from the page--for example, to address capacity issues.

Merge

Combine the data for two selected physical interfaces or CVIs (see page 69).

Add Custom Virtual Interface

Create a custom virtual interface based on the selected physical interface.

Active Interfaces: Router Information

Routers are listed in tables on the Active Interfaces page. The router tables contain the following columns.

Flow Status 

Status of the router at the most recent scheduled polling attempt:

- Red: Any enabled interfaces have not had flow for longer than the **Interface Data Absence Limit**.
- Yellow: Any enabled interfaces have not had flow between 30 minutes and the **Interface Data Absence Limit**.
- Green: All enabled interfaces have had flow in the last 30 minutes.

Router Address

IP address of the router

Router Name

User-assigned name of the router, such as *lab1 router*

Template

Interface naming template that is assigned

Interfaces

Total number of interfaces, interface aggregations, and CVIs for the router. This count includes only the interfaces that have contributed to the collected data at some time. To review all interfaces, go to the Available Interfaces page (see page 71).

Harvester

IP address of the Harvester that collects data from the router

Active Interfaces: Interface Information

Interfaces and custom virtual interfaces (CVIs) are listed on the Active Interfaces page in tables nested under their parent routers. To display the interfaces and CVIs for a router, click the arrow next to the router name. Interface tables contain the following columns.

Traffic Status 

Status of the interface at the most recent poll:

- Red: Inactive interface
- Green: Active interface

Class

Icon for the interface type, such as a physical interface, CVI, or interface aggregation. To display the interface type name, position the cursor over the icon.

Interface Name

User-assigned name of the interface, CVI, or interface aggregation

Description

Optional information for identification

If Index

Index value, which is assigned by the device that sends flows to the interface

Type

Connection type, such as WAN or LAN

In Speed

Inbound speed of the interface, provided the speed is known

Out Speed

Outbound speed of the interface, provided the speed is known

Domain

Domain of the interface. If the environment has only one domain, the Domain column is not visible.

Notes

Link to add, edit, or view notes about an interface. If the note is empty, the Notes icon is dimmed. For example, you could add information about the time zone, business unit, geographical location, alternate bandwidth, circuit ID, or history. Notes icons are dimmed until they are populated and the page is refreshed.

To display the Notes icon, set the Display Notes Field value to True on the Application Settings page.

Search for a Router or Interface

Use the Search function on the Active Interfaces page to locate routers, interfaces, or CVIs. To search for aggregations, use the Search function on the Interface Aggregations page.

Follow these steps:

1. Open the Active Interfaces page:
   a. Select Administration from the NFA console menu.

      The Administration page opens.

   b. Select Interfaces: Physical & Virtual from the Administration menu.

   The Active Interfaces page opens.
2. Enter a text string in the Search field. Search for whole or partial text strings that match the addresses, names, or descriptions of routers, interfaces, or CVIs.

   Note: Do not use wildcards.
3. Click Search.

   The list is filtered to display only the matching entries. If you search for interfaces or CVIs, the search returns a list of routers that contain matching items. When you expand router details, all items appear in the sublists.

To clear the search, click Clear Filter.

Edit Router and Interface Details

Edit the properties of a router or interface. For example, you can edit properties to make corrections or to supply missing information.

Follow these steps:

1. Open the Active Interfaces page:
   a. Select Administration from the NFA console menu.

      The Administration page opens.

   b. Select Interfaces: Physical & Virtual from the Administration menu.

      The Active Interfaces page opens.
2. Locate the routers, interfaces, or other elements that interest you. To display the contents under a router, click the arrow next to the router name.
3. Select the check box next to the items you want to edit. You can edit multiple interfaces and CVIs at the same time or edit multiple routers at the same time. You cannot edit a mix of routers, interfaces, and CVIs at the same time.
4. Click Edit.

   The dialog box opens for editing the selected item or items.
5. Modify the properties as needed, such as:
   - Router name, tenant-domain (if applicable), SNMP profile/version, port, and template

     (CA PC Deployment) The options to edit the SNMP profile and port are enabled once a tenant-domain is selected.
   - Interface name, description, speed, connection type, and domain (if applicable)

   Notes:

   - You can also edit properties for CVIs and aggregations.
   - If you select multiple elements to edit, editing options are restricted to their shared properties.
   - The tenant-domain setting for interfaces does not have to match the setting for the parent router. Changing the domain can affect which operators and reports have access to the related data.
   - Changing the tenant of a router in CA Performance Center can affect which SNMP profiles are available for polling the router's interfaces. This is not applicable to CA NetQoS Performance Center, which uses the same list of SNMP profiles for all routers.

6. Click Save.

   The Edit Router or Edit Interface dialog closes. The changes are shown on the Active Interfaces page.

Notes:

- If you delete an interface without disabling it, the interface will be added again automatically at any time that it begins sending flow.
- For information about deleting interfaces and routers, see:
  - Delete an Interface (see page 63)
  - Delete a Router on the Active Interfaces Page (see page 61)
  - Delete a Router on the Available Interfaces Page (see page 76)

Delete Routers from the Active Interfaces Page

You can delete a router on the Active Interfaces page, but leave it in place on the Available Interfaces page. For example, you may want to delete a router in this way to address capacity issues. Later, you can restore the router (but not its historical data) by enabling its interfaces on the Available Interfaces page.

Deleting a router removes its interfaces, CVIs, 15-minute (historical) data, and traps. The deletion also affects any related aggregations, views, and reports.

Note: If you delete the router from the Available Interfaces page, the system has no further record of the router.

Follow these steps:

1. Verify that the router is no longer sending flows to the product.
2. Disable the router interfaces, if they are not disabled already:
   a. Open the Available Interfaces page:
      - Select Administration from the NFA console menu.

        The Administration page opens.
      - Select System: Enable in the Administration menu.

        The Available Interfaces page opens.
   b. Locate the router by using the Search function or by paging through the table contents.
   c. Select the check box next to the router.

      You can select and disable multiple routers simultaneously provided the selections are on the same page.

d. Click Disable.

   A confirmation message opens.

e. Click Yes.

   The Enabled status for the interfaces changes to 'No.' New data from the interfaces is no longer collected or shown in reports. Data that is already collected is still available for reports.

f. (Optional) Update the Enabled value for the router by refreshing the page--for example, by pressing F5.

3. Open the Active Interfaces page:

   a. Select Administration from the NFA console menu.

      The Administration page opens.

   b. Select Interfaces: Physical & Virtual from the Administration menu.

      The Active Interfaces page opens.

4. Locate the router and select its check box.

5. Click Delete.

   A confirmation message opens.

6. Click Yes.

   The following events result:

   - The confirmation message closes.
   - The router is deleted from the Active Interfaces page and stops consuming capacity.
   - All related interfaces, CVIs, 15-minute (historical) data, and traps are deleted.
   - The router is deleted from any related aggregations.
   - The data for the deleted interfaces no longer appears in the NFA console or in reports.

Note: A router that you delete from the Active Interfaces page reappears if the router interfaces are enabled on the Available Interfaces page and the interfaces begin to send flow again.

Also See:

Delete Routers from the System (see page 76)

Delete Interfaces

If you delete an interface on the Active Interfaces page, its historical data, related CVIs, and traps are deleted. The deletion also affects any aggregations, views, and reports that previously included the interface.

Follow these steps:

1. Disable the interface, if it is not disabled already:
   a. Open the Available Interfaces page:
      - Select Administration from the NFA console menu.

        The Administration page opens.
      - Select System: Enable in the Administration menu.

        The Available Interfaces page opens.

   b. Locate the interface: Use the Search function or expand the contents of the routers.

   c. Select the check box next to the interface.

      You can select and disable multiple interfaces simultaneously, including interfaces with different parent routers. All of your selections must be on the same display page.

   d. Click Disable.

      The Enabled status is changed to 'No' for the interface. New data from the interface is no longer collected or shown in reports, but data that is already collected is still available for reports.
2. Open the Active Interfaces page:
   a. Select Administration from the NFA console menu.

      The Administration page opens.

   b. Select Interfaces: Physical & Virtual from the Administration menu.

      The Active Interfaces page opens.
3. Locate the interface and select its check box.
4. Click Delete.

   A confirmation message warns you about the results of the deletion.

5. Click Yes.

   The following events result:

   - The confirmation message closes.
   - The interface is deleted from the Active Interfaces page, but not from the Available Interfaces page.
   - All related historical data is purged permanently.
   - Data for the deleted interface no longer appears in NFA console views and reports.
   - All related CVIs and traps are deleted.
   - The interface is deleted from any related aggregations.
   - The interface stops consuming capacity on the server that stored its data.

Also See:

Edit Router and Interface Details (see page 60)
Enable or Disable Interfaces (see page 75)
Delete Routers from the Active Interfaces Page (see page 61)
Delete Routers from the System (see page 76)
Steps for Merging Interfaces (see page 69)
About Merging Interfaces (see page 68)

Create Custom Virtual Interfaces

You can create custom virtual interfaces (CVIs) to report on subnet traffic. Create CVIs for a network that is designed in the following way:

- Traffic from a data center is transferred to subnets through a multiprotocol label switching (MPLS) carrier cloud, and
- Flow is enabled on the routers in the data center rather than on the routers on the edge of the cloud.

Without virtual interfaces, you have limited visibility into which subnets actually generate the network traffic for the carrier cloud. Virtual interfaces help ensure that you can collect detailed data about the subnet traffic in this type of configuration.

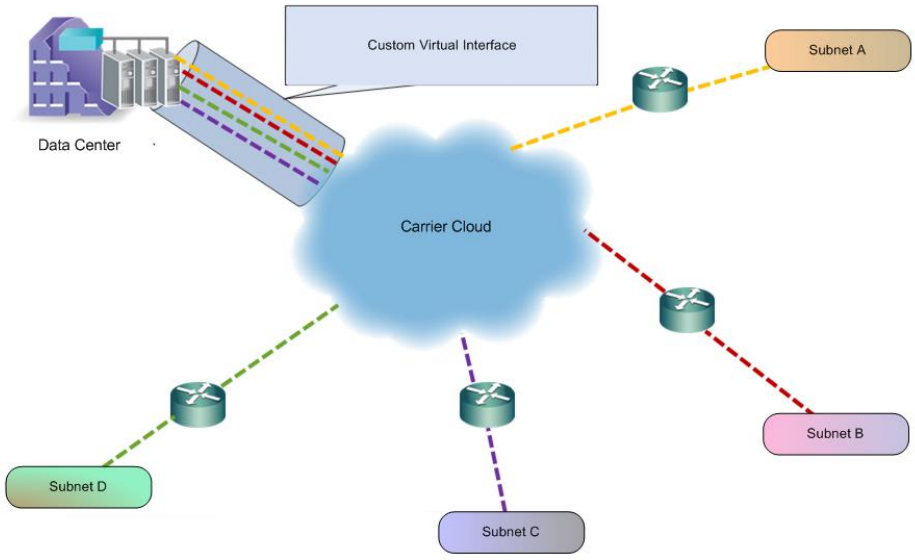

Defining a custom virtual interface (CVI) separates traffic for an interface from other traffic. The CVI in the diagram separates specific traffic from other traffic traveling to and from the data center into the cloud.

How to Configure Custom Virtual Interfaces

Create custom virtual interfaces (CVIs) to separate traffic on a particular interface and subnet from other traffic on the interface.

Follow these steps:

1. Open the Active Interfaces page:
   a. Select Administration from the NFA console menu.

      The Administration page opens.

   b. Select Interfaces: Physical & Virtual from the Administration menu.

      The Active Interfaces page opens.
2. Expand the interface list for the parent router by clicking the arrow next to the router name.
3. Select the check box next to the name of a single interface.
4. Click Add Custom Virtual Interface.

   The Add Custom Virtual Interface dialog opens.
5. Enter values for the following fields:
   - Interface Name: Replace the default value in the Interface Name field with a meaningful name for the interface list.

     Default: Parent interface name, which you must replace.
   - Description: (Optional) Enter a text string to help identify the interface.
   - In Speed and Out Speed: (Optional) Identify the speed of data that is inbound to the parent interface and outbound from the parent interface.
   - Type: (Optional) Select the interface type from the list.
   - Domain: Select a tenant-domain combination from the list or accept the default setting.

     Changing the domain can affect which operators and reports have access to the data. The Domain option is visible only in an environment that contains multiple domains.
   - Subnet: Enter a subnet and mask identification for each subnet filter you want to use for this CVI, then click Add. Use the following format: <subnet IP address/subnet mask>

     The CVI must contain at least one subnet filter.

6. Click Save.

   The Custom Virtual Interface is automatically deployed within one minute. The Class icon for the new CVI distinguishes it from the physical interfaces.

Important! If you delete a parent CVI, all the children CVIs are deleted automatically.

Custom Virtual Interface Priorities

CVI priorities are set according to how narrowly you specify the subnet mask. Subnet masks that are most specific have the highest priority. A CVI that has a single-node subnet mask (such as 192.168.20.2/32) has higher priority than a CVI that has a multiple-node subnet mask (such as 192.0.0.0/8). Establishing priorities in this way helps ensure that CVI traffic remains separate from other traffic, especially for a host on a large subnet.

Note: Reports for traffic between two CVIs with the same priority can be inconsistent.

Example of Custom Virtual Interface Definitions

If you define the following CVIs:

- CVI-A: 192.168.0.0/16
- CVI-B: 192.168.100.0/24
- CVI-C: 192.168.100.123/32

The following reporting is enabled:

- CVI-A reports traffic that involves subnet 192.168.*x*.*x* except traffic that involves subnet 192.168.100.*x*.
- CVI-B reports traffic that involves subnet 192.168.100.*x* except traffic that involves host 192.168.100.123.
- CVI-C reports traffic that involves host 192.168.100.123.

If you add 192.168.100.124/32 as CVI-D, either CVI-C or CVI-D reports the traffic between 192.168.100.123 and 192.168.100.124. Only one of the CVIs reports the traffic.

If you add 192.168.200.0/24 as CVI-E, either CVI-B or CVI-E reports the traffic between 192.168.100.*x* and 192.168.200.*x*. Only one of the CVIs reports the traffic.

Merge Interfaces

You can merge interfaces so they are shown as one interface in reports. You may want to do this whenever major logical changes to a device cause a new interface to be created in the system.

Also See:

Edit Router and Interface Details (see page 60)
Enable or Disable Interfaces (see page 75)
Delete Interfaces (see page 63)

About Merging Interfaces

Consider the following guidelines when you merge interfaces:

- You can merge two interfaces, which can be on different routers.
- The data collection period for the interfaces can contain gaps.
- If you merge interfaces that have overlapping time frames, the overlapping data is discarded. Precedence is given to the newer interface--that is, the interface that has the later start date.

  For example, suppose you merge Interface A and B. Interface A has collected data that starts at 1:00 P.M. and ends at 5:00 P.M. on the same day. The data collected for Interface B starts at 3:00 P.M. The merged data consists of Interface A data from 1:00 to 3:00 P.M. and Interface B data from 3:00 P.M. onward.
- You cannot merge interfaces that started collecting data at the same time.

Example

For example, suppose that your 512-Kbps link that has been running for a year is upgraded to a T1 link (a 1.54-Mbps link). Using the new T1 link causes the interface ifindex to change. For example, the previous ifIndex of 5 changes to 13--the next available ifindex for the T1 link. Other settings are changed or created, such as the ifdescr and ifAlias settings.

These changes cause the program to see the interface as a new interface. You enable the new interface so the program can collect its data.

At this point, the history of the interface is divided. To unify the history, you merge the two versions of the interface. After the merge, the history includes the data that was collected previously from the interface on the slower link and the data from the interface on the new link. The data is combined end-to-end with no overlaps or duplication.

Steps for Merging Interfaces

Follow these steps:

1. Open the Active Interfaces page:
   a. Select Administration from the NFA console menu.

      The Administration page opens.

   b. Select Interfaces: Physical & Virtual from the Administration menu.

      The Active Interfaces page opens.
2. Expand the interface list for a router by clicking the arrow next to the router.

   The view expands to show a list of the interfaces for the selected router.
3. Select the check box next to two interfaces that you want to merge.

   The Merge button at the top of the page is activated and no longer appears dimmed.

   Note: The Merge button is enabled only when you have selected two interfaces.
4. Click Merge.

   The Merge Interface Confirmation dialog opens.
5. Verify that the information in the confirmation dialog is correct:
   - Verify that the interface shown as the source interface (as defined in the fields under the 'Copy data from' label) is the interface from which you want to copy data.
   - Verify that the interface shown as the destination interface (as defined in the fields under the 'Copy data to' label) is the interface to which you want to copy data.
   - Select the 'Delete source interface after merging data' check box to delete the source interface automatically after its data is copied.
6. Click Save.

   The selected interfaces are merged according to the values in the Merge Interface Confirmation dialog box.

Customize the Page

Customize the appearance of the Active Interfaces page as needed to make it easier to find information. You can sort tables, display details for routers, and change the maximum number of interfaces and CVIs that appear in detail pages under each router.

Sort Table Data

To sort the data, click a menu field. For example, to view all the routers with a red status indicator click Flow Status. The results are sorted by the Flow Status timestamp.

Expand the Details for Routers

To expand the details for one of the routers, click the arrow on the left side of the router row. A list of interfaces is shown below the router line.

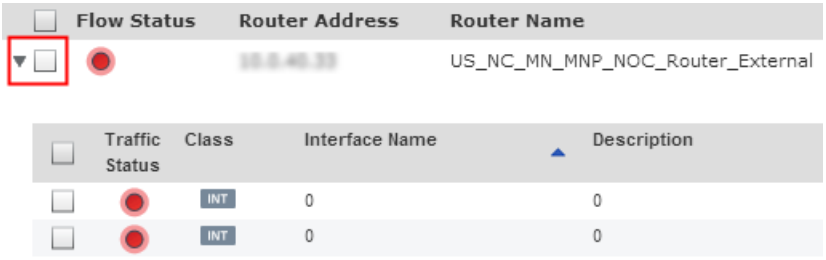

Note: The number of details that are visible in the expanded view is limited to the Max per Page value for the selected router.

To display or hide details for all the routers, interfaces, or CVIs, click Expand All or Collapse All in the upper-right corner.

Change the Number of Visible Details

To change the number of items that are displayed, select a different maximum value from the Max per Page drop-down list. You can change the Max per Page value for the router list. You also can set a Max per Page value for the sublist of interfaces and CVIs under each router.

Available Interfaces Page

The Available Interfaces page displays information about routers and interfaces. The list of routers and interfaces include the interfaces that have never been enabled and have never been the source of collected data. You can enable or disable interfaces, delete end-of-life routers and their interfaces, and perform some polling troubleshooting.

Router Tasks:

- Review the displayed data, including the total number of interfaces, number of enabled interfaces, and the SNMP profile that was used.
- Perform a test poll with the current SNMP profile (Test).
- Look for an SNMP profile (Discover).
- Refresh the polling and interface information for the poller (Refresh).
- Enable or disable all of the interfaces for a router.
- Delete an end-of-life router and its interfaces from the system.

Interface Tasks:

- Review the displayed data, including the most recent time that flow was received and the enabled or disabled status.
- Enable or disable interfaces individually.

Follow these steps:

1. Open the Available Interfaces page:
   a. Select Administration from the NFA console menu.

      The Administration page opens.

   b. Select System: Enable Interfaces in the Administration menu.

      The Available Interfaces page opens.
2. Use the information and options in the following ways, for example:
   - Enable interfaces that you want to begin using.
   - Disable interfaces that you are not using.
   - Delete routers that you will not use again.
   - Review the information that is displayed as described in Router Information (see page 72) and Interface Information (see page 74).

Available Interfaces: Router Information

The Available Interfaces page includes the following options and information for routers.

Test 

Attempts to poll the interfaces for the router by using the settings that are displayed. You can use this option in troubleshooting to test SNMP connectivity. If the SNMP Profile value is missing, the test always fails.

- Success: If the test succeeds, click Refresh.
- Failure: If the test fails, you can click the Discover option to try to find an SNMP profile that works. Alternatively, you can assign a different SNMP profile on the Active Interfaces page and click Test again.

Discover 

Looks for an SNMP profile to use for polling the router's interfaces. You can use the Discover option if the SNMP Profile value is missing and you do not know which profile to use.

- Success: If discovery succeeds click Refresh, then click Test. If the profile is different from the current one, the SNMP Profile value is updated.
- Failure: Any SNMP Profile value that was displayed before discovery is removed. Discovery can fail for several reasons. For example, a valid SNMP profile may not be available, access to the profile may be blocked, or the router may be offline.

Refresh 

Sends updated polling and interface information to the poller. Click the Refresh option after a successful Discover or Test operation.

Enable and Disable

Control whether the interface (or router and its interfaces) is allowed to send flow to CA Network Flow Analysis, as described in Enable or Disable Interfaces (see page 75).

Delete

Removes the router and its interfaces from the system, as described in Delete Routers from the System (see page 76).

Standard Page Options

- Search: Search for routers or interfaces by address or name (see page 59).
- Max per Page: Change the number of items that are displayed (see page 70).

- Expand or collapse router contents (see page 70).
- Sort contents by column (see page 70).

Router Columns

Flow Status 

Status indicator to show whether the most recent regular polling attempt was successful:

- Red: Any enabled interfaces have not had flow for longer than the **Interface Data Absence Limit**.
- Yellow: Any enabled interfaces have not had flow between 30 minutes and the **Interface Data Absence Limit**.
- Green: All enabled interfaces have had flow in the last 30 minutes.

Router Address

IP address of the router.

Router Name

Name of the router.

SNMP Profile

Name of the SNMP profile that was used in the most recent successful polling attempt. If this profile is not successful in the next polling attempt, the poller uses the next available profile until polling succeeds. In this case, the SNMP Profile value is updated. If polling fails for all of the profiles, the SNMP Profile value is blank.

Total Interfaces

Number of interfaces for the router.

Enabled Interfaces

Number of enabled interfaces for the router.

Harvester

IP address of the parent Harvester for the router.

Available Interfaces: Interface Information

The Available Interfaces page includes the following options and information for interfaces.

Enable and Disable

Control whether the interface (or the router and its interfaces) is allowed to send flow to the product, as described in Enable or Disable Interfaces (see page 75).

Standard Page Options: The standard page options as described in Available Interfaces: Router Information (see page 72).

Interface Columns

The ifName, ifAlias, Port Name, and vrfName values are present only if the parent router is configured to provide this information to the poller. This information comes from the interfaces database. If you make changes to the corresponding properties on the Active Interfaces page, your changes are not shown here.

Enabled

Setting to let the interface send flow to the product (Yes) or not send flow (No).

License

Status that shows whether the interface has sent flow to the product (Yes) or has never sent flow (No). If the value is 'No,' the system has no records for the interface.

ifIndex

Identifying index value that is automatically assigned to the interface

ifName

Interface name

ifAlias

Interface alias

Port Name

Interface port name

vrfName

Virtual routing and forwarding name

Speed

Maximum data transmission speed for the interface

Last Flow

Most recent date and time that flow was processed. If flow is being collected currently, the Last Flow value is updated every 15 minutes.

Enable or Disable Interfaces

Interfaces are enabled automatically on the Application Settings page by default. You may want to enable interfaces manually in some cases, such as in the following example scenarios:

- Interfaces have been disabled temporarily and you want to re-enable them.
- The program is not configured to enable newly discovered interfaces automatically. (The Auto-Enable Interfaces option is set to False in the Application Settings page.)

Enabling an interface causes the following events to occur immediately:

- The interface can be used as a filter in Flow Forensics reports.

  The data that is available in Flow Forensics reports includes data that was collected while the interface was disabled.
- Other types of interface data begin to be collected and stored, such as data for hosts, conversations, ToS, and protocols. Once the initial period of data collection is complete, the additional interface data can be displayed in drilldown interface reports, Analysis reports, and Custom reports.

Follow these steps:

1. Open the Available Interfaces page:
   a. Select Administration from the NFA console menu.

      The Administration page opens.

   b. Select System: Enable Interfaces in the Administration menu.

      The Available Interfaces page opens.
2. Click the arrow next to the router that contains the interfaces.

   The view expands to show the interface list. The Enabled status column shows which interfaces are enabled.
3. Select the check box next to one or more interfaces.
4. Click one of the following options:
   - Enable (enable data collection for the interface)
   - Disable (prevent the collection of data for the interface)

   Data collection for the selected interfaces is enabled or disabled immediately.

Delete Routers from the System

If you delete a router on the Available Interfaces page, the router is deleted from the system entirely. The deletion removes the router, its configuration information, 15-minute (historical) data, interfaces, CVIs, and traps. The deletion also affects any aggregations, views, and reports that previously included the interfaces.

Note: If the interfaces from the deleted router begin to send flow again, a new router appears on the Available Interfaces page. If the program is configured to enable new interfaces automatically, the new router and interfaces also appear on the Active Interfaces page. The new router inherits the current tenant-domain setting of its parent Harvester. Configuration settings for the previous router are lost. If your deployment includes CA Performance Center, the router polls by using the SNMP profiles that are assigned to the Harvester tenant.

Follow these steps:

1. Verify that the router is no longer sending flows to CA Network Flow Analysis.
2. Open the Available Interfaces page:
   a. Select Administration from the NFA console menu.

      The Administration page opens.

   b. Select System: Enable Interfaces in the Administration menu.

      The Available Interfaces page opens.
3. Locate the router and select its check box.
4. Click Delete.

   A confirmation message opens.
5. Click Yes.

   The following events result:

   - The confirmation message closes.
   - The router is deleted from the system, the Available Interfaces page, and the Active Interfaces page.
   - The router configuration information is deleted from the NFA console server.
   - All related interfaces, CVIs, 15-minute (historical) data, and traps are deleted.
   - The router is deleted from any related aggregations.
   - The data for the deleted interfaces no longer appears in the NFA console or in reports.

Define Interface Name Templates

To change the rules that determine how interface names and descriptions appear in the NFA console, use any or all of the following procedures:

- Create an interface template and make it the active template for CA Network Flow Analysis.
- Edit the current interface template.
- Refine the interface naming convention for specific NFA console views by using the Application Settings page.

You can create multiple interface templates, but only the currently selected interface template determines the names and descriptions of interfaces in the NFA console, its views, and its printed reports.

(CA PC only) The currently selected interface template also affects interface names and descriptions in some CA Performance Center views of CA Network Flow Analysis data. The following locations in CA Performance Center may display different interface names and descriptions: Interface Pages (Details tab), Inventory pages, and Trend views. To customize interface descriptions in CA Performance Center, apply Interface Description Overrides to specific domains.

Create and Apply a Custom Interface Template

Create a custom interface template to change the way all interface names and descriptions appear in the NFA console.

Follow these steps:

1. Log in as a user who has administrator privileges for CA Network Flow Analysis.
2. Display the Interface Templates page in the NFA console:
   a. Select Administration from the NFA console menu.

      The Administration page opens.

   b. Select System: Templates from the Administration page menu.

      The Interface Templates page opens.
3. Click Add.

   The Interface Templates page displays options for adding an interface template.

4. Specify the interface template settings:
   - Name: Identifying text string that appears in the template list.
   - Interface Name: Properties, text, or properties and text that comprise the name of each interface in the NFA console.
   - Interface Description: Properties, text, or properties and text that appear as the description of each interface in the NFA console.

   Use properties from the following list for both settings:
   - *[DeviceAlias]*: Device (router) name that is displayed in the NFA console
   - *[DeviceName]*: DNS name or IP address of the device (router)
   - *[ifDescr]*: Interface description, which is taken from the original ifDescr value in the SNMP ifEntry table unless the interface description has been customized. If the interface description has been customized on the Active Interfaces page, the customized value is used.
   - *[ifAlias]*: Interface alias
   - *[ifName]*: Name of the interface, which is taken from the original ifName value in the SNMP ifEntry table value unless the interface name has been customized. If the interface name has been customized on the Active Interfaces page, the customized value is used.
   - *[portName]*: Port name, which may be the port number
   - *[ifIndex]:* Unique numerical identifier for the interface as defined in the SNMP ifEntry table
   - *[ifType]:* Interface type as defined in the ifType field of the SNMP ifEntry table
5. Click Submit.

   The new template is created and is added to the template list. The extra options are removed from the Interface Templates page.
6. (Optional) Apply the template: Select the template from the list at the top of the Interface Templates page.

   The template is applied to the NFA console almost immediately. The Performance Center views that use the template reflect the changes at the next synchronization, which occurs within 5 minutes.
7. (Optional) Review the results of changing the template. For example, review the new labels on Interface pages and on the Enterprise Overview page.

Conventions for Interface Templates

Interface template settings consist of properties, plain text, or properties and plain text. The following conventions apply to interface templates:

- Enclose properties in square brackets.
- Separate multiple properties with a pipe symbol. The first property that returns a value is displayed. Specify enough properties to allow for any interfaces that have missing property definitions.
- To display plain text in the interface names or descriptions, include the text without enclosing it in brackets.

Edit an Interface Template

Edit an interface template to change the way all interface names and descriptions appear in the NFA console when the template is selected.

Follow these steps:

1. Display the Interface Templates page in the NFA console:
   a. Select Administration from the NFA console menu.

      The Administration page opens.

   b. Select System: Templates from the Administration page menu.

      The Interface Templates page opens.
2. Select the template from the list at the top of the page.
3. Click Edit.

   A confirmation message opens.
4. Click OK.

   The contents of the Interface Name and Interface Description fields become editable.
5. Specify the interface template settings:
   - Interface Name: Text and/or properties that comprise the names of interfaces in the NFA console.
   - Interface Description: Text and/or properties that appear as the description of interfaces in the NFA console.

   Define both settings with properties from the list in the topic about creating a template (see page 77):

6. Click Update.

   The contents of the fields are updated and are no longer editable.

   The interface names and definitions are displayed in the NFA console almost immediately. The CA Performance Center views that use the template reflect the changes at the next synchronization, which occurs automatically within 5 minutes.

7. (Optional) Review the results of changing the template. For example, review the new labels on Interface pages and on the Enterprise Overview page.

Change the Application Setting for Interface Names

You can use a setting on the Application Settings page to change the naming convention for interfaces. The default setting adds the device name in front of the interface name. This setting affects the way interface names appear in some NFA console report views, such as the Enterprise Overview views, Interface pages, and Custom Report Interface Summaries.

The interface names that result can include unwanted duplications. If a device and an interface are both named Device1, for example, by default the interface name is shown as Device1::Device1. The following example illustrates this repetition.

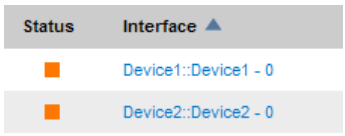

To eliminate the duplication, edit the Show Device Name setting on the Application Settings page.

Follow these steps:

1. Display the Application Settings page:

   a. Select Administration from the NFA console menu.

      The Administration page opens.

   b. Select System: Application Settings from the Administration page menu.

      The Application Settings page opens.

2. Set the Show Device Name value to False.

3. Click Save.

   With the Show Device Name value set to False, device names in reports and views are not prepended to interface names.

Chapter 5: Working with Interface Groups and Aggregations

Options for creating and managing interface groups become available in Performance Center as soon as you register the product as a data source (see page 15). Custom interface groups can help users get the best results from custom reports and analyses. For example, interface groups can help operators set up custom reports that are based on geographic location, interface speed, T1 sites, or load balancing.

You can create and manage interface aggregations in the NFA console. An interface aggregation combines traffic from two or more interfaces so the traffic is reported together. For example, suppose you have two load-balanced circuits and you want to report on their interfaces as a single unit. Aggregations let you report on interfaces as a single unit without setting up a custom report for that purpose. All aggregations appear in the Interface Index under the label *Aggregations*.

Note: You can aggregate interfaces only when their data is collected by the same Harvester. Do not attempt to aggregate interfaces from multiple Harvesters.

Create Interface Aggregations

Create interface aggregations to combine traffic from two or more interfaces so the traffic is reported together.

For example, suppose that you create aggregations for each of several geographical regions. Operators use the aggregations in reports to review and compare the totals for specific regions. Before you created the aggregations, operators had the following problems:

- Data was scattered throughout reports.
- Values were not totaled by region, so quick overviews and comparisons were difficult to make.
- Operators spent too much time designing specialized reports to collect the data and create the totals they needed.

Follow these steps:

1. Display the Interface Aggregations page:
   a. Select Administration from the NFA console menu.

      The Administration page opens.

   b. Select Interfaces: Aggregations from the Administration page menu.

      The Interface Aggregations page opens and shows the list of current interface aggregations.
2. Click New.

   The page changes to show options for adding an aggregation.

   Select interfaces to add to the aggregation by using any of the following methods:

   - *Select routers:* Select the check boxes next to one or more router names. You can select one router, multiple routers, or a combination of routers and individual interfaces from any routers that are not selected. When a router has been selected, you cannot select any of its individual interfaces.
   - *Select interfaces:* Click Expand All or click the arrow icon next to a router name, then select the interfaces from the list. To select all the interfaces for a router, select the check box in the heading row.
   - *Filter the list to find routers or interfaces:* Filter your view by entering a text string in the Search box, then clicking Search. You can search for a full or partial name or address of a router or interface. The router/interface list is filtered to display matching routers or routers that contain matching interfaces.
   - *Filter the list to check your selections:* Click Show Selected to show only the routers that are selected or that contain selected interfaces. You can combine this function with a search string and with Expand All to locate and review your selections quickly.

3. Specify the following values, then click Save:
   - Routers and/or Interfaces: Select one or more of the routers or interfaces that are listed at the bottom of the page. Your selections are added to the aggregation.
   - Aggregation Name: Specify a name for the aggregation.
   - Description: (Optional) Add a notation to help identify the aggregation.
   - In Speed: (Optional) Specify the inbound speed of the selected interfaces.
   - Out Speed: (Optional) Specify the outbound speed of the selected interfaces.

     If you do not specify the In speed and Out speed, both values are set to 0 by default. Inaccurate speeds cause some report results to be inaccurate, such as utilization percentages.
   - Type: (Optional) Select the mode of interface connection from the Type list, such as WAN or Ethernet.

     If you do not specify the interface type, the type is set to Unknown by default.

   The aggregation is automatically deployed within one minute of creation.

Edit Interface Aggregations

Edit a single interface aggregation to choose different interfaces to aggregate or to change the aggregation name, description, speed, or interface type.

Follow these steps:

1. Display the Aggregations Interfaces page:
   a. Select Administration from the NFA console menu.

      The Administration page opens.

   b. Select Interfaces: Aggregations from the Administration page menu.

      The Interface Aggregations page opens and shows a list of the current interface aggregations.
2. Click the check box next to the aggregation that you want to edit, then click Edit.

   The page changes to the Edit Aggregation view, which includes editable options.
3. Make any changes that are needed, such as adding or removing interfaces or changing the aggregation name, description, in/out speed, or type.
4. Click Save.

   Your changes are saved and you return to the list of aggregations.

Delete Interface Aggregations

Delete interface aggregations that are no longer needed.

Follow these steps:

1. Display the Interface Aggregations page:
   a. Select Administration from the NFA console menu.

      The Administration page opens.

   b. Select Interfaces: Aggregations from the Administration page menu.

      The Interface Aggregations page opens and shows the current list of aggregations.
2. Select the check box next to each aggregation that you want to delete.
3. Click Delete.
4. Click Yes in the confirmation message box that opens.

   All of the selected aggregations are deleted automatically.

Chapter 6: Working with Harvesters and DSAs

This section describes tasks for working with Harvesters and with DSAs (in a three-tier deployment).

Use the Harvester page to perform the following tasks:

- Add and delete Harvesters (see page 85)
- Edit Harvester details (see page 86)

If you have a three-tier deployment, use the DSA page to perform the following tasks:

- Add and delete DSAs (see page 30)
- Edit the IP address of a current DSA (see page 88) or a new DSA (see page 89)

Add and Delete Harvesters

Add and delete Harvesters on the Harvester page.

To add a Harvester, complete the steps in the topic Add a Harvester (see page 26).

To delete a Harvester, complete the steps in this topic.

Note: If you delete a Harvester you cannot add the same Harvester instance again successfully unless the installation server has been re-imaged and the Harvester software has been re-installed. Once you delete a Harvester, you cannot recover any of the data that the Harvester collected previously.

Follow these steps:

1. Open the Harvester page:
   a. Select Administration from the NFA console menu.

      The Administration page opens.

   b. Select System: Harvester from the Administration page menu.

      The Harvester page opens and displays a table of information about the current Harvesters.
2. Click Delete in the row for the Harvester that you want to delete.

   A confirmation message opens.

3. Click Yes to confirm that you want to delete the Harvester.

   The Harvester is deleted and is not listed in the Harvester table. The NFA console no longer collects data from the routers that are associated with the deleted Harvester. The data that was collected from those routers previously is not available in reports.

Edit Harvester Details

Edit the details for Harvesters on the Harvester page. You can edit the IP address, description, and tenant-domain setting.

Follow these steps:

1. Select Administration from the NFA console menu.

   The Administration page opens.

2. Select System: Harvester from the Administration page menu.

   The Harvester page opens and displays a table of information about known Harvesters.

3. Click Edit on the Harvester row that you want to edit.

   The Edit Harvester dialog opens.

4. (Optional) Change any of the following settings:

   IP Address

   Address of the Harvester

   Description

   (Optional) Additional information to help identify the Harvester

   Domain

   The tenant-domain association of the Harvester in a multi-domain deployment.

   Changing this setting affects the tenant-domain for any new routers that begin exporting flow data. Existing routers and interfaces retain their previous tenant-domain associations.

   The domain affects which operators and reports have access to the data from routers and interfaces.

   Changing the tenant of a router in CA Performance Center can affect which SNMP profiles are available for polling. This is not applicable to CA NetQoS Performance Center, which uses the same list of SNMP profiles for all routers.

   Default: Default Tenant\Default Domain

   If no custom IP domains have been created, the Harvester table includes only the IP Address and Description columns.

5. Click Save.

   Your changes are saved immediately and appear in the Harvester table.

Edit DSA IP Addresses

Applies to: Three-tier architecture in a distributed deployment: NFA console, Harvester, and DSA components installed on separate servers

Edit the IP Address setting for a Data Storage Appliance (DSA) when you:

- Change the IP address for a current DSA--a DSA that is already in use and that is not moving to a different server
- Start using a DSA on a new server in place of a retired DSA

  Do not add a new DSA instance instead of editing the IP address of the retired DSA. In this case, routers continue to send data to the retired DSA--and that data is not available in reports.

Notes: These tasks apply only to:

- Three-tier distributed deployments of CA Network Flow Analysis. If you have a stand-alone system or two-tier architecture, you do not define or update DSA IP addresses.
- Windows servers. CA Network Flow Analysis supports DSAs on Windows servers only.
- Servers that have static IP addresses and that are not configured to use Dynamic Host Configuration Protocol (DHCP).

Edit the IP Address of a Currently Connected DSA

Edit the DSA IP Address setting when the IP address changes for a currently connected DSA.

Perform this task when the IP address changes on a DSA server. If you are moving the DSA to a new server, skip this task and go to Edit the IP Address for a New DSA (see page 89).

Follow these steps:

1. Log in with an account that has administrator privileges for CA Network Flow Analysis.
2. Stop the current DSA from collecting data:
   a. Open the Windows Services window on the DSA server.
   b. Select NetQoS Reporter/Analyzer Pump Service in the Services list.
   c. Click the Stop link on the left.

      The service stops running.
3. Let the DSA finish processing all of the .rpr files in the following directory: <install_path>\Netflow\datafiles\loaderInput.

   Completing this step lets the DSA retrieve the raw data that was generated while the service was stopped. Processing finishes within 15 minutes. Processing is complete when no .rpr files are left in the loaderInput directory.
4. Change the IP address on the DSA server itself. For more information about this step, refer to the recommendations at microsoft.com.
5. Update the firewall settings if necessary.
6. Display the DSA page in the NFA console user interface:
   a. Select Administration from the NFA console menu.

      The Administration page opens.
   b. Select System: DSA from the Administration menu.

      The DSA page opens and shows a list of the current DSAs.
7. Click Edit on the row for the DSA that you want to edit.

   The Edit DSA dialog opens.
8. Edit the IP Address value so that it matches the IP address of the new DSA server.
9. Click Test Connection.

   If the connection succeeds, the Test Success message appears.
10. Click OK in the Test Success message box.

    The message box closes.

11. Click Save in the Edit DSA dialog.

    The new IP address for the DSA is saved. CA Network Flow Analysis starts using the new IP address when it stores and retrieves some data.

12. Restart the NetQoS Reporter/Analyzer Pump Service in the Windows Services window on the DSA server.

    The DSA starts collecting data again. The DSA should be able to retrieve all the data that was generated while the service was stopped.

Edit the IP Address for a New DSA

Edit the DSA IP Address setting to start using a new DSA in place of a retired DSA.

Perform this task when you move the DSA to a new server. For example, perform this task when a DSA server fails or you move the DSA to a new server for improved performance or storage space. Once you edit the IP address to match the new server, the switch to the new DSA is complete.

Notes:

- If you move a DSA to a new server, the data for the previous DSA is no longer available in reports.
- If you add a DSA instance instead of editing the IP address of a retiring DSA, the retired DSA still appears in the DSA list. Routers apportion data between all the DSAs in the list, so the routers continue to send some data to the retired DSA. Reports do not show the data that is sent to nonfunctional DSAs. To correct this problem, contact CA Support for assistance in deleting the obsolete DSA.
- If the DSA stayed on the same server, skip this task and go to Edit the IP Address of a Currently Connected DSA (see page 88).

Follow these steps:

1. Install the DSA software on its new server and finish configuration--including firewall updates. For more information about this step, see the *CA Network Flow Analysis Installation Guide.*
2. Log in to the NFA console with an account that has administrator privileges for CA Network Flow Analysis.
3. Run and save copies of any reports that you want to use for archiving information from the previous DSA.

4. Display the DSA page in the NFA console user interface:
   a. Select Administration from the NFA console menu.

      The Administration page opens.

   b. Select System: DSA from the Administration menu.

      The DSA page opens and shows a list of the current DSAs.
5. Click Edit on the row for the DSA that you are replacing.

   The Edit DSA dialog opens.
6. Edit the IP Address value so that it matches the IP address of the new DSA server.
7. Click Test Connection.

   If the connection succeeds, the Test Success message appears.
8. Click OK in the Test Success message box.

   The message box closes.
9. Click Save in the Edit DSA dialog.

   The new IP address for the DSA is saved. The new DSA starts collecting, processing, and storing data. The data from the retired DSA is no longer available in reports.

Chapter 7: Creating Names and Groups for Protocols, ToS, and AS Data

This section describes how to create and manage protocol groups, ToS labels, ToS groups, and customized autonomous system (AS) names. You use a number of different pages to perform the functions.

Protocol Group Configuration Page

You can use the Protocol Group Configuration page to view the existing protocol groups and their contents. You can add, edit, and delete custom protocol groups.

Select Groups: Protocol Groups from the Administration page, then use the Protocol Group Configuration page to perform the following main tasks:

- Create a shell protocol group (see page 92).
- Configure the protocol group (see page 93).
- Review and edit the group (see page 94).

ToS Configuration Page

You can use the ToS Configuration page to view and edit the labels and descriptions for ToS values.

Select Define an Application: ToS Names from the Administration page, then use the ToS Configuration page to perform the following task:

- Label ToS values (see page 95).

ToS Group Configuration Page

You can use the ToS Group Configuration page to view the existing ToS groups and their contents. You can add, edit, and delete custom ToS groups.

Select Groups: ToS Groups from the Administration page, then use the ToS Group Configuration page to perform the following main tasks:

- Create a shell ToS group (see page 96).
- Configure the ToS group (see page 97).
- Edit ToS groups (see page 98).
- Delete ToS groups (see page 99).

AS Names Page

You can use the AS Names page to view and edit AS names.

Select Groups: AS Names from the Administration page, then use the AS Names page to perform the following main tasks:

- Review AS names (see page 100).
- Edit AS names (see page 101).

Create Protocol Groups

Protocol groups can be used to filter data in Custom and Analysis reports. As an administrator, you can set up custom protocol groups that contain the protocols for particular types of network traffic.

For example, suppose that operators want to report on network traffic for different types of applications, such as email, videoconferencing, VoIP, and streaming media. No default protocol groups are defined for this purpose, so you create a custom protocol group for each of the application types. Each protocol group includes all the protocol values that are used in your enterprise for the target applications.

Operators can use a custom protocol group without having a thorough knowledge of the way each protocol is used and without adding individual protocol filters to report definitions.

The following topics show CA Network Flow Analysis administrators how to create a shell protocol group, configure the protocol group, then review and edit the protocol group settings.

Create a Shell Protocol Group

Create a shell protocol group that you can configure to filter reports for a particular type of network traffic.

Follow these steps:

1. Log in as a user who has administrator privileges for CA Network Flow Analysis.
2. Display the Protocol Group Configuration page:
   a. Select Administration from the NFA console menu.

      The Administration page opens.

   b. Select Groups: Protocol Groups from the Administration page menu.

      The Protocol Group Configuration page opens and displays information about the currently selected protocol group.

3. Click Add.

   The options for identifying a new protocol group are displayed on the Protocol Group Configuration page.

4. Enter values in the Group Name and Description boxes.

   The name and description appear in the following locations in CA Network Flow Analysis:

   - List of protocol groups on the Protocol Group Configuration page
   - Protocol Group Index that an operator can display from the report wizard for defining an Analysis or Custom report

   A new shell protocol group is created. The new protocol group name is added to the list of protocol groups on the Protocol Group Configuration page.

   Next: Configure the shell protocol group.

Configure the Protocol Group

After you create a shell protocol group, the next step is to configure it to represent a particular type of network traffic.

Follow these steps:

1. Select the protocol group from the 'Select a Protocol Group' list on the Protocol Group Configuration page.
2. Click List.

   Configuration options are added to the Protocol Group Configuration page.

3. (Optional) Review and correct the Domain setting if necessary.

   The Domain setting is displayed only in an environment that has multiple domains.

   The contents of the selected protocol group are displayed with the protocol names for the selected tenant-domain combination. If an Administrator defines domain-specific protocol names, selecting that domain displays the appropriate protocol names.

   The Domain setting does not restrict access to the protocol group or to reports that use the protocol group as a filter.

4. Select the protocols for the group:

   a. Click Add/Remove.

      A dialog opens, which shows a list of available protocols and a list of any protocols that are included in the list currently.

      If the Add/Remove link is not visible, click List.

b. Select the protocols in the top pane that you want to add to the group.

   To select multiple protocols to add simultaneously, use the Shift and Control keys.

   To filter the protocol list, enter a search string in the Filter the Protocol List field box, then click Apply. For example, to show only UDP protocols, enter **udp**.

c. Click Add.

   The selected protocols are added to the protocol list.

d. Select any protocols in the bottom pane that you want to remove from the group.

e. Click Remove.

   The selected protocols are removed from the protocol list.

f. Click Done when you finish configuring the contents of the protocol group.

   To review a list of protocol groups that contain a particular protocol, select the protocol and click 'Jump to Protocol.' The Protocol Configuration page opens and displays all of the protocol lists that contain the selected protocol.

The protocol group is configured with your settings. Operators can select the configured protocol group as a filter for an Analysis or Custom report.

Note: Operators can define reports only if their user settings are set up to enable this capability.

Review and Revise the Protocol Group Settings

After you create a protocol group and operators use it in reports, you may want to make changes. For example, you may want to rename the protocol group to reflect the way operators use it. You also may want to change the list of protocols in the group.

Follow these steps:

1. Select the protocol group from the ‘Select a Protocol Group’ list on the Protocol Group Configuration page.
2. (Optional) Edit the name or description of the protocol group:
   a. Click Edit.

      The Edit a Protocol Group page opens.

   b. Change the values for the Group Name and Description fields to help identify the purpose of the protocol group more clearly.

   c. Click Submit.

      The name and description are updated in the Protocol Group Configuration page list and in the Protocol Group Index.

3. (Optional) Review and revise the list of protocols that are included:

   a. Click List.

      A table is added to the bottom of the page, which lists the protocols.

   b. Select the appropriate domain from the Domain list, if custom protocol names have been defined.

      The Domain setting is displayed only in an environment that has multiple domains.

      For more information about the Domain setting, see step 3 in Configure the Protocol Group (see page 93).

   c. Click Add/Remove and change the contents of the protocol group as needed

      For more information about this step, see step 4 in Configure the Protocol Group (see page 93).

   d. Click Done when you finish configuring the contents of the protocol group.

   The protocol group is configured with your settings. Operators can select the reconfigured protocol group as a filter for an Analysis or Custom report.

Label ToS Values

You can create labels (or descriptions) for ToS values to ensure that operators know which service or application the ToS values represent. If you do not label ToS values, the numeric ToS value is displayed by default.

Note: The ToS labels that you create affect all uses of the ToS in the assigned domain, but do not affect ToS labels in other domains. ToS labels are domain-specific in a deployment that has multiple domains.

Follow these steps:

1. Open the ToS Configuration page:

   a. Select Administration from the NFA console menu.

      The Administration page opens.

   b. Select ToS Names from the Administration page menu.

      The ToS Configuration page opens and displays information about the selected ToS and domain.

2. (Multiple-domain environment) Select the domain that contains the ToS values that you want to edit.

3. Select a ToS value from the list, then click Edit.

   The Description becomes editable.

4. Enter the new description for the ToS value in the Description box and click Save.

   Your change is saved and the Description and Value fields for the selected ToS are updated.

Note: You can also use the ToS Configuration page to add or delete ToS values to ToS groups. Select the domain and the ToS value, then click the Add/Remove link at the top of the ToS groups list. The ToS Group Index dialog opens, which you use to make changes for user-configured groups. (You cannot change the built-in All ToS group.)

Create and Manage ToS Groups

An administrator can create ToS groups that act as filters for report data. The administrator sets up each ToS group to contain the ToS values that characterize a particular type of network traffic. Operators can use the ToS group as a filter in reports instead of adding filters for each ToS value.

For example, suppose that operators want to report on network traffic for the Low Drop Assured Forwarding group. The Low Drop Assured Forwarding group is a group of ToS values that are assigned to applications that should have priority over less critical traffic. The administrator creates a ToS group named Low Drop, which includes all the ToS values for AF11, AF21, AF31, and AF41. Operators can create a report about the application traffic by using the new Low Drop ToS group.

Note: In an environment that includes domains, ToS groups are available in all domains as filters for Analysis reports and Custom reports. The ToS labels that administrators see displayed for the contents of a ToS group are domain-specific, however.

Create a Shell ToS Group

Create ToS groups to help users get optimal results from Analysis and Custom reports quickly.

Follow these steps:

1. Display the ToS Group Configuration page:
   a. Log in as a user who has administrator privileges for CA Network Flow Analysis.
   b. Select Administration from the NFA console menu.

      The Administration page opens.

   c. Select Groups: ToS Groups from the Administration page menu.

      The ToS Group Configuration page opens and displays information specific to the selected ToS group and domain (if applicable).

2. Click Add.

   The ToS Group Configuration page displays options for identifying a new shell ToS group.

3. Enter values in the Name and Description boxes.

   The name and description appear in the following locations:

   - List of ToS groups on the ToS Group Configuration page
   - ToS Group Index that a user can display from the report wizard for defining an Analysis report or Custom report

4. Click Add.

   A new shell ToS group is created. The new ToS group name appears in the list of ToS groups on the ToS Group Configuration page.

   You now can add ToS values to the shell ToS group.

Add ToS Values to the ToS Group

After you create a shell ToS group, the next step is to add values to the ToS group.

Follow these steps:

1. Select the ToS group from the 'Select a ToS group' list on the ToS Group Configuration page.

   - (Optional) Select a tenant-domain combination from the Domain list, provided that you have multiple domains and domain-specific ToS labels.

   The Domain setting is displayed only in a multi-domain environment.

   By selecting a tenant-domain combination, you display any available domain-specific ToS labels, which help you identify the ToS values to use. The ToS group will not be restricted to the selected domain: ToS groups are available in all domains.

2. Click Add/Remove.

   The ToS Index dialog opens. The ToS list displays the ToS index values and the ToS labels that have been defined. In a multi-domain environment, the ToS labels are displayed for the currently selected domain.

3. Select the ToS values to add by clicking their check boxes.

4. Scroll to the bottom of the list and click Save.

   The ToS values are added to the ToS group.

   An operator can select the ToS group from the ToS Group Index when defining an Analysis or Custom report.

   Note: To define reports, the operator must have the necessary role and permission group settings.

Change the Contents of ToS Groups

You can change the list of ToS values included in a ToS group to make corrections or to change the purpose of the group.

Even in a multi-domain environment, ToS groups are independent of domains. The ToS values included in a ToS group do not change when you switch to a different domain.

Follow these steps:

1. Select Administration from the NFA console menu.

   The Administration page opens.

2. Select Groups: ToS Groups from the Administration page menu.

   The ToS Group Configuration page opens and displays information for the currently selected ToS group and tenant-domain combination (if applicable).

3. Select the ToS group to edit from the ‘Select a ToS group’ list.

   The ToS Group Configuration page opens and displays information for the currently selected ToS group and domain (if applicable).

   Note: You cannot edit built-in groups, such as the All ToS group in the Default Domain.

4. (Multi-domain environment only) Select a tenant-domain combination from the Domain list, provided that you have multiple domains and ToS labels. In a multi-domain environment, ToS labels are domain-specific.

   The Domain setting is displayed only if multiple domains exist.

5. Click the Add/Remove link above the list of ToS values that are included in the group.

   The ToS Index dialog opens. The ToS list displays the ToS labels that have been defined.

6. Select or clear check boxes to include only the ToS values you want.

7. Click Save when your changes are complete. The list of ToS values is updated.

Delete ToS Groups

You may want to delete a ToS group that is no longer useful. Once you delete a ToS group, it is not listed in any ToS group lists.

Follow these steps:

1. Select Administration from the NFA console menu.

   The Administration page opens.

2. Select Groups: ToS Groups from the Administration page menu.

   The ToS Group Configuration page opens and displays information for the currently selected ToS group and domain (if applicable).

3. Select the ToS group to delete from the 'Select a ToS group' list.

   Select a configurable ToS group. For example, you cannot delete the All ToS group.

   Note: The Domain setting, if any, has no effect on the list of available ToS groups. If you delete a ToS group, it is deleted for all domains.

4. Click Delete.

   A confirmation message opens.

5. Click OK.

   The selected ToS group is deleted. The list of ToS values is updated.

Customize AS Names

Interface reports that show data about Autonomous System (AS) traffic typically label the AS traffic by name and number. Administrators can customize AS names to make the AS references in reports shorter or more descriptive.

For example, suppose that operators frequently view reports about network traffic for AS 4000000. The label with the default AS name is "UUNET-CANADA - Progressive Communications Services, Inc. d/b/a Acme Business (4000000)." The administrator can customize the AS name so that the report label is "Acme (4000000)."

In a multi-domain environment, AS names are domain-specific. In this type of environment, the AS names that appear in reports are taken from the domain of the interface for the report.

The topics in this section describe how to review and edit the current AS names.

To complete these tasks, make sure that you have configured your NetFlow or NetFlow-compliant flow to support AS reporting. To view meaningful AS data in reports, you must enable it: NetFlow does not export full AS information by default. For information about enabling AS data, refer to the Knowledge Base article TEC562036, "Viewing AS Numbers in Reports," at CA Support (http://ca.com/support).

AS data is shown as AS 0 in reports when any of the following conditions apply:

- Flow is not configured to support AS reporting.
- The data source route is unknown.
- The data originates from within the local system.

Review the Autonomous System Names

Review the AS references in the interface reports for your enterprise and locate any AS numbers in use that have long or unclear AS references.

The interface reports that contain AS references are:

- Single interface AS Next Hop summary table
- Top N AS summary tables, trend charts, and pie charts

Follow these steps:

1. Collect information about which AS names are used in interface reports for your enterprise.

   The list of all AS numbers and names is extensive. You will not want to customize every AS name in the list.

   For example, compile a list of commonly used AS names by examining AS reports or by asking operators which AS numbers they track.

2. Display the AS Names page:

   a. Log in as a user who has administrator privileges for CA Network Flow Analysis.

   b. Select Administration from the NFA console menu.

      The Administration page opens.

   c. Select Define an Application: AS Names from the Administration page menu.

      The AS Names page opens and displays information about the first page of AS values.

3. Locate each AS with a name that you want to customize:
   - To change the number of rows on each page, use the Max per Page option.
   - To view a different page, use the navigation elements at the bottom of the page.
   - To search, enter one or more text strings in the Search box, then click Search.

     Rules for matches:
     - Searches are case-insensitive.
     - All of the specified search strings must be found in either the AS Number or Description column.
     - Enclose the search string in double quotation marks if you want order and completeness to be limiting factors. For example, matching results for *internet back* without quotation marks include *internet global backbone* and *Backbone Internet Service*.
4. Review the AS names that appear in the Description column.

   Once you locate a default AS name that requires customization, you are ready to edit the AS name.

Edit Autonomous System Names

Customize AS names as needed to make interface report labels more user-friendly.

Follow these steps:

1. (Optional) Review and correct the Domain setting if necessary.

   In a multi-domain environment, AS names are domain-specific. In this case, changes to AS names affect reports about interfaces in the selected domain.
2. Select the row for the AS number that you want to edit.
3. Click Edit.

   The Edit AS Number Description dialog opens.
4. Enter the custom name (description) in the New Name box.

   The Old Name value is the official (base) name for the AS, which cannot be edited.

5. Click Save.

   The name that you specified appears in the Description column for the selected AS number. All custom AS names are shown in bold.

   The updated AS names also appear in labels and other references in the following CA Network Flow Analysis Interface views:

   - Single interface AS Next Hop summary table
   - Top N AS summary tables, trend charts, and pie charts

6. Repeat these steps for each AS that you want to customize.

Note: To restore the AS name to its official value, click Reset on the appropriate row, then OK in the confirmation box that opens.

Chapter 8: Making Additional Customizations

You can make additional customizations to ensure that operators get the best results from reports.

Once the product is registered as a data source, you use the Performance Center Console to administrate users, roles, SNMP profiles, and many types of groups.

This section describes how to perform the following customization tasks in the NFA console:

- Create time filters (see page 104)
- Create reporting periods (see page 106)
- Set up Application Mapping (see page 107)
- Work with Reserved Seating (see page 126)
- Set Up Flow Cloning (see page 130)

Also See:

Working with Interfaces and Routers (see page 55)
Working with Harvesters and DSAs (see page 85)
Creating Names and Groups for Protocols, ToS, and AS Data (see page 91)

Create Time Filters

Create time filters to help users create reports that contain streamlined data. For example, suppose that your users need some reports that describe network traffic during business hours. You create a time filter for Monday through Friday from 8:00 A.M. to 5:00 P.M. You may also want to create time filters for the time frames of specific operations that occur in your environment, such as automated backups.

You can also use the time filters that you create to configure traps on the Trap Configuration page.

Note: For more information about using time filters, see the topic Create Traps (see page 145).

The custom time filters are available to users in Time Filter option lists in the following locations:

- Specify Schedule page of the Custom Report wizard or Analysis wizard.
- Options that operators display by clicking the time period setting in a drilldown interface report of any of the following types: Overview, Protocols, ToS, Hosts, Conversations, Flows, Utilization, and AS Numbers.

Follow these steps:

1. Select Administration from the NFA console menu.

   The Administration page opens.

2. Select Time Filters under the Administration: Reporting label in the Administration menu.

   The Time Filter Configuration page opens and displays a list of the currently configured time filters.

3. Click Add.

   The Time Filter Configuration page displays options for adding a time filter.

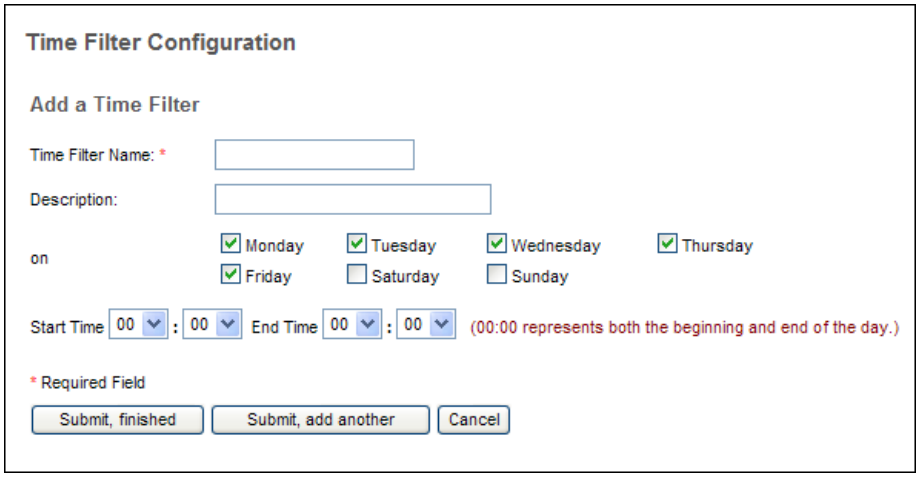

4. Enter values for the following options:

   - Time Filter Name: Define the time filter identifier, which appears in the following lists:
     - Time Filter Configuration page list (for administrators)
     - Time Filter options that are available on the Specify Schedule page of the Custom Report wizard
     - Time Filter options that are available when a user or administrator clicks the time period setting in a drilldown interface report of any of the following types: Overview, Protocols, ToS, Hosts, Conversations, Flows, Utilization, and AS Numbers
   - Description: (Optional) Add information to identify the time filter, which appears in the list of time filters on the Time Filter Configuration page.
   - on: Accept the default settings (Monday through Friday) or select other days for collecting report data.
   - Start Time and End Time: Accept the default settings for the start and end of the daily time span or select custom settings, using the 24-hour clock system. The default setting of 00:00 and 00:00 includes data for all 24 hours of the day.

     For example, to restrict the reporting period to business hours, select 08:00 as the start time and 17:00 as the end time. To set up a filter for backups that run from 11 P.M. Tuesday to 3:00 A.M. Wednesday, select Tuesday, then select 23:00 and 03:00.

5. Save the time filter by clicking one of the following buttons:
   - Submit, finished: Save the time filter and return to the time filter list.
   - Submit, add another: Save the time filter and keep the Add options open so you can configure another time filter.

Note: To delete a time filter, select it from the list of available time filters, click Delete, then confirm the deletion when prompted. To edit a time filter, select it from the list of available filters, change any of the options that are displayed, then click Submit.

Create Reporting Periods

Create reporting periods to get the best results from reports. For example, suppose that your users typically analyze network traffic that occurs over two weeks. To make it easier to display interface data over this timespan, you create a two-week reporting period.

Custom reporting periods are available as Time Period options when a user clicks the time period setting in a drilldown interface report of any of the following types: Overview, Protocols, ToS, Hosts, Conversations, Flows, Utilization, and AS Numbers. For more information about these report types, see the *CA Network Flow Analysis User Guide.*

Follow these steps:

1. Select Administration from the NFA console menu.

   The Administration page opens.

2. Select Reporting Periods under the Administration: Reporting label in the Administration menu.

   The Reporting Periods Configuration page opens and displays a list of the current reporting periods, including the built-in and customized reporting periods.

3. Click Add.

   Fields and options for adding a reporting period are displayed.

4. Specify the following values:
   - Reporting Period: Identifier for the reporting period, which appears in the following lists of available reporting periods:
     - Reporting Periods Configuration page list (visible to administrators)
     - Time Period options available when a user or Administrator clicks the time period setting in a drilldown interface report of any of the following types: Overview, Protocols, ToS, Hosts, Conversations, Flows, Utilization, and AS Numbers.

- Duration: Number of time units and type of time units--years, months, weeks, days, or hours.
- Description: (Optional) Additional notes to help identify the reporting period. The description appears in the Reporting Periods Configuration page list and is visible only to administrators.

5. Save the reporting period by clicking one of the following buttons:
   - Submit, finished: Save the current reporting period and return to the list of available reporting periods.
   - Submit, add another: Save the current reporting period and keep the Add options open so you can continue to configure additional reporting periods.

Note: To delete a reporting period, select it from the list, click Delete, then confirm the deletion when prompted. To edit a reporting period, select it from the list, change any of the options that are displayed, then click Submit.

Set Up Application Mapping

Create application mapping rules to identify traffic in reports. Rules can identify traffic types such as a ToS, host, subnet, or NBAR2 application.

You can use application mapping to combine, differentiate, or more clearly identify traffic in reports:

- Differentiate Traffic: Separate larger traffic blocks by re-mapping traffic sub-types to separate destination ports.

  For example, suppose reports show a large block of FTP traffic on TCP port 20. You want to track the FTP traffic from your internal FTP server separately from internet traffic. To accomplish this, you create a Host application mapping rule, which you name *Internal FTP Traffic.* The Host value of the rule matches the IP address of the internal FTP server. The Port value is 20. You specify 65000 as the Destination Port--a port that does not currently receive any traffic.

  Reports now show traffic from the FTP server on TCP port 65000 with the label *Internal FTP Traffic.* Other TCP port 20 traffic is still labeled *FTP.*

- Combine Traffic: Report traffic of different types as a single unit by re-mapping them to a single destination port.

  For example, suppose your enterprise mail systems use the IMAP and POP protocols. The IMAP mail uses TCP port 443 and the POP mail uses TCP ports 109 and 100. You want reports to show the combined mail traffic, so you create application mapping rules that re-map each type of mail traffic to port 3100. The traffic is combined in reports and is labeled with the rule name, *Mail*. Even though you created several rules, the program uses the same name for all of the rules that map traffic to port 3100.

- Identify Traffic: Use application mapping to re-label traffic without combining or separating it.

Application mapping affects the following reports:

- Enterprise Overview page: Top Protocols report
- Interface page: All reports that show protocol data
- Custom Reporting page and Analysis page: Protocol Index, which you use to select protocols to filter a new or edited report

Notes:

- Application mapping does not affect Flow Forensics reports.

  To continue the example for differentiating types of FTP traffic, the Flow Forensics Session Protocols reports show the FTP traffic as it was before you mapped the FTP sub-category.

  The Flow Forensics reports that display NBAR2 data show the official application name and ID regardless of any application mapping rules that you have.
- Reports show mapping results within the time frame that the rules are in effect. If you create rules today, reports of last week's data do not show the effects of the rules. If you create, delete, or edit application mapping rules within a report time frame, the report may show a dramatic change at the time you made the rule changes.

You can perform the following application mapping tasks:

- Use Administration functions in the NFA console to set up or modify application mapping rules to aggregate or segregate data:
  - Create an All (ToS) application mapping rule (see page 110)
  - Create a Host application mapping rule (see page 112)
  - Create a Subnet application mapping rule (see page 113)
  - Create an NBAR2 application mapping rule (see page 115)
  - Edit application mapping rules (see page 116)
- Configure global settings to support application mapping rules in reports (see page 109)
- Understand how application mapping rule priorities work (see page 109)
- Perform batch Import operations by using a .csv file:
  - Import the default NBAR2 application mapping rules (see page 118)
  - Import custom application mapping rules (see page 119)
  - Import application mapping rule updates (see page 122)
  - Review errors for failed rule imports (see page 124)

Application Mapping Priorities

If incoming flow matches the combination of criteria that are specified in an application mapping rule, the flow is mapped to the rule destination port. If rules conflict, the following priorities apply. A rule takes precedence when its criteria are ranked higher than the criteria of another rule.

- *Priority 1:* Host rule with Host, Protocol, Port, and ToS specified
- *Priority 2:* Host rule with Host, Protocol, and Port specified, but with the ToS set to ALL
- *Priority 3:* All (ToS) rule with ToS specified
- *Priority 4:* Host rule with the Host specified, but with the Protocol and ToS set to ALL
- *Priority 5:* Subnet rule
- *Priority 6:* NBAR2 rule

Configure Global Settings for Application Mapping

Review the global settings that affect application mappings. Make any adjustments that are needed to customize application mapping behavior. The global settings that affect application mapping include the following options:

- *TCP Rebase Port:* Target port for redirected TCP traffic. If an application mapping rule sends TCP traffic to a port that already has native (non-mapped) traffic, the native traffic is redirected to the TCP rebase port.
- *ToS Mask:* Setting to determine whether all ToS bits are used for ToS values.
- *UDP Rebase Port:* Target port for redirected UDP traffic. The UDP rebase port receives redirected native UDP traffic in the way the TCP rebase port receives redirected native TCP traffic.
- *Preserve ToS Map Proto:* Setting to determine whether protocol values are reported in addition to ToS values.

Follow these steps:

1. Open the Application Definitions page:
   a. Select Administration from the NFA console menu.

      The Administration page opens.

   b. Select Application Definitions in the Administration menu.

      The Application Definitions page opens.
2. Click Additional Settings.

   The Application Settings page opens.

3. Review and correct the global settings that affect application mapping:

   - *TCP Rebase Port:* Target port for redirecting native TCP traffic, which is used if an application mapping rule sends TCP traffic to a port that already receives native (non-mapped) traffic.

     For example, suppose that you create an application mapping rule that has port 655 as the destination port. You set the TCP Rebase Port value to 630. If port 655 is receiving native (non-mapped) traffic, the native traffic is redirected to port 630.

     Note: The rebase port is used to avoid merging mapped traffic with other types of traffic. If application mapping rules are configured with unused destination ports, the rebase port is not used: Reports do not show traffic on the rebase port. If your reports show traffic on the rebase port, you may want to specify a new destination port for the mapped traffic. Examine the rebase port traffic to determine which rules are involved.

   - *ToS Mask:* Limits the 8-bit ToS values reported for flow. The default value for this setting is 255, which enables all ToS bits.

   - *UDP Rebase Port:* Redirects native UDP traffic in the same way the TCP Rebase Port setting redirects native TCP traffic.

   - *Preserve ToS Map Proto:* Setting to determine whether protocol values are reported in addition to ToS values. Y(es) retains the protocols that are used and reports ToS data separately. N(o) shows you the ToS values for traffic, but you do not see which protocols are involved. The default setting is Y.

4. Click Save.

   Your changes are saved. The Application Settings page remains open.

Create an All (ToS) Application Mapping Rule

Create an All application mapping rule to combine, separate, or more clearly identify traffic by its ToS (Type of Service) value.

Follow these steps:

1. Open the Application Definitions page:

   a. Select Administration from the NFA console menu.

      The Administration page opens.

   b. Select Application Definitions in the Administration menu.

   The Application Definitions page opens.

2. Verify that Application Mapping is the selected value for Rules.
3. Click Add Rule.

   The Add Application Mapping dialog opens.

4. Select All from the list of rule types at the top of the dialog.

   The Add Application Mapping dialog switches to All (ToS) rule mode.

5. Specify the setting values:

   - ToS: Type of Service (ToS) to use as a filter for the collected data.

     ALL is the ToS value that is shown initially when you open the dialog. If left unchanged, this setting would map traffic for all ToS values to the destination port.

   - Destination Port: Target port that collects the mapped data

     If you specify a destination port that is already used by other rules, the traffic for the related rules will be combined.

   - (Optional) Click Check to run a general check to detect whether the specified port is already receiving data. The check fails if the port is receiving native data--that is, data that is not mapped by application mapping rules. If any native data is on the specified destination port, that data is redirected to the rebase port (see page 109).

   - Name: Identifier for the rule as it is listed on the Application Definitions page

     The rule name also is the label for the mapped traffic in certain reports. If other rules map traffic to the same destination port that you specify for this rule, specify the name that you want to use for the combined traffic.

   - Description: (Optional) Additional descriptive text to identify the rule type and its use, which is displayed only on the Application Definitions page

6. Click Save.

   The dialog closes. The new rule is added to the Application Mapping rule list. If any other rules map traffic to the same port and you specified a new rule name, the other rule names are updated.

7. (Optional) Run reports to verify that the traffic on the designated destination port fits the rule.

8. (Optional) Review the effects of the new or changed application mapping rules on reports, then consider renaming the rule to label the mapped traffic more clearly in reports.

Create a Host Application Mapping Rule

Create a Host application mapping rule to combine, separate, or identify traffic based on its source host. For example, a Host application mapping rule can report the total traffic from a particular server or from an application on the server.

Follow these steps:

1. Open the Application Definitions page:
   a. Select Administration from the NFA console menu.

      The Administration page opens.

   b. Select Application Definitions in the Administration menu.

   The Application Definitions page opens.
2. Verify that Application Mapping is the selected value for Rules.
3. Click Add Rule.

   The Add Application Mapping dialog opens.
4. Select Host from the list of rule types at the top of the dialog.

   The Add Application Mapping dialog switches to Host rule mode.
5. Specify values for the following settings:
   - Host: IP address of the server whose network traffic you are mapping
   - ToS: Type of Service (ToS) to use as a filter for the collected data. To match all ToS values, accept the default value of ALL or leave the ToS box blank.
   - Protocol: Protocol of the data that is affected by the rule, either TCP or UDP
   - Port: Port to use for collecting data. To match all ports, accept the default value of ALL or leave the Port box blank.
   - Destination Port: Target port that collects the mapped data

     If you specify a destination port that is already used by other rules, the traffic for the related rules will be combined.
   - (Optional) Click Check to run a general check to detect whether the specified port is already receiving data. The check fails if the port is receiving native data--that is, data that is not mapped by application mapping rules. If any native data is on the specified destination port, that data is redirected to the rebase port (see page 109).

- Name: Identifier for the rule as it is listed on the Application Definitions page

  The rule name also is the label for the mapped traffic in certain reports. If other rules map traffic to the same destination port that you specify for this rule, specify the name that you want to use for the combined traffic.

- Description: (Optional) Additional descriptive text to identify the rule type and its use, which is displayed only on the Application Definitions page

6. Click Save.

   The dialog closes. The new rule is added to the Application Mapping rule list. If any other rules map traffic to the same port and you specified a new rule name, the other rule names are updated.

7. (Optional) Run reports to verify that the traffic on the designated destination port fits the rule.

8. (Optional) Review the effects of the new or changed application mapping rules on reports, then consider renaming the rule to label the mapped traffic more clearly in reports.

Create a Subnet Application Mapping Rule

Create a Subnet application mapping rule to combine, separate, or more clearly identify traffic that originates from a particular subnet and mask. For example, a Subnet rule can enable reports to show the total traffic for an application.

Follow these steps:

1. Open the Application Definitions page:

   a. Select Administration from the NFA console menu.

      The Administration page opens.

   b. Select Application Definitions in the Administration menu.

      The Application Definitions page opens.

2. Verify that Application Mapping is the selected value for Rules.

3. Click Add Rule.

   The Add Application Mapping dialog opens.

4. Verify that Subnet is selected as the rule type at the top of the dialog. (Subnet is selected by default.)

   The Add Application Mapping dialog displays the options for a Subnet application mapping rule.

5. Specify values for the following settings:
   - Subnet: IP address of the data source, expressed in dotted decimal format. To specify a subnet that matches all addresses, use 0.0.0.0/0 as the subnet and mask.
   - Mask: Mask to apply to the subnet.
   - Protocol: Protocol of the data that is affected by the rule, either TCP or UDP
   - Start Port: Beginning of the port range for collected data, expressed in Base 10 decimal format. The Start Port is included in the port range. The maximum port value that is allowed is 65535.
   - End Port: Last port in the range to use for collecting data. The End Port is included in the port range.
   - Destination Port: Target port that collects the mapped data

     If you specify a destination port that is already used by other rules, the traffic for the related rules will be combined.
   - (Optional) Click Check to run a general check to detect whether the specified port is already receiving data. The check fails if the port is receiving native data--that is, data that is not mapped by application mapping rules. If any native data is on the specified destination port, that data is redirected to the rebase port (see page 109).
   - Name: Identifier for the rule as it is listed on the Application Definitions page

     The rule name also is the label for the mapped traffic in certain reports. If other rules map traffic to the same destination port that you specify for this rule, specify the name that you want to use for the combined traffic.
   - Description: (Optional) Additional descriptive text to identify the rule type and its use, which is displayed only on the Application Definitions page
6. Click Save.

   The dialog closes. The new rule is added to the Application Mapping rule list. If any other rules map traffic to the same port and you specified a new rule name, the other rule names are updated.
7. (Optional) Run reports to verify that the traffic on the designated destination port fits the rule.
8. (Optional) Review the effects of the new or changed application mapping rules on reports, then consider renaming the rule to label the mapped traffic more clearly in reports.

Create an NBAR2 Application Mapping Rule

Create NBAR2 (Next Generation Network-Based Application Recognition) application mapping rules to identify NBAR2 application traffic in reports (see page 107). NBAR2 rules can identify traffic for individual applications, combine traffic for multiple applications, or separate NBAR2 traffic from other traffic.

If multiple rules map traffic to the same destination port, the program gives the rules the same name--the name you specified most recently. The rule name is the label for the NBAR2 traffic in reports.

This topic describes how to create NBAR2 application mapping rules individually on the Applications Definitions page. You also can batch import NBAR2 application rules by using the command line. For more information about command-line import options, see the topics beginning with Import Application Mapping Rules (see page 117).

Notes:

- To display NBAR2 data in reports, your routers must be configured to return IPFIX flows that include the appropriate NBAR2 fields
- Application mapping supports rules for the applications that are identified by the standard Cisco NBAR2 engine (NBAR2 engine 13), but not for applications that are identified by custom NBAR2 engines.

Follow these steps:

1. Open the Application Definitions page:
   a. Select Administration from the NFA console menu.
      The Administration page opens.
   b. Select Application Definitions in the Administration menu.
   The Application Definitions page opens.
2. Verify that Application Mapping is the selected value for Rules.
3. Click Add Rule.
   The Add Application Mapping dialog opens.
4. Select NBAR2 from the list of rule types.
   The dialog switches to NBAR2 rule mode.
5. Specify values for the following settings:
   - NBAR2 Application ID: The application ID that is defined by the standard NBAR2 engine.
     Specify the application ID correctly or the rule does not function as expected. The application IDs are included in the nbar2.csv file in the following folder: <install_path>/reporter/racmd.

- Destination Port: Target port that collects the mapped data

  If you specify a destination port that is already used by other rules, the traffic for the related rules will be combined.

  (Optional) Click Check to run a general check to detect whether the specified port is already receiving data. The check fails if the port is receiving native data--that is, data that is not mapped by application mapping rules. If any native data is on the specified destination port, that data is redirected to the rebase port (see page 109).

- Name: Identifier for the rule as it is listed on the Application Definitions page

  The rule name also is the label for the mapped traffic in certain reports. If other rules map traffic to the same destination port that you specify for this rule, specify the name that you want to use for the combined traffic.

- Description: (Optional) Additional descriptive text to identify the rule type and its use, which is displayed only on the Application Definitions page

Note: The NBAR2 Engine ID value is pre-populated and cannot be edited. The value is 13, the standard NBAR2 engine.

6. Click Save.

   The dialog closes. The new rule is added to the Application Mapping rule list. If any other rules map traffic to the same port and you specified a new rule name, the other rule names are updated.

7. (Optional) Run reports to verify that the traffic on the designated destination port fits the rule.

8. (Optional) Review the effects of the new or changed application mapping rules on reports, then consider renaming the rule to label the mapped traffic more clearly in reports.

Edit Application Mapping Rules

Edit an Application Mapping rule when you want to change the specified source port or port range, destination port, protocol, host, ToS, subnet, mask, rule name, rule description, or rule type.

Follow these steps:

1. Open the Application Definitions page:

   a. Select Administration from the NFA console menu.

      The Administration page opens.

   b. Select Application Definitions in the Administration menu.

      The Application Definitions page opens.

2. Click the check box next to the rule.

3. Click Edit.

   The Edit Application Mapping dialog opens.

4. Make the needed changes.

5. Click Save.

   If the changes are successful, the dialog closes. The list of rules is updated to reflect your changes. If multiple rules map traffic to the same destination port and you change one of the rule names, the other names are updated to match. This name is used to label the combined traffic in reports.

Note: You can delete one or more rules by selecting their check boxes and clicking Delete.

Import Application Mapping Rules

You can create or update application mapping rules by importing a .csv file that is formatted for the rule type. You can perform several types of application mapping rule imports:

- Import custom application mapping rules (see page 119)
- Import the default NBAR2 application mapping rules (see page 118)
- Import application mapping rule updates (see page 122)

Importing rules has the following effects and characteristics:

- After the import, the new or updated rules are shown in the Application Definitions page list.
- Reports show mapped traffic with labels that match the new rule names, as described in Set Up Application Mapping (see page 107).
- Perform this type of import operation locally--you cannot perform the import remotely.
- You can work on the application mapping rules in a Microsoft Excel spreadsheet, then export the rules to .csv format.
- If you import the default NBAR2 rule set, each NBAR2 application is mapped to a separate port by default. The default port range is port 65001 and above.

Note: For information about import failures and potential error messages, see Mapping Rule Import Errors (see page 124).

Import the Default NBAR2 Application Mapping Rules

You can add a complete set of default NBAR2 application mapping rules in a batch import operation. You perform this task on the command line by using the nbar2.csv file that comes with the product.

Note: The NBAR2 application mapping rules that you import represent the current set of NBAR2 application definitions as of the product release time frame.

Follow these steps:

1. Log in to the NFA console server or stand-alone server as a user who is a member of the Administrators group.
2. Open a command prompt.
3. Navigate to the directory that contains the nbar2.csv file: Enter the following command:
   `cd <install_path>\reporter\racmd`

   where:

   <install_path> is the product installation path. The default path is C:\CA\NFA.

   racmd is the directory that contains the nbar2.csv file. The file is written to this directory when you install the product.
4. Enter the following command:
   `racmd -import nbar2.csv`

   where:

   nbar2.csv is the name of the application mapping rule file that you want to import. The command string is based on the expectation that the racmd command and .csv file are in their default location. If you have moved the .csv file, include the fully qualified path (the path and file name).

   If any errors occur during the import, error messages are shown (see page 124). If no message is returned, the import succeeded with no problems.
5. (Optional) Verify that the rules are listed on the Applications Definitions page:
   a. Select Administration from the NFA console menu.

      The Administration page opens.

   b. Select Application Definitions in the Administration menu.

      The Application Definitions page opens. The new rules are shown in the Application Definitions page list.

6. (Optional) Verify that the mapped NBAR2 application traffic is labeled appropriately in the following locations:
   - Enterprise Overview page: Top Protocols report
   - Interface page: All reports that show protocol data
   - Custom Reporting page and Analysis page: Protocol Index, which you use to select protocols to filter a new or edited report
7. (Optional) Combine traffic reporting for selected applications if you like.
   a. Identify a set of applications that you want to see reported as combined traffic.
   b. Edit the corresponding application rules to send the data to a single destination port.

      By default, each NBAR2 application is mapped to a separate port. The default ports are above 65000.
   c. Use an appropriate name for the rules: Use a name that reflects the type of applications that are included.

      Reports use this name to label the combined traffic.

      If you update a rule name, all rules that use that destination port are renamed to match the new definition. The rule name and label is set when you edit the final rule name.

Import Custom Application Mapping Rules

You can create application mapping rules of a single type by importing a properly formatted .csv file. Example .csv files are provided, which show the required fields to include in the import file for each rule type.

Follow these steps:

1. Log in to the NFA console server or stand-alone server as a user who is a member of the Administrators group.
2. Open a command prompt.
3. Open the example import file for your rule type:
   - All (ToS) rule - tos.csv
   - Host rule with a specified protocol - server-protocol.csv
   - Host rule without a specified protocol - server.csv
   - Subnet rule - subnet.csv
   - NBAR2 rule - nbar2.csv

4. Follow the format in the example file for your rule type:
   - The first row in the file is the column name row, which identifies the fields. Leave the first row exactly as it is shown in the example file. Do not change the spelling or order of the column name row.
   - Add a row below the top row for each rule that you want to import.
   - Enter values for each required field.

     Separate the field values with commas. Do not include a comma inside any value string. A comma signals the import utility to go the next field.

     All fields are required except for the desc (Description) field. To specify a blank desc value, enter only a comma (with no space).
   - The import file columns correspond to the following columns on the Application Definitions page:
     - name = Name: Rule name
     - desc = (Optional) Description
     - protocolName - Protocol
       To specify all protocols, enter the value -1.
     - tos = ToS (Type of Service)
       To specify all ToS, enter the value -1.
     - ip = IP/Subnet: Host IP address
     - mask = IP/Subnet: Subnet addition to the host IP address
     - newPort = Destination Port
     - beginPort = Start Port: Starting port in a port range or port number for a server-protocol rule
     - endPort = End Port: Ending port in a port range
     - applicationID = NBAR2 Application ID

     Some fields apply only to particular rule types, as shown in the table at the end of these steps.
5. Go to the directory that contains the .csv file. The following command shows the default location:
   `cd <install_path>\reporter\racmd`

   where:

   <install_path> is the product installation path. The product installation path is C:\CA\NFA by default.

   racmd is the directory that contains the .csv import file. The file is written to this directory when you install the product.

6. Perform the import by entering the following command:
   `racmd -import nbar2.csv`

   where:

   nbar2.csv is the name of the application mapping rule file that you want to import. The command string is based on the expectation that the racmd command and .csv file are in their default location. If you have moved the .csv file, include the fully qualified path (the path and file name).

   If any errors occur during the import, error messages are shown (see page 124). If no message is returned, the import succeeded with no problems.

7. (Optional) Verify that the rules are listed on the Applications Definitions page:

   a. Select Administration from the NFA console menu.

      The Administration page opens.

   b. Select Application Definitions in the Administration menu.

      The Application Definitions page opens.

8. (Optional) Verify that the application traffic is labeled appropriately in the following locations:

   - Enterprise Overview page: Top Protocols report
   - Interface page: All reports that show protocol data
   - Custom Reporting page and Analysis page: Protocol Index, which you use to select protocols to filter a new or edited report

9. (Optional) Combine or separate traffic for selected applications if you like, as discussed in Set Up Application Mapping (see page 107).

The following table lists the fields that apply to each rule type. List the fields in the order shown in the import file example, not the order shown in the table.

Rule Type	name	desc	protocolName	tos	ip	mask	newPort	beginPort	endPort	applicationID
All - tos	Y	Y		Y			Y			
Host - server	Y	Y			Y		Y			
Host - server-protocol	Y	Y	Y	Y	Y		Y	Y		
NBAR2 - nbar2	Y	Y					Y			Y
Subnet - subnet	Y	Y	Y		Y	Y	Y	Y	Y	

Import Application Mapping Rule Updates

You can create application mapping rules of a single type by importing a properly formatted .csv file. Example .csv files are provided, which show the fields to include in each type of import file.

Follow these steps:

1. Log in to the NFA console server or stand-alone server as a user who is a member of the Administrators group.
2. Get the rule ID for the existing rules that you want to update.
   a. Open a command prompt and go to the directory that contains the .csv file. The following command shows the default location:
   `cd <install_path>\reporter\racmd`

   where:

   <install_path> is the product installation path. The product installation path is C:\CA\NFA by default.

   racmd is the directory that contains the .csv import file. The file is written to this directory when you install the product.

   b. Export the rule definitions by entering the following command:
   `racmd -export csv`

   The export file is named getapplicationmapping_<timestamp>.csv. The file is located in the current directory.

   The command returns the status message: Creating csv file. When the operation is complete, the command prompt reappears.

   Note: Rule IDs are specific to the current stand-alone or NFA console system.
3. Open the export file in a spreadsheet or text editor.

   The export file contains a line for each current application mapping rule, which begins with the rule ID. The line also includes extra information that you can ignore.
4. Locate and make a note of the ID for each rule that you want to edit.
5. Prepare the import file:
   a. Open a copy of the example import file for your rule type:
      - All (ToS) rule - tos.csv
      - Host rule with a specified protocol - server-protocol.csv
      - Host rule without a specified protocol - server.csv
      - Subnet rule - subnet.csv
      - NBAR2 rule - nbar2.csv

b. Add the appID column and rule values, as shown in the following examples:

Example: First two lines in an import file to add NBAR2 rules

```
name,desc,newPort,applicationid
youtube,Youtube video streaming,65035,82
```

where:

name = Rule name
desc = Rule description
newPort = Destination port
applicationID = NBAR2 application ID

Example: First two lines in an import file to update NBAR2 rules

```
appID,name,desc,newPort,applicationid
35,YouTube,Youtube video streaming,65035,82
```

where:

appID = Rule ID

If you update NBAR2 application mapping rules, do not change the NBAR2 application ID value. If you change this value, the rule will not function as expected.

c. Other than adding the appID column and values, follow the import mapping guidelines listed in step 4 of Import Custom Application Mapping Rules (see page 119).

d. Save the import file.

We recommend that you save the import file to the same directory that contains the import command: <install_path>\reporter\racmd.

6. Go to the directory that contains the .csv file:

```
cd <install_path>\reporter\racmd
```

where:

<install_path> is the product installation path. The product installation path is C:\CA\NFA by default.

racmd is the directory that contains the .csv import file. The file is written to this directory when you install the product.

7. Enter the following command:

```
racmd -import nbar2_updated.csv
```

where:

nbar2_updated.csv is the name of the update file that you just created. If you have moved the .csv file, include the fully qualified path (the path and file name).

If any errors occur during the import, error messages are shown (see page 124). If no message is returned, the import succeeded with no problems.

8. (Optional) Verify that the updated rules are listed correctly on the Applications Definitions page:
   a. Select Administration from the NFA console menu.

      The Administration page opens.

   b. Select Application Definitions in the Administration menu.

      The Application Definitions page opens.
9. (Optional) Verify that the application traffic is labeled correctly in the following locations:
   - Enterprise Overview page: Top Protocols report
   - Interface page: All reports that show protocol data
   - Custom Reporting page and Analysis page: Protocol Index, which you use to select protocols to filter a new or edited report

Error Messages for Importing Rules

This topic describes some of the error messages that may be displayed. When importing is complete, the command returns error messages for any failed imports, then lists the fields for the failed rules.

If the rule import file contains column definition errors, the import operation fails. If the file contains errors in rule definitions, the import operation skips each faulty rule and continues with the next rule.

Notes about the Description field:

- All of the rule fields must be populated with values except the optional Description (desc) field. The rule line must include the comma for the desc field, however.
- If Description field values contain internal commas, the rule import fails. The import utility does not support commas inside Description values.

A column named 'XXX' already belongs to this DataTable

The import file contains a duplicate of the column named in the error message. The column name is shown in place of 'XXX'.

Troubleshooting: Delete the duplicated column. Follow the format in the appropriate example file as described in Import Custom Application Mapping Rules (see page 119).

An invalid ApplicationMapping

One of the following problems occurred while you were updating existing rules:

- The import file specifies the appID of the original rule incorrectly.
- The specified appID corresponds to a rule that has been deleted.

Troubleshooting: Verify that the appID is entered correctly and that its rule still exists. You can export the current rules to a .csv file as described in Import Application Mapping Rule Updates (see page 122). The export file includes the appIDs for all of the current rules.

Application Mapping object is not valid: An existing record is already in use

Both of the following conditions were met:

- You are using the file format for importing a new rule.
- One of the rule definitions in the import file has the same field values as an existing rule.

Troubleshooting: If this error occurs rarely, you may want to edit the rules individually on the Application Definitions page. To update a group of existing rules by using the racmd command, use the import file format that is described in Import Application Mapping Rule Updates (see page 122).

Application Mapping object is not valid: An invalid protocol was entered

The import file may have a misspelled column name or may include an invalid column. This error may also occur if a rule has an invalid value for a required field (a field other than desc).

Troubleshooting: Verify that the following elements are correct in the import file:

- Field values: Rule definitions have supported field values. The values are entered correctly.
- Format: The correct columns are included in the file. The column names are spelled correctly.

Application Mapping object is not valid: IP cannot be blank

The error may occur if either of the following conditions is met:

- An NBAR2 rule is missing the NBAR2 application ID value.
- A Host (server) rule is missing the server IP address value.

Troubleshooting: Add the missing value to the rule definition.

Application Mapping object is not valid: Mask must be between 0 and 32

The Subnet rule has an unsupported Mask value.

Troubleshooting: Specify a Mask value between 0 and 32.

Application Mapping object is not valid: ToS must be between 0 and 255

The All (ToS) rule has an unsupported ToS value.

Troubleshooting: Specify a ToS value between 0 (all ToS) and 255.

Error: Missing parameter name from -params list

The rule has no specified value for one of the required fields.

Troubleshooting: Verify that all of the required field values are included. All field values are required except for desc.

To display help for the command, enter:
`setapplicationmapping`

Work with Reserved Seating

You can create Reserved Seating rules to help ensure that reports include the port and protocol combinations that interest you, regardless of traffic volume or rates. The rules create 'reserved seats' for the ports that are used by those protocols so the data is sure to be included in reports.

For example, during an application rollout you want to watch the traffic for a particular application, but the Top N Protocols reports for interfaces do not show the traffic for the application. The protocol that the application uses is not included in the Top N Protocol group--the group of protocols with the highest traffic volume or utilization rate. You create a Reserved Seating rule to collect data for the specific protocol and port that the application uses. The protocol now is included in the Top N Protocols reports.

Note: The data that CA Network Flow Analysis typically collects is described in the topic How Data Is Collected (see page 165).

Create Reserved Seating Rules

Create Reserved Seating rules to help ensure that reports include the port and protocol combinations that interest you, regardless of traffic volume or rates.

Follow these steps:

1. Open the Application Definitions page:
   a. Select Administration from the NFA console menu.

      The Administration page opens.

   b. Select Application Definitions in the Administration menu.

      The Application Definitions page opens.

2. Select Reserved Seating from the Rules list.

   The Application Definitions page switches to Reserved Seating mode and displays a list of the current Reserved Seating rules.

3. Click Add Rule.

   The Add Reserved Seating dialog opens.

4. Specify ports as follows:

   - Protocol: Protocol of the data that is affected by the rule, either TCP or UDP
   - Port: Target port for the Reserved Seating rule. Enter the port number in the Port box, a value from 0 through 65535, expressed in Base 10 decimal format. If you do not enter a value, port 0 is assigned. The port and protocol combination must be unique--that is, it cannot match any other Reserved Seating rule.

     Data of the specified protocol type is reported for this port regardless of traffic rate or volume.
   - Description: (Optional) Identifying text for the Reserved Seating rule. The description appears in the list of Reserved Seating rules on the Application Definitions page.

5. Click Save.

   If you entered a valid port and protocol combination and you have not yet reached the maximum number of rules, the dialog closes. The new rule appears in the list of Reserved Seating rules.

6. Repeat this process for each Reserved Seating rule you want to add.

   You can specify a maximum of 50 Reserved Seating rules.

Edit Reserved Seating Rules

Edit a Reserved Seating rule to make corrections in the specified port, protocol, or description.

Follow these steps:

1. Open the Application Definitions page:

   a. Select Administration from the NFA console menu.

      The Administration page opens.

   b. Select Application Definitions in the Administration menu.

      The Application Definitions page opens.

2. Select Reserved Seating from the Rules list.

   The Application Definitions page switches to Reserved Seating mode and displays a list of the current Reserved Seating rules.

3. Click the check box next to the rule you want to edit, then click Edit.

   The Edit Reserved Seating dialog opens.

4. Make the needed changes in the Protocol, Port, and Description values, then click Save.

   If the changes are successful, the Edit Reserved Seating dialog closes.

Delete Reserved Seating Rules

Delete Reserved Seating rules when they are no longer useful.

Follow these steps:

1. Open the Application Definitions page:
   a. Select Administration from the NFA console menu.

      The Administration page opens.

   b. Select Application Definitions in the Administration menu.

      The Application Definitions page opens.

2. Select Reserved Seating from the Rules list.

   The Application Definitions page switches to Reserved Seating mode and displays a list of the current Reserved Seating rules.

3. Click the check box next to one or more rules you want to delete. To select the check boxes for all the rules, click the check box in the heading row.

4. Click Delete.

   A confirmation message opens.

5. Click Yes.

   The list of Reserved Seating rules is updated to reflect your deletions.

Work with Port Priorities

By default, CA Network Flow Analysis defines the server port and protocol as the lower number in the flow record.

Source port: 80

Destination port: 8000

In this case, NFA will determine that the lower port is 80, and therefore http.

When the server port (tcp / udp) is a high number, the port priority can be used.

For example:

Server port: 8888

Client port: 6000

By default, NFA would use the lower port (6000) as the server port. The Port Priority functionality will tell NFA to use 8888 as the server port when it is present in the data.

Create Port Priority Rules

Create Port Priority rules to help ensure that the correct protocols are identified for each range of ports.

Follow these steps:

1. Open the **Application Definitions** page:
    a. Select **Administration** from the NFA console menu. The **Administration** page opens.
    b. Select **Application Definitions** in the **Administration** menu. The **Application Definitions** page opens.
2. Select **Port Priority** from the **Rules** list.

    The **Application Definitions** page switches to **Port Priority** mode and displays a list of the current Port Priority rules.
3. Click **Add Rule**.

    The **Add Port Priority** dialog opens.

4. Specify ports as follows:
   - Protocol: Protocol of the data that is affected by the rule, either TCP or UDP
   - **Start Port**: Target starting port for the Port Priority rule. Enter the port number in the Start Port box, a value from 0 through 65535, expressed in Base 10 decimal format. If you do not enter a value, port 0 is assigned.
   - **End Port**: Target ending port for the Port Priority rule. Enter the port number in the End Port box, a value from 0 through 65535, expressed in Base 10 decimal format. If you do not enter a value, port 0 is assigned.

     Note: The start port, end port, and protocol combination must be unique--that is, it cannot match any other Port Priority rule.
   - **Description**: (Optional) Identifying text for the Port Priority rule. The description appears in the list of Port Priority rules on the Application Definitions page.
5. Click **Save**.

   If you entered a valid start port, end port, and protocol combination and you have not yet reached the maximum number of rules, the dialog closes. The new rule appears in the list of Port Priority rules.
6. Repeat this process for each Port Priority rule you want to add.

   You can specify a maximum of 50 Port Priority rules.

Set Up Flow Cloning

You can use the Flow Cloner feature to forward flow data from a flow-enabled Harvester to another collection device, such as a Harvester in a different deployment. For example, the Flow Cloner could send flows to an Intrusion Detection System (IDS). By using the Flow Cloner, you can send the same data to two collection devices without burdening your routers with sending the data twice.

Once you have the Flow Cloner installed and configured, the flows going to the Harvester are forwarded whenever the CA NFA Flow Cloner service is running. The service starts by default whenever the server is rebooted. You can change this setting to run the service on demand. The configuration file must identify at least one destination IP address or the service will not start.

The Flow Cloner listens for packets in promiscuous mode, then forwards them to the IP addresses that you designate. In this mode, the Flow Cloner passes the packets along to any other process that is listening for them. A Harvester that is co-installed with a running Flow Cloner sees all the packets that are destined for it.

Install the Flow Cloner on the Harvester server in a distributed deployment or on the single server in a stand-alone deployment.

Note: The Flow Cloner has not affected Harvester performance significantly during testing. If you use the Flow Cloner on a high-flow Harvester server, we recommend monitoring performance.

Prerequisites for Flow Cloner Installation

Before you install the Flow Cloner, verify that your installation server meets the following prerequisites:

- The server is configured and installed or upgraded as described in the *CA Network Flow Analysis Installation Guide* or *Upgrade Guide.*
- The server already has software installed and configured to act as one of the following components:
  - Windows Harvester server in a CA Network Flow Analysis 9.3.0 distributed deployment
  - Single server in a stand-alone deployment.
- The server has at least 12.8 MB of disk space available on the target drive for the Flow Cloner.
- You exited from all other programs.
- No other user is logged in to the server.

Install the Flow Cloner

Use the steps in this topic to install the Flow Cloner.

Follow these steps:

1. Log in to the Windows-based Harvester installation server as a user with administrator privileges.

   The installation server must have the CA Network Flow Analysis 9.3.0 Harvester software installed.

2. Locate the installation program file: <install_path>\setup\FlowClonerSetup9.3.0.exe.

3. Start the installation program: For example, double-click the FlowClonerSetup9.3.0.exe file in Windows Explorer.

4. Click Next in the Welcome screen that opens.

   The Pre-Installation Summary screen opens and shows the installation path and disk space requirements. The Flow Cloner will be installed in the same root installation directory that is used for the Harvester or stand-alone software.

5. Click Install.

   The Install Complete screen opens when the installation is complete.

6. Click Done.

   The installation program closes. An installation log file named FlowCloner_Install_<timestamp>.log is created in the root installation directory.

Next: Configure the Flow Cloner options (see page 132).

Configure Flow Cloner Options

To configure the Flow Cloner, modify its default initialization (.ini) file. The .ini file contains a header line followed by a line for each destination host--each host that will receive the cloned packets. You must specify at least one destination host. If you do not specify values for the header fields, the default values are used.

For information about configuration file conventions, including how to comment out lines or fields, see Conventions for Flow Cloner Configuration Files (see page 135).

Follow these steps:

1. Log in to the Flow Cloner installation server as a user who has administrator privileges.
2. Open the following file in a text editor:
   <install_path>\Netflow\FlowCloner\flowclonedef.ini.

   The .ini file has a header line followed by a line for each host that will receive packets.

3. Customize the header line:

   The header content must be contained in a single line--the first non-commented and non-blank line in the file.

   - To use the default value for the input NIC, replace the entire header line with the following token:

     ```
     /use defaults
     ```

     You can follow the */use defaults* token with a comment, as shown in the following example:

     ```
     /use defaults ; use first available NIC and port 9995 to listen and send flows on the first available NIC
     ```

     The program uses the first available NIC. The hosts listen for the original flows and cloned flows on port UDP 9995. The */use defaults* token takes effect only if the header does not contain any other tokens.

   - (Optional) To specify the listening port, enter the /port= token, followed by the port number. The Harvester that receives the original flows listens on UDP 9995 unless you use the /port token to specify a different port.

     Default: UDP 9995

   - (Optional) To specify the destination port, enter the /dest port= token, followed by the port number. The hosts that receive the cloned flows listen on UDP 9995 unless you use the /dest port token to specify a different port. All of the hosts listen for the cloned flows on the same port.

     Default: UDP 9995

   - (Optional) To specify the Input NIC, enter the */listen ip*= token, followed by the IP address for the NIC on which the Flow Cloner listens for packets.

     Default: First functional IP address of the host

   For examples, see Examples of Flow Packet Cloner Configuration Files (see page 134).

4. Specify one or more hosts that will receive the cloned packets:

   Enter each host on a separate line, which consists of the *dest ip*= token and IP address of the destination host. You can put the destination host lines in any order.

   Example:

   ```
   /dest ip=10.0.0.100 ; send cloned packets to 10.0.0.100
   ```

   If the IP address is missing, the line is ignored.

5. Save and close the FlowCloneDef.ini file.

6. Start the CA NFA Flow Cloner service on the Harvester server.

   The Flow Cloner is enabled and attempts to forward packets to each valid destination that you specified. Flow cloning continues until you stop the CA NFA Flow Cloner service manually.

   The CA NFA Flow Cloner service is configured to start automatically on reboot and start sending cloned flow data. To operate the Flow Cloner only on demand, change this configuration in the Services window. The service can run only if the configuration file identifies at least one destination IP address.

Also see:

- Examples of Flow Cloner Configuration Files (see page 134)
- Conventions for Flow Cloner Configuration Files (see page 135)
- Characteristics of Cloned Packets (see page 136)

Examples of Flow Cloner Configuration Files

The following examples show the contents of a modified FlowCloneDef.ini file for a single-NIC server and for a multiple-NIC server environment.

Example: Default Input NIC and Port

In this example, the Flow Cloner uses the first available NIC to send the cloned data to hosts 192.100.0.100 and 192.160.0.100. The destination hosts listen for the data on UDP port 9995.

The example has a line for each host that receives the cloned data. These lines begin with the */dest ip=* token, followed by the host IP address. A semicolon precedes the comments at the end of the line.

```
/use defaults; take default setup
/dest ip=192.100.0.100; send cloned packets to ddev1
/dest ip=192.160.0.100; send cloned packets to gdev3
```

Example: Custom Input NIC and Port

In this example, the Flow Cloner uses a specified NIC (190.0.0.100) to send the cloned data to the hosts (10.0.0.000, 192.100.0.100, and 192.160.0.100). The Harvesters listen for the original incoming data on UDP port 9994.

Note: If you specify a custom */port* value, it must match the harvester listening port and the destination port that you used to configure the flow. The port values must match or no flow data is cloned.

```
/listen ip=190.0.0.100; Use 190.0.0.100 to send the cloned packets
/port=9994; Listen on port 9994 instead of the default port 9995

/dest ip=10.0.0.000 ; send cloned packets to 10.0.0.000

/dest ip=192.100.0.100; send cloned packets to ddev1

/dest ip=192.160.0.100; send cloned packets to gdev3
```

Example: Custom Input NIC and Destination Port

In this example, the Flow Cloner uses a specified NIC (190.0.0.100) to send the cloned data to host 10.0.0.000. The Harvesters that receive the cloned data listen on UDP port 9996.

```
/listen ip=190.0.0.100; Use 190.0.0.100 to send the cloned packets
/dest port=9996; Listen for the cloned data on port 9996

/dest ip=10.0.0.000 ; send cloned packets to 10.0.0.000
```

Conventions for Flow Cloner Configuration Files

Observe the following conventions for FlowCloneDef.ini files:

- Header line -- Begin the FlowCloneDef.ini file with a single line, which contains all of the header token-value entries.
- Destination host lines -- Follow the header line with a line for each destination host. Use the following format /dest ip=<X.X.X.X>, where <X.X.X.X> is the IP address of the destination host. You can put the destination host lines in any order, but the dest ip= lines must follow the header line.

- Tokens -- Precede all values in the file with a token that identifies which type of value is defined. Precede each token with a **|** or **/** character to signal that a token follows.

  Tokens are ignored if they are:

  - Unsupported
  - Entered incorrectly
  - Have no specified values

  The */use defaults* token is ignored if the header line also contains one or more token-value entries.

- Commenting Out Content:
  - Entire line -- Enter a semicolon (**;**) at the beginning of the line. The entire line is ignored.
  - Content from a particular point to the start of the next token -- Enter a semicolon (**;**) at the start of the content to ignore.
  - Everything from a particular point to the end of the line -- Enter two semicolons (**;;**) at the start of the content to ignore.

Characteristics of Cloned Packets

Cloned IP packets have the following characteristics:

- ToS value of cloned packets = 0 (zero), regardless of the original packet value. This characterizes the ToS value of the cloned packets that contain the flows. The ToS values of the traffic that the flows describe are not affected.
- Sender port = 10,000, regardless of the sender port of the original packet
- Time To Live (TTL) value = 255, regardless of the TTL value of the original packet
- Fragment ID - Arbitrary number

Chapter 9: Maintenance and Data Collection

Maintain CA Network Flow Analysis to ensure that all features are performing well. You will want to perform some maintenance tasks regularly. For example, when a user sets up a trap, the Administrator must deploy that trap.

Note: To complete the tasks that are described in this section, the role for your NFA console user account must be either Administrator or Power User.

View System Status

View the System Status page at any time to check the overall status of CA Network Flow Analysis components.

Follow these steps:

1. Display the System Status page:
   - Select Administration from the NFA console menu.

     or

   - Select System Status from the Administration menu.

   The System Status page opens and shows a quick overview of the status of CA Network Flow Analysis components. Status symbols identify any components that have generated warnings. The number of warnings appears in parentheses after the component label.
2. Click any component that has a warning symbol to see details.

   A table of warnings appears.

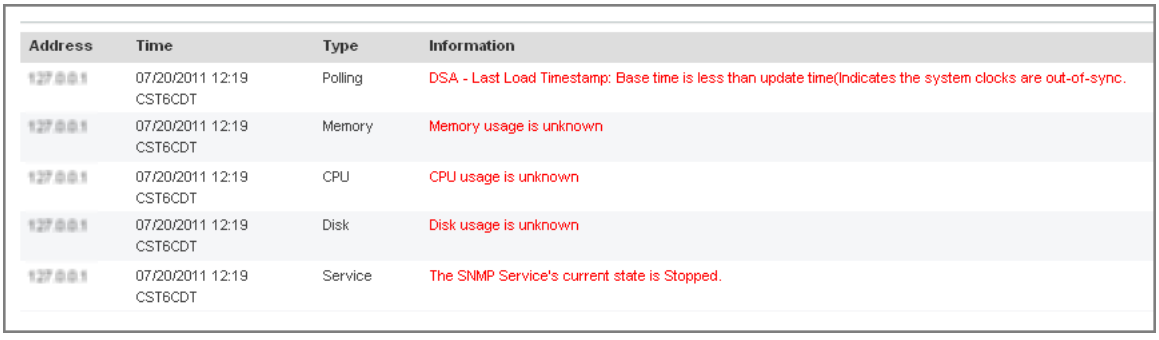

Address	Time	Type	Information
127.0.0.1	07/20/2011 12:19 CST6CDT	Polling	DSA - Last Load Timestamp: Base time is less than update time(Indicates the system clocks are out-of-sync.
127.0.0.1	07/20/2011 12:19 CST6CDT	Memory	Memory usage is unknown
127.0.0.1	07/20/2011 12:19 CST6CDT	CPU	CPU usage is unknown
127.0.0.1	07/20/2011 12:19 CST6CDT	Disk	Disk usage is unknown
127.0.0.1	07/20/2011 12:19 CST6CDT	Service	The SNMP Service's current state is Stopped.

3. Review any warnings that you find for each component.

Configure Application Settings

You can configure a wide range of settings on the Application Settings page.

Follow these steps:

1. Select Administration from the NFA console menu.

   The System Status page opens.

2. Select Administration: System: Application Settings from the Administration menu on the left.

   The Application Settings page opens.

3. (Optional) Change any of the following settings as needed, then click Save.

   Interface Data Absence Limit

   Specifies how long the program waits before it flags an interface as inactive, starting from the Last Flow value on the Available Interfaces page. When the limit is reached, the interface status changes in the following locations:

   - Interface Index: Active column value changes to No.
   - Active Interfaces page: Traffic Status value changes to Inactive/Red.

   Default: 4 Hours

   TCP Rebase Port

   Specifies the target port for TCP traffic that is redirected by an application mapping rule. TCP traffic that you do not want to go to a target port goes to the TCP Rebase Port instead. Other settings that affect application mapping behavior are UDP Rebase Port, ToS Mask, and Preserve ToS Map Proto. For more information, see the topics under Set Up Application Mapping (see page 107).

   Default: 9000

   ToS Mask

   Specifies the number of bits that application mapping rules use for matching ToS values. The default value of 255 sets the program to look for matches throughout all ToS values. Other settings that affect application mapping behavior are TCP Rebase Port, UDP Rebase Port, and Preserve ToS Map Proto. For more information, see the topics under Set Up Application Mapping (see page 107).

   Default: 255

UDP Rebase Port

Specifies the target port for UDP traffic that is redirected by an application mapping rule. UDP traffic that you do not want to go to a target port goes to the UDP Rebase Port instead. Other settings that affect application mapping behavior are TCP Rebase Port, ToS Mask, and Preserve ToS Map Proto. For more information, see the topics under Set Up Application Mapping (see page 107).

Default: 8000

Auto-Enable Interfaces

Specifies whether newly discovered interfaces are enabled automatically (True) or are disabled (False). If you want to control which interfaces are reported and consume licenses, you may want to set the value to False. In this case you enable the interfaces manually on the Available Interfaces page (see page 75). This setting affects the Enabled status for new interfaces. Interfaces that have already been discovered are not affected by changes to this setting.

Default: True

Default Time Zone

Sets the time zone for running Custom and Analysis reports. For example, a Custom report that has a reporting period of 1 day marks the start and end of the day according to the Default Time Zone. The time zone of the operator who runs the report is not used.

Default: GMT

DNS Domains

Removes the specified suffixes from host names in NFA console views and reports. If you include .my_company.com, for example, this suffix is not shown in the host names that appear in any views or reports. To specify multiple entries, separate the entries with commas and without intervening spaces.

Default: <no default>

Show Trendline Zeroes

Specifies whether trendline reports show data with interconnecting lines.

If the value is True, trend lines connect data points in reports such as the Multi Trend Summary and Stacked Trend on the Interface pages. Fill pattern is shown beneath the lines.

If the value is False, the reports show only the data points that are real. The trend line ends at the last data point and starts at the next data point. The reports show gaps wherever data points are missing. The reports do not have boundary lines for fill pattern, so the fill pattern is missing.

Default: False

From Address

Specifies the email address of the NFA Administrator, which is used as the From value when reports are emailed. If this setting is not configured properly, users cannot send scheduled or on-the-fly reports. These functions also require a properly configured SMTP Server value.

Default: <no default>

SMTP Server

Specifies the IP address of the SMTP mail server that is used for emailing reports. If this setting is not configured properly, users cannot send scheduled or on-the-fly reports. These functions also require a properly configured From Address value.

Default: <no default>

Licensed Devices

Records the total number of licenses that you purchased from CA. This value is used to calculate the percentage of licenses in use, which is displayed on the About page. The License Utilization percentage is accurate only if the Licensed Devices value is accurate.

Default: 50

Preserve ToS Map Proto

Specifies whether protocol traffic for ToS-based application-mapped data is combined (N) or is shown as separate data streams that are labeled with the original protocol designators (Y).

For example, suppose the value is Y and you map TCP, UDP, and some other IP protocol traffic to one port. To continue the example, suppose you drill in to a link in the Enterprise Overview: Top Host view for a host that has the mapped data. In this case the Stacked Protocol Trend and Protocol Trend views show and label the protocol traffic separately--whether the protocol traffic is for TCP, UDP, or some other IP protocol.*

If the Preserve ToS Map Proto value is N and the Stacked Protocol Trend views show related protocol traffic, all of the protocols for the mapped traffic are combined in a single traffic stream that has a TCP label.

* Stacked Protocol Trend and Protocol Trend views show protocol traffic that meets the following conditions: (1) The traffic passes the minimum threshold and (2) The protocol volume is high enough to place it in the Top N group.

Other settings that affect application mapping are TCP Rebase Port, UDP Rebase Port, and ToS Mask. For more information, see the topics under Set Up Application Mapping (see page 107).

Default: Y

Pump Broadcast/Multicast

Specifies whether interface views and reports include (True) or hide (False) broadcast/multicast traffic.

Default: True

Reporter IP

Specifies the IP address of the NFA console. The DSA in a 3-tier distributed deployment uses this IP address to contact the NFA console. If the setting is incorrect, the DSA cannot retrieve data files and your reports cannot display 15-minute data.

Default:

- Stand-alone deployment: Loopback IP address of the stand-alone server
- Distributed 2-tier deployment (No DSA): Loopback IP address of the NFA console
- Distributed 3-tier deployment: After you add a DSA and respond to the prompt to verify the IP address of the NFA console, the program updates the Reporter IP setting to match the IP address of the NFA console.

Report Service Polling Delay

Specifies the number of seconds between checks to see if reports have finished running. The status Complete is displayed when both of the following conditions are met:

- The Report Service check confirms that the report is finished.
- You refresh the report list on the Custom Reporting page, Analysis page, or Flow Forensics page.

Default: 15

Router Domains

Removes the specified suffixes from router names that appear in the NFA console views and reports. To specify multiple entries, separate the entries with commas and without intervening spaces.

Default: <no default>

Show Aggregations

Specifies whether to include interface aggregations in Enterprise Overview page views. If the value is True, interface aggregations are included in the views.* To be included in Enterprise Overview page views, the aggregations must have enough traffic to pass the minimum thresholds and to rank in the Top N group.

Default: False

Show Device Name

Specifies whether the interface name format starts with the device name (True) or omits it (False). This setting affects the interface names that appear in views and reports, such as the Enterprise Overview views, Interface page reports, and Custom Report Interface Summaries. For more information, see Change the Application Setting for Interface Names (see page 80).

Default: True

Display Notes Field

Displays (True) or hides (False) the Notes icon for interface rows on the Active Interfaces page. If the Notes icon is visible, you can click it to add, edit, or view additional information about an interface, as described in Active Interfaces: Interface Information (see page 58).

Default: False

Trap Destination

IP address or DNS name of the target server for sending the traps that are shown as events in the Performance Center Console--on the Events Display page (CA PC) or the Event List page (NPC). Traps can be displayed as events only when this setting is configured correctly. Set the Trap Destination value to match the IP address of one of the following servers:

- (CA PC) NFA console or stand-alone server that is registered as a data source for Performance Center
- (NPC) Event Manager server

For more information, see Configure Trap Destinations (see page 148).

Default: IP address of the NFA console (distributed deployment) or stand-alone server (stand-alone deployment)

How to Monitor the Components

Watchdog Services let you monitor CA Network Flow Analysis components. The Watchdog Services poll each server in your CA Network Flow Analysis configuration once every two hours to determine the status of all components. You can establish thresholds, an email address for receiving messages, and other settings for the Watchdog Services to ensure that you are notified of issues with the components as soon as possible.

Note: For more information about the services included on each server in your CA Network Flow Analysis configuration, see Service Management (see page 152).

Edit Watchdog Service Settings

Edit the Watchdog settings to change configuration values such as thresholds, trap settings, polling settings, notification address, and community strings.

Follow these steps:

1. Display the Watchdog Settings page:
   a. Select Administration from the NFA console menu.

      The Administration page opens.

   b. Select Health: Watchdog Settings from the Administration page menu.

      The Watchdog Settings page opens and displays the current settings.

2. Edit the Watchdog Service settings:

   Community String

   SNMP string that the Watchdog services use to verify the identity of components in a distributed deployment. The community string is used for gathering information from Harvesters (and from DSAs in a three-tier architecture deployment). Use the same community name throughout the CA Network Flow Analysis deployment:

   - Watchdog Settings page
   - SNMP service on each Windows server
   - snmpd.conf file on each Linux server

   Default: public

   CPU Threshold

   Threshold for CPU utilization. You are notified by email when the CPU threshold on a server is exceeded on any server and an SNMP trap notification is generated, provided that the address and string are set.

   Default: 80 percent CPU utilization

   Disk Threshold

   Threshold for disk utilization. If the disk threshold on a server is exceeded, you are notified by email and an SNMP trap notification is generated, provided that the address and string are set.

   Default: 80 percent disk utilization

   Email Address

   Destination email address to use for email notifications when thresholds are exceeded. To notify multiple recipients, separate the addresses with commas. The Email Address setting has no default value.

   Default: (none)

Memory Threshold

Threshold for memory utilization. You are notified by email when the memory threshold on a server is exceeded, provided that the address and string are set.

Default: 80 percent memory utilization

SNMP Retries

Number of times the program attempts to poll an SNMP device. A high number of SNMP Retries can affect performance, depending on your network configuration.

Default: 2

SNMP Timeout

Number of seconds before an SNMP poll times out.

Default: 5

System Check Interval

Number of minutes between Watchdog system checks.

Default: 60

Trap Community String

SNMP string to use for sending traps to a third-party trap receiver. Use one of the community names that the trap receiver is configured to accept.

Default: public

Trap Destination

IP address of the server that receives SNMP traps from the Watchdog Services. The traps are generated when thresholds for CA Network Flow Analysis component performance are violated.

Default: (none)

3. Click Save when you finish editing the settings.

Work with Traps

Create traps to notify network management when certain events occur. Administrators can create traps and must deploy any traps that users create.

CA Network Flow Analysis traps can be integrated with external fault management packages (see page 149).

Create Traps

Follow these steps:

1. Display the Add a Trap Configuration page:
   a. Select Administration from the NFA console menu.

      The Administration page opens.

   b. Select Alerts from the Administration page menu.

      The Trap Configuration page opens and displays a list of the current traps.

   c. Click Add.

      The Add a Trap Definition page opens.

2. Define the trap description and the interfaces that the trap monitors:
   - Description: Identification for the new trap. Include information about the protocol and threshold, such as 'Microsoft SQL Monitor utilization above 65%.'
   - Select Interface or Select Interface Group: Link to an index of interfaces or interface groups, which you use to select which interfaces that the trap monitors.

     Note: If the associated interface or interface group is deleted after you deploy the trap, an "Unknown interface group error" will occur. To eliminate this error, associate the trap with a different interface or interface group.

3. (Optional) Create a Utilization-based trap by providing the following information:
   - Thresholds: Utilization option: Setting to make the trap threshold type utilization. Use the context-sensitive fields that are added for this trap type.
   - Utilization: In field: Percentage of utilization that acts as the inbound threshold value
   - Utilization: In: minutes list: Number of minutes of inbound threshold violation that generates a trap
   - Utilization: Out field: Percentage of utilization that acts as the outbound threshold value
   - Utilization: Out: minutes list: Number of minutes of outbound threshold violation that generates a trap
   - Protocol: Protocol for the monitored traffic. Select the protocol in the Protocol Index dialog that opens. The trap will apply to the selected interfaces and protocol.

- ToS: (Optional) Type of service (ToS) of the monitored traffic The trap will apply to the selected interfaces, protocol, and ToS. To monitor all ToS traffic, select All.
- Time Filter: (Optional) Time period for monitoring if you choose to limit monitoring in this way. Select one of the available time filters from the list. The Time Filter list includes all the time filters that the Administrator created.

4. (Optional) Create a Rate-based trap by providing the following information:

   For example, suppose that you create a Rate-based trap and establish thresholds for both the inbound and outbound traffic. The trap is triggered when the threshold is met during the defined period.

   - Thresholds: Rate: Setting to base the trap threshold on data transfer rate. Use the context-sensitive fields that are added for this trap type.
   - Rate: In: Inbound data transfer speed that acts as a threshold value. Define the number of bits, kilobits, megabits, or gigabits per second, according to the unit of measurement selected.
   - Rate: Units of Measure: Unit of measurement for data transfer. Select bits per second (bps), kilobits per second (kbps), megabits per second (Mbps), or gigabits per second (Gbps).
   - Rate: In: minutes: Number of minutes of inbound threshold violation that generates a trap
   - Rate: Out: Threshold for outbound data transfer. Enter a number of bits, kilobits, megabits, or gigabits per second--according to the unit of measure for the inbound threshold.
   - Rate: Out: minutes: Number of minutes of outbound threshold violation that generates a trap
   - Protocol: Protocol for the monitored traffic. Select the protocol in the Protocol Index dialog that opens. The trap will apply to the selected interfaces and protocol.
   - ToS: (Optional) Type of service (ToS) of the monitored traffic The trap will apply to the selected interfaces, protocol, and ToS. To monitor all ToS traffic, select All.
   - Time Filter: (Optional) Time period for monitoring if you choose to limit monitoring in this way. Select one of the available time filters from the list. The Time Filter list includes all the time filters that the Administrator created.

5. (Optional) Create a Volume-based trap by providing the following information:

   - Thresholds: Volume: Setting to base the trap threshold on data volume. Use the context-sensitive fields that are added for this trap type.
   - Volume: In: Inbound data volume that acts as a threshold. Define the number of bytes, kilobytes, megabytes, gigabytes, or terabytes, according to the unit of measurement selected.

- Volume: Units of Measure: Unit of measurement for data transfer. Select bytes, kilobytes (kB), megabytes (MB), gigabytes (GB), or terabytes (TB).
- Volume: Out: Inbound data volume that acts as a threshold. Enter a number of bytes, kilobytes, megabytes, gigabytes, or terabytes--according to the unit of measure for the inbound threshold.
- Volume: Out: minutes: Number of minutes of outbound threshold violation that generates a trap.
- Protocol: Protocol for the monitored traffic. Select the protocol in the Protocol Index dialog that opens. The trap will apply to the selected interfaces and protocol.
- ToS: (Optional) Type of service (ToS) of the monitored traffic The trap will apply to the selected interfaces, protocol, and ToS. To monitor all ToS traffic, select All.
- Time Filter: (Optional) Time period for monitoring if you choose to limit monitoring in this way. Select one of the available time filters from the list. The Time Filter list includes all the time filters that the Administrator created.

6. (Optional) Create a Flow-based trap by providing the following information:
   - Thresholds: Flows: Setting to base the trap threshold on the number of flows. Use the context-sensitive fields that are added for this trap type.
   - Flows: Total: Number of inbound flows that acts as a threshold value. Define the number by flows per minute, thousands of flows per minute, or millions of flows per minute, according to the unit of measurement selected.
   - Flows: Units of Measure: Unit of measurement for data transfer. Select flows/minute, thousands of flows/minute, or millions of flows/minute.
   - Flows: minutes: Number of minutes of threshold violation that generates a trap
   - Time Filter: (Optional) Time period for monitoring if you choose to limit monitoring in this way. Select one of the available time filters from the list. The Time Filter list includes all the time filters that the Administrator created.
7. Click Submit when the selections are complete.

   The trap is deployed.

   Note: Traps are sent to the destination that is defined in the application settings.

To delete a trap definition, click the trap description in the Trap Configuration page trap list. Click the Delete button that appears in Edit mode. Verify the deletion when prompted.

To edit a trap definition, click the trap description in the Trap Configuration page trap list. Make updates in the Trap Configuration page in Edit mode. Click Submit. The trap is updated.

Configure Trap Destinations

Configure destinations for the traps that you and operators create (see page 145).

- Set the trap destination in the Application Settings page to display traps as events in the Performance Center Console.
- If you also want to generate Watchdog traps, configure the trap destination on the Watchdog Settings page in the NFA console.

Follow these steps to enable traps to be displayed in the Performance Center Console:

1. Display the Application Settings page:
   a. Select Administration from the NFA console menu.

      The Administration page opens.

   b. Select Application Settings from the menu on the left.

      The Application Settings opens.
2. Locate the Trap Destination field at the bottom of the page.
3. Verify that the Trap Destination field value at the bottom of the pages has matches the appropriate IP address:
   - (CA PC) The NFA console or stand-alone server that is registered as a data source
   - (NPC) The Event Manager server

   To check this value in CA Performance Center, make sure the Trap Destination IP address matches the Host Name value in CA Performance Center. If the host IP address is not included in the data source name, check the Host Name value by completing the following steps:
   a. Select Admin, Data Sources in the Performance Center Console.

      The Manage Data Sources page opens.

   b. Right-click the CA Network Flow Analysis data source and select Edit from the menu.

      The Edit Data Source dialog opens.

   c. Check the Host Name value.

Follow these steps to set the destination for Watchdog traps:

1. Display the Watchdog Settings page:
   a. Select Administration from the NFA console menu.

      The Administration page opens.

   b. Select Watchdog Settings from the menu on the left.

      The Watchdog Settings page opens and displays the current settings.

2. Verify that the Watchdog settings are correct, including the Trap Destination:

   Enter the IP address of the server that hosts the trap receiver. The Watchdog Services send SNMP to the Trap Destination address.

   For information about the other Watchdog settings, see Edit Watchdog Service Settings (see page 143).

3. Click Save when you finish editing the settings.

   The settings are saved. Messages are sent out for any new traps that are generated.

Optional Tasks:

- Verify that the events are displayed on the Events Display page (CA PC) or the Event List page (NPC). If the page does not show the events as expected, verify that the following conditions are met:
  - The logs show that events have been generated and forwarded to the Event Manager.
  - The Event Manager host name is resolvable by the DNS server for CA Network Flow Analysis.
  - The Trap Destination value on the Application Settings page in the NFA console matches the IP address of the NFA console or stand-alone server that is registered as a data source.
- Configure Traps for External Programs (see page 150)

How to Set Up Traps for External Fault Management Programs

You can integrate traps that CA Network Flow Analysis generates with other network management programs. The traps provide the following information, which can be useful in other programs:

- Interface that is affected by the threshold
- Protocol that is affected by the threshold
- Direction of traffic that is affected by the threshold (in, out, or both)
- Threshold that was crossed (for example, 1.23 Mb/s)
- Actual observed traffic that triggered the trap (for example, 1.30 Mb/s)
- Whether the notification indicates a newly observed threshold violation or a cleared threshold violation. For a cleared threshold violation, the reason for clearance is included. The reasons are:
  - Rate fell below the threshold.
  - No data was observed for the last time period.
  - A time filter prevented data from being considered.

Configure Traps for External Programs

Configure CA Network Flow Analysis traps to collate events with external programs and to supply useful data to the external programs.

Follow these steps:

1. Access the MIB file.

   The MIB contains information about the contents of the traps that CA Network Flow Analysis sends. The MIB file is located in the following directory: <install_path>\reporter\MIB.

2. Compile the MIB into the network management program, and configure the program to match CA Network Flow Analysis traps logically.

   A trap is sent when the program observes a new threshold violation and again when the condition is cleared. Configure the network management program to collate these two events, so you clear the warning condition when the program sends the clear condition trap. To accomplish the collation configure the program to look at the following values in the MIB:

   - NetQoSTrafficFlowEntry
   - NetQoSTrafficFlowDataEventStart, which signifies that the trap is new
   - NetQoSTrafficFlowDataEventStop, which signifies that the trap has been cleared

   For more information about defining rules for intelligent trap handling, refer to the documentation for the network management program.

3. Create and deploy CA Network Flow Analysis traps, if you have not already done so, as described in Create Traps (see page 145).

4. Verify that the trap destination contains the IP address of the server that hosts the fault management program.

   Traps are sent to the destination established on the Application Settings page.

Troubleshoot Issues with Integrated Traps

Review trap settings to correct problems with the CA Network Flow Analysis trap data that external fault management programs receive.

Follow these steps:

1. Verify that the trap destination is set to the IP address of the host server for the fault management program.
2. Verify that the traps are configured properly. For more information, see Create Traps (see page 145).

If you have checked these settings and you still do not receive traps as expected, contact CA Support for help.

Expire Stale Addresses

IP addresses can be automatically assigned by Dynamic Host Configuration Protocol (DHCP), and addresses can become outdated. You can expire stale (outdated) IP addresses to schedule them for refresh. You also can change the frequency of DNS name updates when you edit an address.

The program checks the DNS server and refreshes the DNS name for with each expired address.

Follow these steps:

1. Display the Address-Hostname Configuration page:
   a. Select Administration from the NFA console menu.

      The Administration page opens.

   b. Select Reporting: Addresses from the Administration page menu.

      The Address-Hostname Configuration page opens.
2. (Multi-domain environment only) Select the tenant-domain combination that contains the address you want to refresh.
3. Click List.

   A list opens, which shows the addresses.
4. Click the check box next to each address that you want to expire.
5. Click Expire, then confirm the expiration when the prompt opens.

   The selected addresses are scheduled to have their DNS resolution refreshed. The refresh process is typically run within 5 minutes.

Service Management

To start, stop, or check the status of services, you can use the Windows Services window or the Linux Service Configuration window.

The CA Network Flow Analysis and CA Anomaly Detector services are described briefly in the following table. The Stand-Alone column shows which services reside on a two-tier stand-alone server. The other server columns show where the services reside in a two-tier or three-tier distributed deployment. The CA Anomaly Detector services run on the installation server for that program, which may be the stand-alone, NFA console, or separate CA Anomaly Detector server.

Service	Stand-Alone	Harvester	Console	DSA	Anomaly Detector
CA NFA Collection and Poller Webservices (Linux: nfa_collpollws) Provides web service interfaces for the NFA console to communicate with the Harvester and Poller (Stand-alone and Harvester servers).	Yes	Yes			
CA NFA Data Retention (Linux: nfa_dataretention) Manages trimming to enforce volume and date limits for retaining data files. 2-tier: 15-minute, 1-minute, and Flow Forensics data (Stand-alone/Harvester) 3-tier: 15-minute data on DSA servers; 1-minute and Flow Forensics data on Harvester servers	Yes	Yes		Yes	
CA NFA DNS/SNMP Proxies (Linux: nfa_proxies) Handles SNMP and DNS requests from the Poller. Uses port 8081 by default.	Yes	Yes			
CA NFA DSALoader 3-tier: Processes 15-minute data (.rpr files) into .rpa15 files for storage on DSA servers.				Yes	
CA NFA File Server (Linux: nfa_filewebservice) 2-tier: In a distributed deployment, hosts the web service on the Harvesters to handle file requests from the NFA console Pump service. 3-tier: Hosts the web service for requesting file transfers between the DSA and the NFA console servers.	Yes	Yes	Yes (3-tier)		
CA NFA Harvester (Linux: nfa_harvester) Runs the Harvester/collector process.	Yes	Yes			

Service	Stand-Alone	Harvester	Console	DSA	Anomaly Detector
CA NFA Host Resolver Service Performs hostname lookup for any host that has anomalies detected by CA Anomaly Detector.	Yes				Yes
CA NFA Hunter Tracker Service Starts and stops the AnomalyDetector process for CA Anomaly Detector.	Yes				Yes
CA NFA Poller (Linux: nfa_poller) Initiates and handles SNMP requests to the devices that export flow to the product.	Yes	Yes			
CA NFA Pump 3-tier: Retrieves .rpr files from the NFA console.				Yes	
CA NFA Reaper (Linux: nfa_reaper) Runs the Reaper process, which processes incoming files and writes out files for 1-minute, 15-minute, and some Enterprise Overview data.	Yes	Yes			
CA NFA RibSource Provides a static interface--a RIB or Report Information Base interface--to provide the NFA data that appears in Performance Center views.	Yes		Yes		Yes
CA Performance Center SSO Runs the Single Sign-On authentication software, so users can navigate between the NFA console and Performance Center without signing on again.	Yes		Yes		
NetQoS MySql Runs the standard MySQL instance using port 3308 and stores MySQL configuration data.	Yes	Yes	Yes	Yes	Yes
NetQoS NQMySql (Linux: nfa_mysqlCSE) Runs the Custom Storage Engines and provides an interface to run queries against the 1-minute, 15-minute, and Flow Forensics data.	Yes	Yes		Yes	
NetQoS Reporter Manager Service Runs maintenance threads in the background to handle the interoperation of the components.	Yes		Yes		
NetQoS Reporter/Analyzer General Services Maintain log retention in the reporter/Logs directory and handles data retention for Enterprise Overview data.	Yes		Yes		

Service	Stand-Alone	Harvester	Console	DSA	Anomaly Detector
NetQoS Reporter/Analyzer Pump Service Tasks include: Collecting and processing Enterprise Overview data from Harvesters. 3-tier: Also collecting .rpr files from Harvesters and copying them to a staged folder for a DSA to retrieve.	Yes		Yes		
NetQoS Reporter/Analyzer Query Services Act as the interface for Custom reports and Analyses.	Yes		Yes		
NetQoS ReporterAnalyzer Report Service Executes Custom reports, Analyses, Flow Forensics, and Site to Site reports.	Yes		Yes		
NetQoS Reporter/Analyzer Watchdog Polls components for status information, in order to provide administrators with warnings. Also runs simple database integrity checks.	Yes		Yes		

Service Logs

The service logs are described in the following table. The initial part of each path is the CA Network Flow Analysis installation path, such as the default locations--C:\CA\NFA on a Windows server or /opt/CA/NFA on a Linux Harvester server. The table lists the service log file names and locations on Windows and Linux servers, along with any available configuration file to control the log level setting.

Service / Log Information	Stand-Alone	Harvester	Console	DSA
CA NFA Collection and Poller Webservices (Linux: nfa_collpollws) Log: \Netflow\Logs\collpollws-wrapper.log	Yes	Yes		
CA NFA Data Retention (Linux: nfa_dataretention) Log: \Netflow\Logs\dataretention-wrapper.log	Yes	Yes		Yes
CA NFA DNS/SNMP Proxies (Linux: nfa_proxies) Log: \Netflow\Logs\proxies-wrapper.log	Yes	Yes	Yes (2-tier)	Yes
CA NFA DSALoader Log: \Netflow\Logs\dsaLoaderErrors-<yyyy-mm-dd>.log				Yes
CA NFA File Server (Linux: nfa_filewebservice) Log: \Netflow\Logs\fileserver-wrapper.log	Yes	Yes	Yes (3-tier)	

Service / Log Information	Stand-Alone	Harvester	Console	DSA
CA NFA Harvester (Linux: nfa_harvester) Log: \Netflow\Logs\harvester-wrapper.log	Yes	Yes		
CA NFA Hunter Tracker Service Log: Logs\ADLog<yyyy-mm-dd>	Yes			
CA NFA Poller (Linux: nfa_poller) Log: \Netflow\Logs\poller-wrapper.log	Yes	Yes		
CA NFA Pump Log: \Netflow\logs\pumpservice-wrapper.log				Yes
CA NFA Reaper (Linux: nfa_reaper) Log: \Netflow\Logs\RealtimeReaperErrors<yyyy-mm-dd>.log		Yes		
CA NFA RibSource Log: \Reporter\RIB\NFA\webapps\NFARS\WEB-INF\logs\ application.log	Yes		Yes	
CA Performance Center SSO Logs: \Portal\SSO\logs\SSOService.log \Portal\SSO\logs\wrapper.log	Yes		Yes	
DNS Proxy Web Service Log: \ProxyServices\Logs\DnsProxyWSLog<yyyy-mm-dd>.log Log level setting (NFA console only): \ProxyServices\Web.config <add key="LogLevel" value="6"/> Default: 4	Yes		Yes	
NetQoS MySql (Linux: mysql) Log: \MySql\data\<server_name>.err	Yes	Yes	Yes	Yes
NetQoS NQMySql (Linux: nfa_mysqlCSE) Log: \Netflow\Logs\oursql_error.log		Yes	Yes (2-tier)	Yes
NetQoS Reporter Manager Service Log: \Reporter\Logs\Manager Service Log<yyyy-mm-dd>.log Log level setting: \Reporter\ReporterAnalyzer.ManagerService\bin\ ReporterManagerService.exe.config <setProperties Severity="6" /> Default: 4	Yes		Yes	
NetQoS Reporter Manager Service thread logs in the \Reporter\Logs\ directory: MigrationService: MigrationLog<yyyy-mm-dd>.log PollerSyncService: CollectorSyncServiceLog<yyyy-mm-dd>.log System Maintenance Service: ManagerService_MaintenanceLog<yyyy-mm-dd>.log	Yes			

Service / Log Information	Stand-Alone	Harvester	Console	DSA
NetQoS Reporter/Analyzer General Services Log: \Reporter\Logs\nqservErrors<yyyy-mm-dd>.log	Yes			
NetQoS Reporter/Analyzer Pump Service Log: \Reporter\Logs\PumpLog<yyyy-mm-dd>.log Log level setting: \Reporter\NetQoS.ReporterAnalyzer.PumpService\bin\ NetQoS.ReporterAnalyzer.PumpService.exe.config <setProperties Severity="6"/> Default value: 4	Yes			
NetQoS Reporter/Analyzer Query Services ('Reporting services') Log: \Reporter\Logs\nqreporterErrors<yyyy-mm-dd>.log Automatic Groups task log: \Reporter\Logs\ManagerService_AutomaticGroupsLog <yyyy-mm-dd>.log DNS task log: \Reporter\Logs\ManagerService_DnsLog<yyyy-mm-dd>.log Maintenance tasks log: \Reporter\Logs\ManagerService_MaintenanceLog <yyyy-mm-dd>.log	Yes			
NetQoS Reporter/Analyzer Report Service Log (Flow Forensics log): \Reporter\Logs\ReportServiceLog<yyyy-mm-dd>.log Log level setting: \Reporter\NetQoS.ReporterAnalyzer.ReportService\ bin\NetQoS.ReporterAnalyzer.ReportService.exe.config <setProperties Severity="6" /> Default value: 4 To log each poll in addition to regular event logging: <add key="VerboseLogging" value="True"/>	Yes			
NetQoS Reporter/Analyzer Watchdog Log: \Reporter\Logs\WatchdogLog<yyyy-mm-dd>.log Log level setting: \Reporter\ NetQoS.ReporterAnalyzer.WatchdogService\bin\ WatchdogService.exe.config <setProperties Severity="6"/> Default value: 4	Yes			
Multiple processes Log: \Reporter\Logs\ConsoleErrorsLog<yyyy-mm-dd>.log	Yes			
ProductSyncWS Web Service Log: \Reporter\Logs\ProductSyncWSLog<yyyy-mm-dd>.log	Yes		Yes	

How to Back Up and Restore Data

Use the procedures in this section to back up and restore the CA Network Flow Analysis databases. You can back up CA Network Flow Analysis databases manually, as described in this section, or back up the databases automatically. To back up databases automatically, add the appropriate directories to your regular backup routine, such as:

- Customized configuration files: Custom locations (Any server)
- Customized data_retention database: <install_path>\MySql\data\data_retention (Stand-alone or Harvester servers)
- Enterprise Overview data and NFA console configuration data: <install_path>\MySql\data\reporter (NFA console server)
- Historical (15-Minute) Data on a stand-alone server or a two-tier distributed deployment: <install_path>\Netflow\datafiles\ReaperArchive15 (Stand-alone or Harvester server)
- Historical (15-Minute) Data on a three-tier distributed deployment: <install_path>\MySql\data\nqrptr (DSA server)
- Harvester Configuration Data: <install_path>\mysql\data\harvester (Stand-alone or Harvester server)
- Poller Data: <install_path>\mysql\data\poller (Stand-alone or Harvester server)

Important:

- Run backups concurrently. If you restore data from backups that have different timestamps, problems can result. Ensure that your backed-up data files are timestamped with the same hour.
- Store backups to a remote location to guard against the possibility of a hardware or operating system failure on the main server. For example, back up the databases to an administrative share or mapped network drive.

Consider the general recommendations for backup frequency (see page 158).

Before you back up large databases, consider how much disk space the backups will require and the lifespan of the files you are backing up. You probably will not want to back up the raw flow data (.NFA files) or back up the entire datafiles directory.

Databases to Back Up

CA Network Flow Analysis uses several databases to store configuration data, 15-minute (historical) data, high resolution (1-minute) data, and Enterprise Overview data. This topic describes databases to consider for backups and gives general recommendations for backup frequency. Data is used differently in different environments, however, so your priorities for safeguarding against data loss may vary.

All Servers

- Customized configuration files: Back up any other customized configuration files--files that you customized or that were customized by CA Support.

  Location: Files with an .config, .conf, or .ini extension that are located anywhere in the CA Network Flow Analysis installation path

  Recommendation: Daily backup

Single Sign-On Program

(CA PC deployments) Back up any customized Single Sign-On configuration settings on the SSO server, which may be the Performance Center server. If you lose customized Single Sign-On settings, you may not be able to log in.

- Single Sign-On (SSO) configuration files, as described in the "Back Up Single Sign-On Configuration Files" topic of the *CA Performance Center Installation Guide.*

NFA Console (Distributed Deployment)

- Reporter database: Contents include the previous 24 hours of Enterprise Overview data, settings, and synchronization information.

  Location: <install_path>\mysql\data\reporter

  Recommendation: Weekly backup

Harvester (Distributed Deployment)

- Historical (15-minute) data (ReaperArchive15 database on a two-tier distributed deployment): The contents of this database include the 15-minute data that is stored for the reporting routers and interfaces.

  Location: <install_path>\Netflow\datafiles\ReaperArchive15

  Recommendation: Weekly backup

- Harvester and poller configuration files: The poller and harvester configuration data are essential to perform the relational mapping that provides access to the 15-minute data. The poller configuration data provides information about devices and interfaces to enable polling, such as persistent IDs for interfaces.

  Locations: <install_path>\MySql\data\harvester and <install_path>\MySql\data\poller

  Recommendation: Daily backup

- (Optional) Flow Forensics data (HarvesterArchive database):

  Location: <install_path>\Netflow\datafiles\HarvesterArchive

  Many administrators do not back up Flow Forensics data because of its short storage life--a maximum of 24 hours.

- (Optional) High-resolution (1-minute) data files (ReaperArchive database): Many administrators do not back up 1-minute data because its high volume requires long backups. When you decide whether to back up these files, consider these factors:

  - Value that this single day of data has in your particular circumstances
  - Long back up time

  The 1-minute data is stored for one month.

  Location: <install_path>\Netflow\datafiles\ReaperArchive

- Customized Data Retention configuration files: If you have customized any data retention settings, back up the data retention configuration data.

  Note: It is unusual to customize data retention settings, except with the assistance of CA Support. Changes to data retention settings can cause problems as the demands on drive space rise.

  Location: <install_path>\MySql\data\data_retention

  Recommendation: Daily backup

DSA (Three-Tier Distributed Deployment)

- Historical (15-minute) data (ReaperArchive15 database on a three-tier distributed deployment): The contents of this database include the 15-minute data that is stored for the reporting routers and interfaces.

  Location: <install_path>\Netflow\datafiles\ReaperArchive15

  Recommendation: Weekly backup

- Historical (15-minute) data (MySQL nqrptr database): This directory contains the DSA settings.

  Location: <install_path>\MySql\data\nqrptr

  Recommendation: Weekly backup

Stand-Alone Server

- Reporter database: The contents of this database include the previous 24 hours of Enterprise Overview data, settings, and synchronization information.

  Location: <install_path>\mysql\data\reporter

  Recommendation: Weekly backup

- Historical (15-minute) data (ReaperArchive15 database): The contents of this database include the 15-minute data that is stored for the reporting routers and interfaces.

  Location: <install_path>\Netflow\datafiles\ReaperArchive15

  Recommendation: Weekly backup

- Harvester and poller configuration files: The poller and harvester configuration data are essential to perform the relational mapping that provides access to the 15-minute data. The poller configuration data provides information about devices and interfaces to enable polling, such as persistent IDs for interfaces.

  Locations: <install_path>\MySql\data\harvester and <install_path>\MySql\data\poller

  Recommendation: Daily backup

- (Optional) Flow Forensics data (HarvesterArchive database):

  Location: <install_path>\Netflow\datafiles\HarvesterArchive

  Many administrators do not back up Flow Forensics data because of its short storage life--a maximum of 24 hours by default.

- (Optional) High-resolution (1-minute) data files (ReaperArchive database): Many administrators do not back up 1-minute data because its high volume requires long backups. When you decide whether to back up these files, consider these factors:

  - Value that this single day of data has in your particular circumstances
  - Long back up time

  The 1-minute data is stored for one month by default.

  Location: <install_path>\Netflow\datafiles\ReaperArchive

- Customized Data Retention configuration files: If you have customized any data retention settings, back up the data retention configuration data.

  Note: It is unusual to customize data retention settings, except with the assistance of CA Support. Changes to data retention settings can cause problems as the demands on drive space rise.

  Location: <install_path>\MySql\data\data_retention

  Recommendation: Daily backup

Stop the Services

To prepare for backing up the databases, stop the services on all of the Windows servers in your CA Network Flow Analysis deployment.

Follow these steps:

1. Open the Services window: Click Start, Control Panel, Administrative Tools, Services.
2. Stop the Harvester service (NetQoS Harvester service or CA NFA Harvester) on each Harvester server.
3. Wait 15 minutes for data file processing to complete.
4. Stop the CA Network Flow Analysis services on each server:

Service	Stand-Alone	Harvester	Console	DSA
CA NFA Collection and Poller Webservices (Linux: nfa_collpollws)	Yes	Yes		
CA NFA Data Retention (Linux: nfa_dataretention)	Yes	Yes		Yes
CA NFA DNS/SNMP Proxies (Linux: nfa_proxies)	Yes	Yes	Yes	Yes
CA NFA DSALoader				Yes
CA NFA File Server (Linux: nfa_filewebservice)	Yes	Yes	Yes (3-tier)	
CA NFA Harvester (Linux: nfa_harvester)	Yes	Yes		
CA NFA Poller (Linux: nfa_poller)	Yes	Yes		
CA NFA Pump				Yes
CA NFA Reaper (Linux: nfa_reaper)		Yes		
CA NFA RibSource	Yes		Yes	
NetQoS MySql (Linux: mysql)	Yes	Yes	Yes	Yes
NetQoS NQMySql (Linux: nfa_mysqlCSE)	Yes	Yes	Yes	Yes
NetQoS Reporter Manager Service	Yes		Yes	
NetQoS Reporter/Analyzer General Services	Yes		Yes	
NetQoS Reporter/Analyzer Pump Service	Yes		Yes	
NetQoS Reporter/Analyzer Query Services	Yes		Yes	
NetQoS Reporter/Analyzer Report Service	Yes		Yes	

Service	Stand-Alone	Harvester	Console	DSA
NetQoS Reporter/Analyzer Watchdog	Yes		Yes	

The services and data collection stop. The data files are processed within 15 minutes.

5. Check the following directory on the NFA console server:

   <install_path>\Netflow\datafiles\HarvesterWork

   When the HarvesterWork directory is empty, back up the databases (see page 162).

Back Up the Databases

When you have stopped the services on the CA Network Flow Analysis component servers and have verified that the HarvesterWork directories are empty, you are ready to perform the backups.

Important:

- Run backups concurrently. If you restore data from backups that have different timestamps, problems can result. Ensure that your backed-up data files are timestamped with the same hour.
- Store backups to a remote location to guard against the possibility of a hardware or operating system failure on the main server. For example, back up the databases to an administrative share or mapped network drive.

Follow these steps:

1. Determine which CA Network Flow Analysis databases and files to back up.

   The databases and files to back up for CA Network Flow Analysis are listed in the following table and are described in Databases to Back Up (see page 158).

Database	Stand-Alone Server	Harvester Servers (Distributed)	NFA Console Server (Distributed)	DSA (3-Tier Distributed)
reporter	Important		Important	
harvester	Important	Important		
poller	Important	Important		
ReaperArchive15	Recommended	Recommended		Recommended
Customized Files	Important	Important	Important	Important

Database	Stand-Alone Server	Harvester Servers (Distributed)	NFA Console Server (Distributed)	DSA (3-Tier Distributed)
Customized data_retention	Important if customized	Important if customized		Important if customized
HarvesterArchive	Optional, rarely backed up	Optional, rarely backed up		
ReaperArchive	Optional, rarely backed up	Optional, rarely backed up		
nqrptr				Important

2. Copy each of the target directories or files to a remote location.

   The following list identifies database locations:

   - Customized configuration files: Various locations
   - Customized data_retention database: <install_path>\MySql\data\data_retention
   - harvester database: <install_path>\MySql\data\harvester
   - HarvesterArchive database: <install_path>\Netflow\datafiles\HarvesterArchive
   - nqrptr database: <install_path>\MySql\data\nqrptr
   - poller database: <install_path>\MySql\data\poller
   - ReaperArchive database: <install_path>\Netflow\datafiles\ReaperArchive
   - ReaperArchive15 database: <install_path>\Netflow\datafiles\ReaperArchive15
   - reporter database: <install_path>\MySql\data\reporter
   - Single Sign-On (SSO) configuration files: Back up the following files and directories on the SSO server, which may be the Performance Center server:
     - <install_path>\Portal\SSO\start.ini file
     - <install_path>\Portal\SSO\etc directory
     - <install_path>\Portal\SSO\conf\wrapper.conf file
     - <install_path>\Portal\SSO\webapps\sso\configuration directory

3. Restart the NetQoS Mysql service and the dependent CA Network Flow Analysis services:
   a. Right-click the NetQoS Mysql service in the Services window and click Restart. The Restart Other Services confirmation message opens and lists all the dependent CA Network Flow Analysis services that will also restart.
   b. Click Yes.
   c. The confirmation message closes and the services stop and start again.
4. Restart the CA Network Flow Analysis services (see page 161).

Restore the Databases

You can restore data from a backup of a CA Network Flow Analysis database. The steps in this topic are for restoring data on a Windows server, but the database path is the same for a Linux server.

Follow these steps:

1. Log in with administrator rights.
2. Stop the CA Network Flow Analysis services (see page 161).
3. Restore each of the target directories or files from its remote location to its original location:
   - Customized configuration files: Various locations (Any server)
   - Customized data_retention database: <install_path>\MySql\data\data_retention (Stand-alone or Harvester server)
   - harvester database: <install_path>\MySql\data\harvester (Stand-alone or Harvester server)
   - nqrptr database: <install_path>\MySql\data\nqrptr (DSA server)
   - poller database: <install_path>\MySql\data\poller (Stand-alone or Harvester server)
   - HarvesterArchive database: <install_path>\Netflow\datafiles\HarvesterArchive (Stand-alone or Harvester server)
   - ReaperArchive database: <install_path>\Netflow\datafiles\ReaperArchive (Stand-alone or Harvester server)
   - ReaperArchive15 database: <install_path>\Netflow\datafiles\ReaperArchive15 (Stand-alone or Harvester server)
   - reporter database: <install_path>\MySql\data\reporter (Stand-alone, NFA console, or DSA server)
1. Restart the CA Network Flow Analysis services that are listed in the topic Stop the Services (see page 161).

Recommendations for Preserving Data Integrity

To ensure data integrity and to prevent the corruption of the CA Network Flow Analysis databases, implement the following recommendations on the servers or hardware systems you use for running CA Network Flow Analysis components:

- Exclude the following directories from real-time antivirus scans: C:\Windows\Temp and <install_path> and all its subdirectories.
- Do not implement drive space compression.
- Do not install any third-party software, except for the following types of software: antivirus, system management, and time synchronization.
- Install important Microsoft updates. Use discretion when you decide whether to install optional updates.
- Defragment the hard disk drive infrequently. Frequent defragmentation is not necessary: CA Network Flow Analysis components typically write sequentially to their database so data fragmentation is minimized.

Data Collection

CA Network Flow Analysis supports two types of deployment architectures:

- Two-tier architecture stand-alone and distributed deployments
- Three-tier architecture distributed deployments

The product components work together in both architectures to collect, process, and store flow data; display the data in reports; and generate traps, events, and scheduled reports.

Two-Tier Architecture

A two-tier architecture deployment consists of the NFA console and one or more new Harvesters. A new Harvester handles traditional Harvester tasks and the tasks that a DSA performs in a three-tier deployment.

The deployment footprint is reduced, so you have fewer component types to upgrade and maintain. Data handling requires fewer steps and files are not transferred as frequently.

A stand-alone deployment always uses the two-tier architecture: the stand-alone server hosts the NFA console and new Harvester software.

Three-Tier Architecture

A three-tier architecture deployment consists of the NFA console, one or more Harvesters, and one or more DSAs.

You may prefer the three-tier architecture if your DSA servers have lower latency than the Harvester servers. Harvester servers also have lighter storage demands.

The three-tier architecture is available only on certain types of distributed deployments, as described in the *CA Network Flow Analysis Upgrade Guide* topic "Software Versions That Are Supported for Upgrade."

Data Collectio in a Two-Tier Deployment

The NFA console and Harvester components work together in the following way on a 2-tier architecture deployment:

New Harvester

Distills the raw flows from the routers, parses the data, and compiles historical (15-minute) data and 1-minute data. The Harvester stores the following data:

- Raw flow data
- 1-minute resolution data
- Historical (15-minute) data

Note: It is important for the Harvester server to have adequate storage space, as it stores more types of data than on a three-tier architecture deployment.

NFA console

Gathers data from the Harvester as needed and displays report data in the web interface. Stores Enterprise Overview data.

Supplies a web-based user interface for administrative tasks and reports and stores the Enterprise Overview data. Reports consist of the following data:

- Enterprise Overview data that is stored locally
- 1-minute data and 15-minute data that is stored on the Harvester

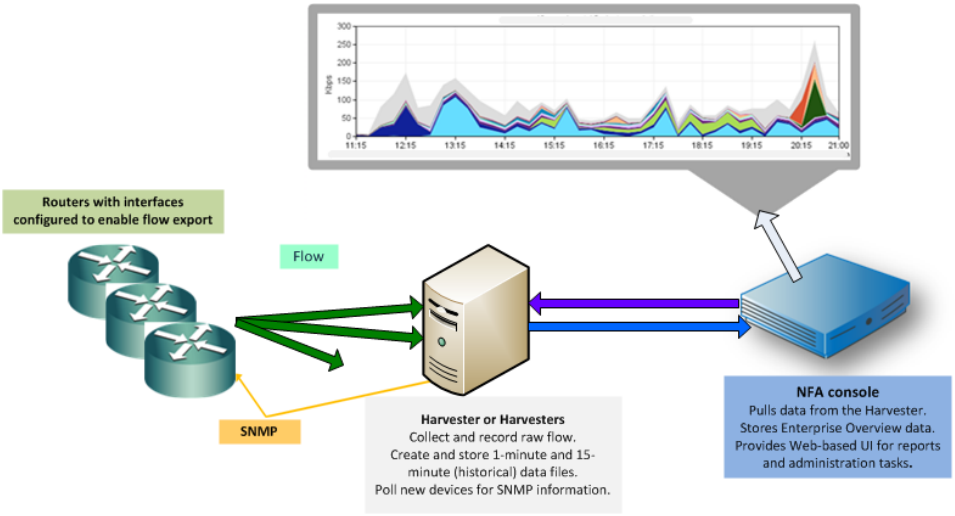

Data Collection in a Three-Tier Deployment

The NFA console, Harvester, and DSA components work together in the following way on a 3-tier architecture deployment:

Harvester

Distills the raw flows from the routers, parses the data, and records it.

Stores the 1-minute resolution data and the raw flow data. Raw flow data is stored for the previous 24 hours.

NFA console

Gathers the Enterprise Overview and 15-minute (historical) data from the Harvester and aggregates the 15-minute data.

Stores the Enterprise Overview data.

Supplies a web-based user interface for administrative tasks and reports. Reports consist of the following data:

- Enterprise Overview data that is stored locally
- 1-minute data that is stored on the Harvester
- 15-minute data that is stored on the DSA

DSA

Collects the 15-minute resolution (historical) data from the NFA console and stores it.

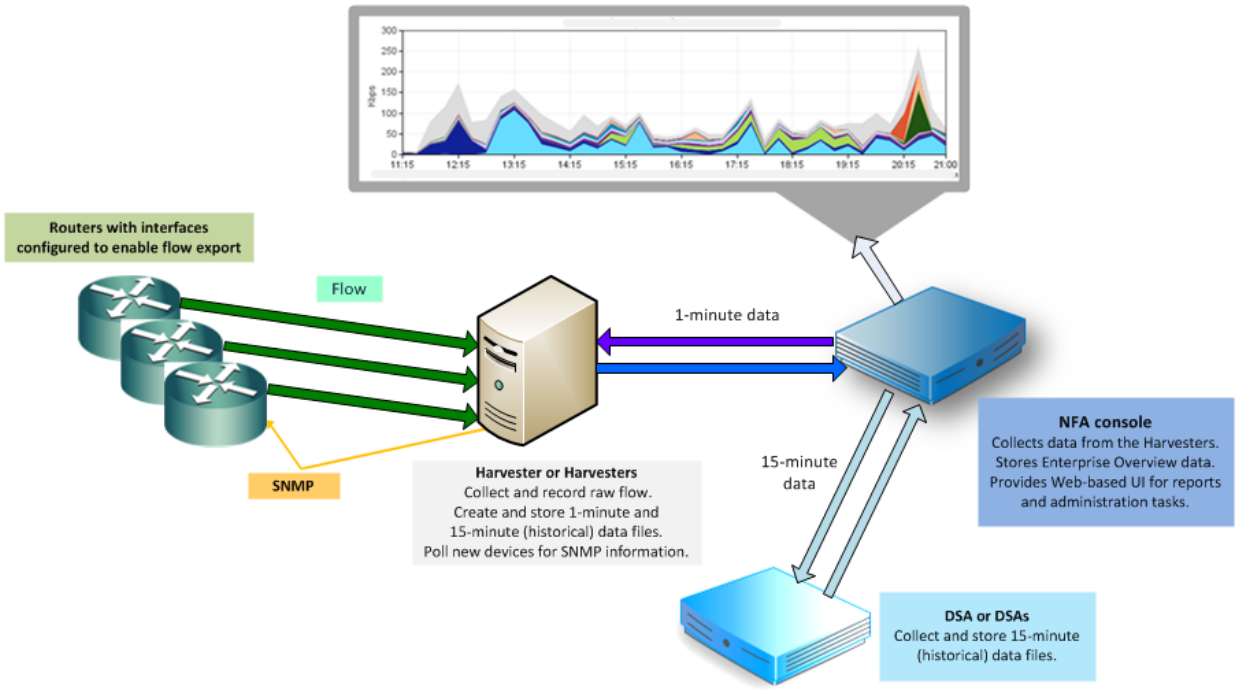

Enterprise Overview Data

The Enterprise Overview data are collected and displayed on the Enterprise Overview page. For each report, a maximum of 12 interfaces, protocols, or hosts is included from across the enterprise. Enterprise Overview views are based on the previous 24-hours of data collection. The data are displayed in the following reports:

- Interface Utilization

  Selected by: User-defined thresholds for interface capacity utilization.

- Top Interfaces - In (inbound traffic) and Top Interfaces - Out (outbound traffic)

  Selected by: Amount of interface capacity that was utilized during the report period. The utilization calculation is based on data volume and interface speeds, The interface speeds are derived from SNMP polling.

- Top Protocols

  Selected by: Volume of traffic that used the protocols.

  - Top Interfaces for Protocol

    Selected by: Interface traffic volume for the protocol that you clicked in the Top Protocols view.

  - Top Hosts for Protocol

    Selected by: Host traffic volume for the protocol that you clicked in the Top Protocols view.

- Top Hosts

  Selected by: Volume of total traffic that the hosts generated.

  - Top Interfaces for Host

    Selected by: Interface traffic volume for the host that you clicked in the Top Hosts view.

  - Top Protocols for Host

    Selected by: Protocol traffic volume for the host that you clicked in the Top Hosts view.

Data Lifespan: Enterprise Overview data is retained for 1 month by default.

Storage Location: The data is stored on the NFA console in both two-tier and three-tier deployments.

15-Minute Data

The following types of data are collected and stored as 15-minute resolution data by default.

Overall Traffic

- Total IP, TCP, and UDP traffic that travels in and out of each interface
- Flows and bytes of flow traffic that travels in and out of each interface
- Total flow traffic from each router

While 15-minute data reflect totals for each interface, Enterprise Overview data reflects totals for all interfaces. For example, Enterprise Overview data includes the Top N protocols for all interfaces. The 15-minute data includes the Top N protocols for each interface.

Hosts

Traffic to and from the top 50 IP hosts

Conversations

Traffic for the top 50 IP conversations

Protocols

- *Overall:* Top 100 protocols in and out of each interface

  Reserved Seating rules let you include specific protocols in the top protocols group regardless of their traffic volume.
- *Protocol-Hosts:* Traffic to and from the top 10 hosts for the top 20 protocols
- *Protocol-Conversations:* Traffic for the top 10 conversations for the top 20 protocols

ToS

- *Overall:* Total IP traffic in and out of each interface for each ToS value
- *ToS-Protocols:* Traffic for the top 20 protocols for the top 5 non-zero ToS values
- *ToS-Hosts:* Traffic for the top 10 hosts for the top 5 non-zero ToS values
- *ToS-Conversations:* Traffic for the top 10 conversations for the top 5 non-zero ToS values

AS

- Top 100 source and destination AS values
- Top 20 next hop addresses for each AS value

AS traffic is stored for 100 days by default.

Data Lifespan: Each data type has a default maximum storage period, as shown in the following list:

- Overall traffic: 372 days or 12.4 months
- Protocols (Top 100): 372 days or 12.4 months
- Hosts, Protocol-Hosts, and ToS-Hosts: 67 days
- Conversations (Top 50), Protocol-Conversations, and ToS-Conversations: 67 days
- ToS and ToS-Protocols: 372 days or 12.4 months
- AS: 100 days

Minimum Thresholds: To be included in a 15-minute data report, data traffic must meet or exceed the default minimum thresholds for the 15-minute period:

- Protocols: 50 KB
- Hosts: 100 KB
- Conversations: 100 KB
- ToS: No minimum
- AS: No minimum

Storage Location: Historic (15-minute) data is stored on the new Harvester in a two-tier deployment and on the DSA on a three-tier deployment.

Reports: 15-minute data is often used for examining trends and performing capacity analysis. This data appears in the following reports:

- Interface reports that show more than 2 hours of data
- Custom reports
- Analysis reports
- Performance Center views and reports that show more than 2 hours of CA Network Flow Analysis data

1-Minute Data

The following types of data are stored as 1-minute resolution data by default.

Hosts

Traffic to and from the top 300 IP hosts

Conversations

Traffic to and from the top 300 IP conversations

Protocols

- *Overall:* Top 150 protocols for traffic that travels in and out of each interface
- *Hosts:* Traffic to and from the top 25 hosts that use the top 25 protocols
- *Conversations:* Traffic for the top 25 conversations that use the top 25 protocols

ToS

- *Protocols:* Traffic for the top 25 protocols that use the top 5 ToS values
- *Hosts:* Traffic for the top 25 hosts that use the top 5 ToS values
- *Conversations:* Traffic for the top 25 conversations that use the top 5 ToS values

Data Lifespan: The default maximum storage period for 1-minute data is 1 month.

Storage Location: The 1-minute data is stored on the Harvester in both two-tier and three-tier deployments.

Reports: This type of data is often used for troubleshooting and granular analysis. The data appears in the following reports, provided that the reports are configured to show time ranges of 2 hours or less:

- Interface views and reports
- Performance Center views and reports that show less than 2 hours of CA Network Flow Analysis data

Raw Data

Raw data is used for even more granular analysis in Flow Forensics reports. Raw data is stored on the Harvester for a maximum of 24 hours by default.

Glossary

15-minute (historical) data

15-minute (historical) data is longer-range information that is collected for each interface. The information includes the protocols, hosts, and conversations for each interface. Summary data is also collected for the ToS, the top protocols for the top ToS values, and the top hosts and conversations for the top ToS values. The data is stored in a MySQL database named nqrptr, which is at <install_path>\MySql\data\nqrptr.

1-minute (high-resolution) data

1-minute (high-resolution) data is detailed information that is collected from each Harvester and is provided to the NFA console for display in views and reports. The data includes top protocols for each interface; traffic for the top hosts and conversations; top conversations for the top protocols; and top protocols, hosts, and conversations for the top ToS values. The 1-minute data is stored on the Harvester server in a database at <install_path>\Netflow\data\archive.

Administrator

An A*dministrator,* in the context of this document, is a person who is responsible for administering the product in the NFA console. An Administrator also manages elements in the Performance Center Console that are related to CA Network Flow Analysis, such as SNMP profiles, groups, users, and roles.

application mapping

Application mapping is a rule-based technique for combining the traffic for an application to facilitate reporting for the application. Application mapping rules are based on factors that can include the traffic origin (host, subnet and mask, and/or port), ToS, and protocol.

Autonomous System

Autonomous System (AS) refers to a connected group of Internet Protocol (IP) routing prefixes. The IP routing prefixes have a single, clearly defined routing policy and are controlled by one or more network operators. Meaningful AS data is available in reports only when routers and interfaces are configured to export it.

baseline

A *baseline* is a record of typical behavior, which is computed from past behavior. Baselines help you compare changes over time and predict future data or performance. Comparing current values to baseline projections is useful for determining whether current values are typical. The baseline in a trend plot is computed by using data from the six weeks before the selected date range, excluding the data point already in the trend plot.

conversation

A *conversation* is a session of subnet-to-subnet or user-to-user (host-to-host) traffic. The NFA console displays conversation information, so you can find out whether a particular conversation is causing a traffic spike on an interface, for example. You can create and run reports to identify the top volume-based conversations.

custom virtual interface

A *custom virtual interface* (CVI) is an abstract representation of a network interface, which corresponds to one or more subnets of actual physical interfaces. CVIs can give you visibility into network traffic for a carrier cloud. Set up CVIs for data center traffic that is transferred to subnets through an MPLS carrier cloud when flow is enabled on the routers in the data center.

dashboards

Dashboards are dynamic report-building pages in the Performance Center Console. Dashboards are accessible from the Dashboards tab (CA PC) or Reports tab (NPC). Each dashboard is a collection of views that present data from registered data sources on a single web page. The layout, views, time interval, and group context of each dashboard can be customized.

data sources

Data sources are the products that provide data for display in the Performance Center Console. Data sources also provide some configuration data that is stored in Performance Center. CA Network Flow Analysis is designed to be a data source for Performance Center.

drilldown report

A *drilldown report* is a more detailed report that you display by clicking a link in a report. You can open a drilldown report by clicking an interface name in an Enterprise Overview page report, for example. Properly credentialed users also can drill down from Performance Center views to detailed reports in the NFA console.

DSA (Data Storage Appliance)

A *DSA (Data Storage Appliance)* is a component in a three-tier architecture deployment of CA Network Flow Analysis. Each DSA collects 15-minute resolution (historical) data from the NFA console and stores it. In a two-tier architecture deployment, the 15-minute data is processed and stored on the Harvester.

firewall

A *firewall* server acts as a gateway between a local area network (LAN) and a large network that is not secure--such as the Internet. A firewall server typically runs a software package that inspects inbound and outbound packets, and decides whether to allow the packets to pass.

flow

A *flow* is a set of IP packets that pass a network observation point during a certain time interval. In CA Network Flow Analysis 9.3.0, flow may consist of NetFlow v5, v7, or v9 or one of the following flow types that conforms to the standards for NetFlow v5, v7, or v9: sFlow version 5; or IPFIX, J-Flow, cFlow, or Huawei NetStream flow .

In order for data from non-sampled flows to appear in reports of 15-minute (historical) data, the following minimum fields are required:

- One of the following: 1 - IN_BYTES, 85 - IN_PERMANENT_BYTES, 231 - FW_INITIATOR_OCTETS, or 232 - FW_RESPONDER_OCTETS
- All of the following: 4 - PROTOCOL, 7 - L4_SRC_PORT, 8 - IPV4_SRC_ADDR, 10 - INPUT_SNMP, 11 - L4_DST_PORT, 12 - IPV4_DST_ADDR, and 14 - OUTPUT_SNMP

group

A *group* is a collection of managed items that is organized in a tree structure. A global administrator can use Performance Center to create custom groups of the managed items that an operator can see. These managed items can be applications, servers, networks, routers, and interfaces, for example.

Harvester

A *Harvester* is a component in a distributed deployment of CA Network Flow Analysis, which collects raw flows from the routers. In a two-tier architecture deployment, the Harvester processes and stores the 1-minute and 15-minute data. In a three-tier architecture deployment, the Harvester processes and stores the 1-minute data. The NFA console retrieves and processes the 15-minute data.

IIS

IIS is the Web server that is part of the Microsoft Windows Server application. IIS consists of several services, including Simple Mail Transfer Protocol (SMTP). In versions of IIS before 5.0, IIS is an abbreviation for Internet Information Server. In version 5.0 and later, IIS is an abbreviation for Internet Information Services.

interface

An *interface* is a point of connection, such as a Serial, Frame Relay, Fast Ethernet, ATM, or PVC interface. CA Network Flow Analysis reports on any logical interface that is enabled on a supported router that has flow enabled. The NFA console displays the interfaces that are monitored in your environment.

IP domains

IP domains are logical collections of data from different devices and networks. Domains let your enterprise conduct separate monitoring of IP addresses with associated interfaces or monitor applications that belong to separate customer networks. A global administrator can monitor IP domains from a single Console, but operators view data only for the domains that they have permission to view. Administrators create custom IP domains in the Performance Center Console. Administrators can use the NFA console to assign Harvesters, routers, interfaces, CVIs, and some other elements to IP domains.

LDAP

LDAP, or Lightweight Directory Access Protocol, is a software protocol for locating organizations, individuals, and other resources, such as files and devices in a network. LDAP is based on a client/server model. The LDAP client makes a Transmission Control Protocol (TCP) connection to an LDAP server, and then sends requests and receives responses over this connection.

NetFlow

NetFlow is a transaction between two hosts, which uses a unique pair of port numbers and IP addresses and which includes certain network traffic information. A Cisco router can be configured to export flow information by sending UDP packets that contain flow statistics to one or more collectors such as the Harvesters. CA Network Flow Analysis supports NetFlow versions 5, 7, and 9 and sFlow version 5. CA Network Flow Analysis also supports IPFIX, J-Flow, cFlow, and Huawei NetStream that complies with the standards for NetFlow v5, v7, or v9.

NFA console

The *NFA console* is a component in a distributed deployment of CA Network Flow Analysis, which provides a web-based user interface for reports and for some administrative functions. The NFA console creates reports from Enterprise Overview data, which is stored locally and from the 1-minute resolution data and 15-minute resolution data that it retrieves from other components.

Performance Center

Performance Center is a term this documentation uses to refer to CA Performance Center and CA NetQoS Performance Center collectively. CA Network Flow Analysis is designed to be used with one of these programs. Page names or functions that are specific to a Performance Center version may be identified by the full program name or acronym. *CA PC* is used as an acronym for CA Performance Center and *NPC* is used for CA NetQoS Performance Center.

permission groups

Permission groups define the scope of the managed items that each user or operator can monitor. Administrators can create and assign custom groups of items to match each user's area of responsibility, such as applications, servers, networks, routers, and interfaces. Administrators assign permission groups in Performance Center to give users access to default or custom groups.

product privilege

A *product privilege* is a type of permission that is associated with a user account in Performance Center. The product privileges grant access to features in the Performance Center Console, the NFA console, and any other data sources. The administrators who manage user accounts assign product privileges in the Performance Center Console.

report

A *report* is a display of collected data, which you view in the NFA console from the Enterprise Overview, Interfaces, Custom Reporting, Flow Forensics, and Analysis pages. You can print or save reports in PDF format. You can also export reports as comma-separated value (CSV) files. An Administrator can set up some reports to be sent by email at scheduled intervals.

reporting information base (RIB)

The r*eporting information base (RIB)* is a system of web services and XML files that describe and provide the data for views and dashboards in the CA Performance Center Console. This data originates from data sources, such as CA Network Flow Analysis. The RIB capability provides an operating environment for cross-product, federated, and third-party reporting. RIB uses a single data access web service with SQL-like capabilities.

reporting period

A *reporting period* is a user-specified time range for data to be included in a CA Network Flow Analysis report. The time options vary with each report type, but the report period could consist of hours, days, weeks, or months.

Reserved Seating

Reserved Seating is a rule-based technique for ensuring that reports include the traffic that interests you, even if the traffic volume or rate is low. The rules create 'reserved seats' in reports for data that matches the target ports and protocols.

role

A *role* controls access to product features in the NFA console and the Performance Center Console. In a well-planned deployment, roles let users access the features they need to perform their duties. Roles also restrict access to features that operators and administrators do not need. The administrator who manages user accounts assigns roles in the Performance Center Console.

Single Sign-On

Single Sign-On is the authentication scheme that provides one-time login to authenticate users in the suite of related products. Once users are authenticated, they can navigate among the products without signing in again.

SMTP

SMTP (Simple Mail Transfer Protocol) is the Transfer Control Protocol/Internet Protocol (TCP/IP) protocol that is used for sending and receiving e-mail in data networks.

SNMP

SNMP (Simple Network Management Protocol) is a network management protocol that is used almost exclusively in data networks. SNMP is a method for monitoring and controlling network devices, as well as managing configurations, statistics collection, performance, and security.

SNMP profiles

SNMP profiles are definitions that contain the information for using SNMP securely to query device MIBs (Management Information Bases). Each connection to a device is made by using an SNMP profile. Administrators create SNMP profiles as needed in the Performance Center Console. In a multi-tenant CA Performance Center environment, SNMP profiles are tenant-specific. In this type of environment, each Harvester uses one of the SNMP profiles that are set up for its parent tenant.

synchronization

Synchronization, or global synchronization, is a Performance Center process that exchanges configuration and other data with CA Network Flow Analysis. For example, if an administrator creates user accounts or SNMP profiles, the associated data is pushed down to the NFA console through synchronization. Synchronization occurs every 5 minutes automatically. Administrators also can perform a full or partial synchronization on demand.

trap

A *trap* is a message that indicates a threshold has been reached or that another user-defined condition has occurred. An SNMP agent sends traps to the NFA console or to a network management system (NMS). The Watchdog agent defines a number of traps for system and application management.

trend line

A *trend line* is a projection of the future performance of an element that is based on data from past performance. CA Network Flow Analysis constructs the trend line as the best straight line through the data points of the baseline period.

Web user interface

The CA Network Flow Analysis web user interface appears as the NFA console, which lets an operator access CA Network Flow Analysis views and reports from a web browser. Administrators for CA Network Flow Analysis use this interface to perform a number of administrative functions.

Index

1

A

B

C

D

E

F

G

H

I

L

M

N

P

Q

R

S

T

U

V

W