



# Network Visibility Module

---

- [About Network Visibility Module, on page 1](#)
- [How to Use NVM, on page 3](#)
- [Collection Parameters for NVM, on page 4](#)
- [NVM Profile Editor, on page 7](#)
- [About Flow Filters, on page 11](#)
- [Customer Feedback Module Gives NVM Status, on page 12](#)

## About Network Visibility Module

Because users are increasingly operating on unmanaged devices, enterprise administrators have less visibility into what is going on inside and outside of the network. The Network Visibility Module (NVM) collects rich flow context from an endpoint on or off premise and provides visibility into network connected devices and user behaviors when coupled with a Cisco solution such as Stealthwatch, or a third-party solution such as Splunk. The enterprise administrator can then do capacity and service planning, auditing, compliance, and security analytics. NVM provides the following services:

- Monitors application use to enable better informed improvements (expanded IPFIX collector elements in nvzFlow protocol specification: <https://developer.cisco.com/site/network-visibility-module/>) in network design.
- Classifies logical groups of applications, users, or endpoints.
- Finds potential anomalies to help track enterprise assets and plan migration activities.

This feature allows you to choose whether you want the telemetry targeted as opposed to whole infrastructure deployment. The NVM collects the endpoint telemetry for better visibility into the following:

- The device—the endpoint, irrespective of its location
- The user—the one logged into the endpoint
- The application—what generates the traffic
- The location—the network location the traffic was generated on
- The destination—the actual FQDN to which this traffic was intended

When on a trusted network, AnyConnect NVM exports the flow records to a collector such as Cisco Stealthwatch or a third-party vendor such as Splunk, which performs the file analysis and provides a UI

interface and reports. The flow records provide information about the capabilities of the user, and the values are exported with ids (such as `LoggedInUserAccountType` as 12361, `ProcessUserAccountType` as 12362, and `ParentProcessUserAccountType` as 12363). For more information about Cisco Endpoint Security Analytics (CESA) built on Splunk, refer to <http://www.cisco.com/go/cesa>. Since most enterprise IT administrators want to build their own visualization templates with the data, we provide some sample base templates through a Splunk app plugin.

## NVM on Desktop AnyConnect

Historically, a flow collector provided the ability to collect IP network traffic as it enters or exits an interface of a switch or a router. It could determine the source of congestion in the network, the path of flow, but not much else. With NVM on the endpoint, the flow is augmented by rich endpoint context such as type of device, the user, the application, etc. This makes the flow records more actionable depending on the capabilities of the collection platform. The exported data provided with NVM which is sent via IPFIX is compatible with Cisco NetFlow collectors as well as other 3rd party flow collection platforms such as Splunk, IBM Qradar, LiveAction. See platform-specific integration documentation for additional information, For example, Splunk integration is available via

<https://www.cisco.com/c/en/us/support/docs/security/anyconnect-secure-mobility-client/200600-Install-and-Configure-Cisco-Network-Visi.html>.

An AnyConnect profile for NVM gets pushed from the ISE or ASA headend if this feature is enabled. On the ISE headend, you can use the standalone profile editor, generate the NVM service profile XML, upload it to ISE, and map it against the new NVM module, just as you do with Network Access Manager. On the ASA headend, you can use either the standalone or ASDM profile editor.

NVM gets notified when the VPN state changes to connected and when the endpoint is in a trusted network.



---

**Note** If you are using NVM with Linux, make sure that you have completed the preliminary steps in [Using NVM on Linux](#).

---

## Standalone NVM

For those who do not have an AnyConnect deployment or are using another VPN solution, you can install the NVM standalone package for your NVM needs. This package works independently but provides the same level of flow collection from an endpoint as the existing AnyConnect NVM solution. If you install the standalone NVM, the active processes (such as the Activity Monitor on macOS) indicate the use.

Standalone NVM is configured with the [NVM Profile Editor](#), and Trusted Network Detection (TND) configuration is mandatory. Using the TND configuration, NVM determines if the endpoint is on the corporate network and then applies the appropriate policies.

Troubleshooting and logging is still done by AnyConnect DART, which can be installed from the AnyConnect package.

## Deployment Modes

You can deploy NVM in one of two ways: 1) with the AnyConnect package or 2) with the Standalone NVM package (on AnyConnect desktop only). Refer to the *Deploy AnyConnect* chapter for steps to deploy as part of the AnyConnect package. Otherwise, you can initially install the Standalone NVM without the complete AnyConnect package by downloading the following packages:

- anyconnect-win-[version]-nvm-standalone-k9.msi (for Windows)
- anyconnect-macos-[version]-nvm-standalone.dmg (for macOS)
- anyconnect-linux64-[version]-nvm-standalone.tar.gz (for Linux)

The Standalone NVM does not depend on VPN for its functioning; therefore, you can deploy it on the endpoint without having to install VPN.

If standalone NVM is already installed, you can seamlessly migrate to a full AnyConnect installation of the same or higher version, and all NVM data files and profiles will be retained.

To upgrade to an NVM standalone configuration, you must use an out-of-band method (such as SMS) with an NVM profile. If you need both VPN and NVM functionality on the endpoint, we recommend that you deploy the AnyConnect package to install both VPN and NVM, as a separate installation is not recommended. Installation fails in the following scenarios:

- downgrading standalone NVM
- installing an older version of AnyConnect VPN with NVM where a newer version of standalone NVM already existed. This scenario would result in uninstallation of standalone NVM.
- installing any version of standalone NVM where AnyConnect VPN with NVM already existed

## NVM on Mobile AnyConnect

The Network Visibility Module (NVM) is included in the latest version of the Cisco AnyConnect Secure Mobility Client for Android, Release 4.0.09xxx, available in the Google playstore. NVM is supported on Samsung devices running Samsung Knox version 2.8 or later. No other mobile devices are currently supported.

Network Visibility on Android is part of the service profile configurations. To configure NVM on Android, an AnyConnect NVM profile is generated by the AnyConnect NVM Profile Editor, and then pushed to the Samsung mobile device using Mobile Device Management (MDM). The AnyConnect NVM Profile Editor from AnyConnect release 4.4.3 or later is required to configure NVM for mobile devices.

### Guidelines

- NVM is supported on Samsung devices running Samsung Knox version 3.0 or later. No other mobile devices are currently supported.
- On mobile devices, connectivity to the collector is supported over IPv4 or IPv6.
- Data collection traffic on Java based apps is supported.

## How to Use NVM

You can use NVM for the following scenarios:

- To audit a user's network history for potential exfiltration after a security incident occurred.
- To see how system or administrative rights impact what network connected processes are running on a user's machine.
- To get a list of all devices running a legacy OS.

- To determine what application in your network is running the highest network bandwidth.
- To determine how many versions of Firefox are being used in your network.
- To determine what percentage of Chrome.exe connections are IPv6 in your network.

## Collection Parameters for NVM

The following parameters are collected at the endpoint and exported to the collector:

**Table 1: Endpoint Identity**

Parameter	Description / Notes
Virtual Station Name	Device name configured on the endpoint (for example, Boris-Macbook)  Domain joined machines will be in the form <machinename>.<domainname>.<com> (for example, CESA-WIN10-1.mydomain.com)  Empty for Android; not provided by Samsung.
UDID	Universally Unique Identifier. Uniquely identifies the endpoint corresponding to each flow. This UDID value is also reported by HostScan in Desktop, and ACIDex in Mobile.
OS Name	Name of the operating system on the endpoint (for example, WinNT)
OS Version	Version of the operating system on the endpoint (for example, 6.1.7601)
OS Edition	The OS edition, such as Windows 8.1 Enterprise Edition
SystemManufacturer	Endpoint manufacturer (for example, Lenovo, Apple, and so on)
System Type	Set to <code>arm</code> for Android.  <code>x86</code> or <code>x64</code> for other platforms.
Agent Version	Version of NVM client software running on the endpoint. Typically of the form <code>major_v.minor_v.build_no</code>

**Table 2: Interface Information**

Parameter	Description / Notes
Endpoint UDID	Same as UDID.
Interface UID	Unique ID for an interface metadata.
Interface Index	The index of the network interface as reported by the OS.

Parameter	Description / Notes
Interface Type	Interface type, such as wired, wireless, cellular, VPN, and so on.
Interface Name	Network interface/adapter name as reported by the OS.
Interface Details List	State and SSID, attributes of InterfaceDetailsList. Indicate the network state of the interface (trusted or untrusted), and the SSID of the connection.
Interface MAC address	MAC address of the interface. Desktop only. Empty for Android (not supported).

**Table 3: Flow Information**

Parameter	Description / Notes
Source IPv4 Address	IPv4 address of the interface from where the flow was generated on the endpoint.
Destination IPv4 Address	IPv4 address of the destination to where the flow was generated from the endpoint.
Source Transport Port	Source port number from where the flow was generated on the endpoint.
Destination Transport Port	Destination port number to where the flow was generated from the endpoint.
Source IPv6 Address	IPv6 address of the interface from where the flow was generated on the endpoint. Empty for Android (not supported).
Destination IPv6 Address	IPv6 address of the destination to where the flow was generated from the endpoint. Empty for Android (not supported).
Start Sec End Sec	The absolute timestamp of the start or end of the flow in seconds.
Start Msec End Msec	The absolute timestamp of the start or end of the flow in milliseconds.
Flow UDID	Same as UDID.
Logged In User	The logged in username on the physical device, in the form Authority\Principal Empty for Android (not supported).
Logged In User Account Type	Account type of the logged in user. Empty for Android (not supported).

Parameter	Description / Notes
Process ID	Process ID of the process that initiated the network flow.
Process Name	Name of the executable generating the network flow on the endpoint.
Process Hash	Unique SHA256 hash for the executable generating the network flow on the endpoint.
Process Account	The fully qualified account, in the form Authority\Principle, under whose context the application generating the network flow on the endpoint was executed. Empty for Android (not supported).
Process Account Type	Account type of the process account. Empty for Android (not supported).
Process Path	Filesystem path of the process that initiated the network flow Empty for Android (not supported).
Process args	Command line arguments of the process that initiated the network flow, excluding the process path. Empty for Android (not supported).
Parent Process ID	Process ID of the parent of the process that initiated the network flow.
Parent Process Name	Name of the parent process of the application generating the network flow on the endpoint.
Parent Process Hash	Unique SHA256 hash for the executable of the parent process of the application generating the network flow on the endpoint. Set to 0 for Android.
Parent Process Account	The fully qualified account, in the form Authority\Principle, under whose context the parent process of the application generating the network flow on the endpoint was executed. Empty for Android (not supported).
Parent Process Account Type	Account type of the parent process account. Empty for Android (not supported).
Parent Process Path	Filesystem path of the parent of the process that initiated the network flow. Empty for Android (not supported).
Parent Process Args	Command line arguments of the parent of the process that initiated the network flow, excluding the parent process path. Empty for Android (not supported).

Parameter	Description / Notes
DNS Suffix	Configured on the interface associated with the flow on the endpoint.
L4ByteCountIn	The total number of bytes downloaded during a given flow on the endpoint at layer 4, not including L4 headers.
L4ByteCountOut	The total number of bytes uploaded during a given flow on the endpoint at layer 4, not including L4 headers.
Destination Hostname	Actual FQDN that resolved to the destination IP on the endpoint
Interface UID	Same as interface UID in interface information table. Used to identify the interface information for this flow from the interface records sent along with UDID.
Module Name List	List of 0 or more names of the modules hosted by the process that generated the flow. This can include the main DLLs in common containers, such as dllhost, svchost, rundll32, and so on. It can also contain other hosted components, such as the name of the jar file in a JVM.  Empty for Android (not supported).
Module Hash List	List of 0 or more SHA256 hashes of the modules associated with the Module Name List.  Empty for Android (not supported).

## NVM Profile Editor

In the profile editor, configure the IP address or FQDN of the collection server. You can also customize the data collection policy choosing what type of data to send, and whether data is anonymized or not.

Network Visibility Module can establish connection with a single stack IPv4 with an IPv4 address, a single stack IPv6 with an IPv6 address, or a dual stack IPv4/IPv6 to the IP address as preferred by the OS.

The mobile Network Visibility Module can establish a connection using IPv4 only. IPv6 connectivity is not supported.



**Note** The Network Visibility Module sends flow information only when it is on the trusted network. By default, no data is collected. Data is collected only when configured as such in the profile, and the data continues to be collected when the endpoint is connected. If collection is done on an untrusted network, it is cached and sent when the endpoint is on a trusted network. If you are sending collection data to Stealthwatch 7.3.1 and prior releases (or something other than Splunk or similar SIEM tool), cache data is sent once on a trusted network but not processed. For Stealthwatch applications, refer to the [Stealthwatch Enterprise Endpoint License and NVM Configuration Guide](#).

If TND is configured in the NVM profile, then the trusted network detection is done by NVM and does not depend on VPN to determine if the endpoint is in a trusted network. Also, if VPN is in a connected state, then the endpoint is considered to be on the trusted network, and the flow information is sent. The NVM-specific system logs show TND use.

When configuring TND directly in the NVM profile, an administrator-defined trusted server and certificate hash determine whether the user is on a trusted or untrusted network. Administrators configuring TND for the core VPN profile would alternatively configure the Trusted DNS Domains and Trusted DNS Servers in the core VPN profile: [AnyConnect Profile Editor, Preferences \(Part 2\)](#).

- **Desktop or Mobile**—Determines whether you are setting up NVM on a desktop or mobile device. **Desktop** is the default. Mobile will be supported in the future.
- **Collector Configuration**
  - **IP Address/FQDN**—Specifies the IPv4 or IPv6 IP address/FQDN of the collector.
  - **Port**—Specifies at which port number the collector is listening.
- **Cache Configuration**
  - **Max Size**—Specify the maximum size the database can reach. The cache size previously had a pre-set limit, but you can now configure it within the profile. The data in the cache is stored in an encrypted format, and only processes with root privileges are able to decrypt the data.  
Once a size limit is reached, the oldest data is dropped from the space for the most recent data.
  - **Max Duration**—Specify how many days of data you want to store. If you also set a max size, the limit which reaches first takes precedence.  
Once the day limit is reached, the oldest day's data is dropped from the space for the most recent day. If only Max Duration is configured, there is no size cap; if both are disabled, the size is capped at 50MB.
- **Periodic Template**—Specify the period interval at which templates are sent out from the endpoint. The default value is 1440 minutes
- **Periodic Flow Reporting**(Optional, applies to desktop only)—Click to enable periodic flow reporting. By default, NVM sends information about the flow at the end of connection (when this option is disabled). If you need periodic information on the flows even before they are closed, set an interval in seconds here. The value of 0 means the flow information is sent at the beginning and at the end of each flow. If the value is  $n$ , the flow information will be sent at the beginning, every  $n$  seconds, and at the end of each flow. Use this setting for tracking long-running connections, even before they are closed.



- **Throttle Rate**—Throttling controls at what rate to send data from the cache to the collector so that the end user is minimally impacted. You can apply throttling on both real time and cached data, as long as there is cached data. Enter the throttle rate in Kbps. The default is 500 Kbps.

The cached data is exported after this fixed period of time. Enter 0 to disable this feature.

- **Collection Mode**—Specify when data from the endpoint should be collected by choosing collection mode is off, trusted network only, untrusted network only, or all networks.
- **Collection Criteria**— You can reduce unnecessary broadcasts during data collection so that you have only relevant data to analyze. Control collection of data with the following options:
  - **Broadcast packets and Multicast packets** (Applies to desktop only)—By default, and for efficiency, broadcast and multicast packet collection are turned off so that less time is spent on backend resources. Click the check box to enable collection for broadcast and multicast packets and to filter the data.
  - **KNOX only** (Optional and mobile specific)—When checked, data is collected from the KNOX workspace only. By default, this field is not checked, and data from inside and outside the workspace is collected.
- **Data Collection Policy**—You can add data collection policies and associate them with a network type or connectivity scenario. You can apply one policy to VPN and another to non-VPN traffic since multiple interfaces can be active at the same time.

When you click Add, the Data Collection Policy window appears. Keep these guidelines in mind when creating policies:

- By default, all fields are reported and collected if no policy is created or associated with a network type.
- Each data collection policy must be associated with at least one network type, but you cannot have two policies for the same network type.
- The policy with the more specific network type takes precedence. For example, since VPN is part of the trusted network, a policy containing VPN as a network type takes precedence over a policy which has trusted as the network specified.
- You can only create a data collection policy for the network that applies based on the collection mode chosen. For example, if the **Collection Mode** is set to **Trusted Network Only**, you cannot create a **Data Collection Policy** for an **Untrusted Network Type**.
- If a profile from an earlier AnyConnect release is opened in a later AnyConnect release profile editor, it automatically converts the profile to the newer release. Conversion adds a data collection policy for all networks that exclude the same fields as were anonymized previously.
- **Name**—Specify a name for the policy you are creating.
- **Network Type**—Determine the collection mode, or the network to which a data collection policy applies, by choosing VPN, trusted, or untrusted. If you choose trusted, the policy applies to the VPN case as well.
- **Flow Filter Rule**—Defines a set of conditions and an action that can be taken to either Collect or Ignore the flow when all conditions are satisfied. You can configure up to 25 rules, and each rule can define up to 25 conditions. Use the up and down buttons to the right of the Flow Filter Rules list to adjust the priority of rules and give them higher consideration over subsequent rules. Click **Add** to set up the component of a flow filter rule.

- **Name**—The unique name of the flow filter rule.
- **Type**—Each filter rule has a Collect or Ignore type. Determine the action (Collect or Ignore) to apply if the filter rule is satisfied. If collect, the flow is allowed when conditions are met. If ignore, the flow is dropped.
- **Conditions**—Add an entry for each field that is to be matched and an operation to decide if the field value should be equal or unequal for a match. Each operation has a field identifier and a corresponding value for that field. The field matches are case sensitive unless you apply case-insensitive operations (EqualsIgnoreCase) to the rule set when you are setting up the filter engine rules. After it has been enabled, the input in the Value field set under the rule is case insensitive.

- **Include/Exclude**

- **Type**—Determine which fields you want to **Include** or **Exclude** in the data collection policy. The default is **Exclude**. All fields not checked are collected. When no fields are checked, all fields are collected.
- **Fields**—Determine what information to receive from the endpoint and which fields will be part of your data collection to meet policy requirements. Based on the network type and what fields are included or excluded, NVM collects the appropriate data on the endpoint.

For AnyConnect release 4.4 (and later), you can now choose Interface State and SSID, which specifies whether the network state of the interface is trusted or untrusted.

- **Optional Anonymization Fields**—If you want to correlate records from the same endpoint while still preserving privacy, choose the desired fields as anonymized, and they are sent as the hash of the value rather than actual values. A subset of the fields is available for anonymization.

Fields marked for include or exclude are not available for anonymization; likewise, fields marked for anonymization are not available for include or exclude.

- **Data Collection Policy for Knox (Mobile Specific)**—Option to specify data collection policy when mobile profile is selected. To create Data Collection Policy for Knox Container, choose the **Knox-Only** checkbox under Scope. Data Collection policies applied under Device Scope applies for Knox Container traffic also, unless a separate Knox Container Data Collection policy is specified. To add or remove Data Collection Policies, see Data Collection Policy description above. You can set a maximum of 6 different Data Collection Policies for mobile profile: 3 for Device, and 3 for Knox.
- **Acceptable Use Policy (Optional and mobile specific)**—Click **Edit** to define an Acceptable Use Policy for mobile devices in the dialog box. Once complete, click **OK**. A maximum of 4000 characters is allowed.

This message is shown to the user once after NVM is configured. The remote user does not have a choice to decline NVM activities. The network administrator controls NVM using MDM facilities.

- **Export on Mobile Network (Optional and Mobile Specific)**—Specifies whether the exporting of NVM flows is allowed when a device is using a mobile network. If enabled (the default value), an end user can override an administrator when an Acceptable User Policy window is displayed or later by enabling the **Settings > NVM-Settings > > Use mobile data for NVM** checkbox in an AnyConnect Android application. If you uncheck the Export on Mobile Network checkbox, NVM flows are not exported when the device is using a mobile network, and an end user cannot change that.

- **Trusted Network Detection**—This feature detects if an endpoint is physically on the corporate network. The network state is used by NVM to determine when to export NVM data and to apply the appropriate Data Collection Policy. Click **Configure** to set the configuration for Trusted Network Detection. An SSL probe is sent to the configured trusted headend, which responds with a certificate, if reachable. The thumbprint (SHA-256 hash) is then extracted and matched against the hash set in the profile editor. A successful match signifies that the endpoint is in a trusted network; however, if the headend is unreachable, or if the certificate hash does not match, then the endpoint is considered to be in an untrusted network.




---

**Note** When operating from outside your internal network, TND makes DNS requests and attempts to establish an SSL connection to the configured server. Cisco strongly recommends the use of an alias to ensure that the name and internal structure of your organization are not revealed through these requests by a machine being used outside your internal network.

---

If TND is not configured in the NVM profile and the VPN module is installed, then NVM uses the [TND feature of VPN](#) to determine if the endpoint is in a trusted network. TND configuration in the NVM profile editor includes the following:

1. **https://**—Enter the URL (IP address, FQDN, or port address) of each trusted server and click **Add**.




---

**Note** Trusted servers behind proxies are not supported.

---

2. **Certificate Hash (SHA-256)**—If the SSL connection to the trusted server is successful, this field is populated automatically. Otherwise, you can set it manually by entering the SHA-256 hash of the server certificate and clicking **Set**.
3. **List of Trusted Servers**—You can define multiple trusted servers with this process. (The maximum is 10.) Because the servers are attempted for trusted network detection in the order in which they are configured, you can use the **Move Up** and **Move Down** buttons to adjust the order. If the endpoint fails to connect to the first server, it tries the second server and so on. After trying all of the servers in the list, the endpoint waits for ten seconds before making another final attempt. When a server authenticates, the endpoint is considered within a trusted network.

Save the profile as `NVM_ServiceProfile.xml`. You must save the profile with this exact name or NVM fails to collect and send data.

## About Flow Filters

The addition of flow filters extends the current data collection policy from being just field centric, where the action is configured for a given field in each flow. With flow filter, you can create and apply rules to collect or ignore entire flows (as opposed to only particular fields), thus monitoring only the traffic of interest and potentially reducing storage requirements.

### Rule Conditions

- A rule is a match only if *all* of the conditions specified in the rule are satisfied when matched against the flow data.

- The first rule that is satisfied is applied on the flow.
- The rest of the data collection policy (include/exclude fields, anonymized fields) is also applied on the flow if it is allowed by the filter policy.
- With instances of multiple rules,
  - No action is taken on the flow if no rule matches the flow data. The default behavior is followed, which is to collect the flow.
  - If a rule matches the flow data, the action specified in that rule of the flow is applied. Subsequent rules are not checked. The order of rules, as designated in the [NVM Profile Editor](#) Flow Filter Rule parameter, indicates the priority if multiple matches could occur.

### Use of Wildcard, CIDR, and Escape Sequence Support

When entering rule conditions, you can define a wider range of field values using wildcard characters or CIDR notations, in the case of IP addresses. Also, you can use certain escape sequences in the field value. For IP fields, the CIDR/slash notation can specify the IP address that the rule should match. For example, "192.30.250.00/16" would match all addresses that have the routing prefix "192.30.0.0" derived by applying the subnet mask of "255.255.0.0." For text fields, you can use wildcards (\* and ?) and escape sequences (\\*, \?, and \\) to catch a wider range of inputs. For example, logged in user "Jane\*" would match all user names that start with "Jane."

### Sample Configurations to Achieve Flow Filtering Scenarios

To drop all UDP traffic on a particular port (such as port 53), configure a flow filter rule with type *Ignore* and two conditions:

- Condition 1: Specify that the Flow Protocol *Equals* UDP.
- Condition 2: Specify that the port number *Equals* 53.

To only collect traffic originating from only one particular process (such as Tor Browser), configure a filter rule with type *Ignore* to drop all other flows by adding one condition:

- Condition 1: Specify that the process name *Not Equals* Tor Browser.

To only collect traffic originating from only one particular IP in a subnet, configure two rules:

- Rule 1: Set a rule of type *Collect* with the condition that the IPv4 source address *Equals* 192.168.30.14.
- Rule 2: Set a second rule of type *Ignore* with the condition that the IPv4 source *Equals* 192.168.30.0/24.

## Customer Feedback Module Gives NVM Status

Part of the Customer Feedback Module collection provides data about whether NVM is installed or not, the number of flows per day, and the DB size.