Cisco Application Visibility and Control User Guide for IOS Release 15.4(1)T and IOS XE Release 3.11S

November 2013

Americas Headquarters
Cisco Systems, Inc.
170 West Tasman Drive
San Jose, CA 95134-1706
USA
http://www.cisco.com
Tel: 408 526-4000
800 553-NETS (6387)
Fax: 408 527-0883

Text Part Number: OL-30581-01

THE SPECIFICATIONS AND INFORMATION REGARDING THE PRODUCTS IN THIS MANUAL ARE SUBJECT TO CHANGE WITHOUT NOTICE. ALL STATEMENTS, INFORMATION, AND RECOMMENDATIONS IN THIS MANUAL ARE BELIEVED TO BE ACCURATE BUT ARE PRESENTED WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED. USERS MUST TAKE FULL RESPONSIBILITY FOR THEIR APPLICATION OF ANY PRODUCTS.

THE SOFTWARE LICENSE AND LIMITED WARRANTY FOR THE ACCOMPANYING PRODUCT ARE SET FORTH IN THE INFORMATION PACKET THAT SHIPPED WITH THE PRODUCT AND ARE INCORPORATED HEREIN BY THIS REFERENCE. IF YOU ARE UNABLE TO LOCATE THE SOFTWARE LICENSE OR LIMITED WARRANTY, CONTACT YOUR CISCO REPRESENTATIVE FOR A COPY.

The Cisco implementation of TCP header compression is an adaptation of a program developed by the University of California, Berkeley (UCB) as part of UCB's public domain version of the UNIX operating system. All rights reserved. Copyright © 1981, Regents of the University of California.

NOTWITHSTANDING ANY OTHER WARRANTY HEREIN, ALL DOCUMENT FILES AND SOFTWARE OF THESE SUPPLIERS ARE PROVIDED "AS IS" WITH ALL FAULTS. CISCO AND THE ABOVE-NAMED SUPPLIERS DISCLAIM ALL WARRANTIES, EXPRESSED OR IMPLIED, INCLUDING, WITHOUT LIMITATION, THOSE OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT OR ARISING FROM A COURSE OF DEALING, USAGE, OR TRADE PRACTICE.

IN NO EVENT SHALL CISCO OR ITS SUPPLIERS BE LIABLE FOR ANY INDIRECT, SPECIAL, CONSEQUENTIAL, OR INCIDENTAL DAMAGES, INCLUDING, WITHOUT LIMITATION, LOST PROFITS OR LOSS OR DAMAGE TO DATA ARISING OUT OF THE USE OR INABILITY TO USE THIS MANUAL, EVEN IF CISCO OR ITS SUPPLIERS HAVE BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

Cisco and the Cisco logo are trademarks or registered trademarks of Cisco and/or its affiliates in the U.S. and other countries. To view a list of Cisco trademarks, go to this URL: www.cisco.com/go/trademarks. Third-party trademarks mentioned are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (1110R)

Any Internet Protocol (IP) addresses and phone numbers used in this document are not intended to be actual addresses and phone numbers. Any examples, command display output, network topology diagrams, and other figures included in the document are shown for illustrative purposes only. Any use of actual IP addresses or phone numbers in illustrative content is unintentional and coincidental.

Cisco Application Visibility and Control User Guide for IOS Release 15.4(1)T and IOS XE Release 3.11S
© 2013 Cisco Systems, Inc. All rights reserved.

Preface

Revised: November 21, 2013, OL-30581-01

This preface describes the objectives, audience, organization, and conventions used in this guide and describes related documents that have additional information. It contains the following sections:

Objective

Scope

This guide provides an overview of Cisco Application Visibility and Control (AVC) and explains how to configure various Cisco AVC features for routers operating Cisco IOS or Cisco IOS XE.

Some information may not apply to your particular router model.

This guide does not provide step-by-step setup procedures for operating AVC with each management and reporting package. Refer to the documentation for your management and reporting tools, such as Cisco Prime Infrastructure or third-party tools, for step-by-step setup information.

Audience

This guide is intended for Cisco equipment providers, partners, and networking teams who are technically knowledgeable and familiar with Cisco routers and Cisco IOS software and features.

Organization

This guide is organized into the following sections.

Table 1 Organization

Chapter	Name	Description
Chapter 1	Business Overview	Describes how the Cisco AVC solution can address challenges faced by enterprise network administrators.
Chapter 2	Technology Overview	Overview of the Cisco AVC solution, including benefits, features, architecture, and interoperability.
Chapter 3	AVC Licensing and Feature Activation	Describes Cisco AVC licensing and feature activation, including temporary feature activation without a license.
Chapter 4	AVC Configuration	Describes configuration within the Cisco AVC solution, including examples.
Chapter 5	Troubleshooting	Procedures for resolving configuration issues.
Chapter 6	AVC Notes, Limitations, and Caveats	Important limitations and caveats.
Appendix A	AVC Supported Platforms and Interfaces	Platforms that support Cisco AVC, and interfaces that AVC supports.
Appendix B	AVC Feature History	Highlights of new features and optimizations in recent AVC releases.
Appendix C	References	Related documentation.
Glossary	Glossary	Glossary of terms used in this guide.

Conventions

Table 2 lists the command conventions used in this documentations to convey instructions and information.

Table 2 Command Conventions

Convention	Description
bold font	Commands and keywords.
italic font	Variables for which you supply values.
[ ]	Optional keywords or arguments appear in square brackets.
{x | y | z}	Choice of required keywords appear in braces separated by vertical bars. You have to select one.
`screen font`	Examples of information displayed on the screen.
`boldface screen font`	Examples of information you have to enter.
< >	Nonprinting characters, for example: passwords, appear in angle brackets in contexts where italics are note available.
[ ]	Default responses to system prompts appear in square brackets.

Note Means *reader take note*. Notes contain helpful suggestions or references to additional information and material.

Caution This symbol means *reader be careful*. In this situation, you might do something that could result in equipment damage or loss of data.

Tip Means *the following information will help you solve a problem*. The tips information might not be troubleshooting or even an action, but could be useful information, similar to a Timesaver.

Related Documentation

For more information, see Appendix C, "References," or see:

http://www.cisco.com/go/avc

Obtaining Documentation and Submitting a Service Request

For information on obtaining documentation, submitting a service request, and gathering additional information, see the monthly *What's New in Cisco Product Documentation*, which also lists all new and revised Cisco technical documentation:

http://www.cisco.com/en/US/docs/general/whatsnew/whatsnew.html

Subscribe to the *What's New in Cisco Product Documentation* as an RSS feed and set content to be delivered directly to your desktop using a reader application. The RSS feeds are a free service. Cisco currently supports RSS Version 2.0.

Business Overview

Revised: November 21, 2013, OL-30581-01

Introduction

Enterprise networks are carrying a growing volume of both business and recreational web traffic. Often business applications, including cloud applications such as Cisco WebEx, use the same HTTP and HTTPS protocols used by recreational web traffic. This complicates the task of optimizing network performance.

To optimize network performance and define policy for each of the applications utilizing the network, administrators need detailed visibility into the different types of applications running on the network.

The Cisco Application Visibility and Control (AVC) solution offers truly innovative and powerful capabilities of application awareness in enterprise networks. AVC incorporates into the routing devices application recognition and performance monitoring capabilities traditionally available as dedicated appliances. This integrated approach simplifies network operations, maximizes the return on network investments, and reduces the total cost of ownership.

With application awareness built into the network infrastructure, plus visibility into the performance of applications running on the network, AVC enables per-application policy for granular control of application bandwidth use, resulting in a better end user experience.

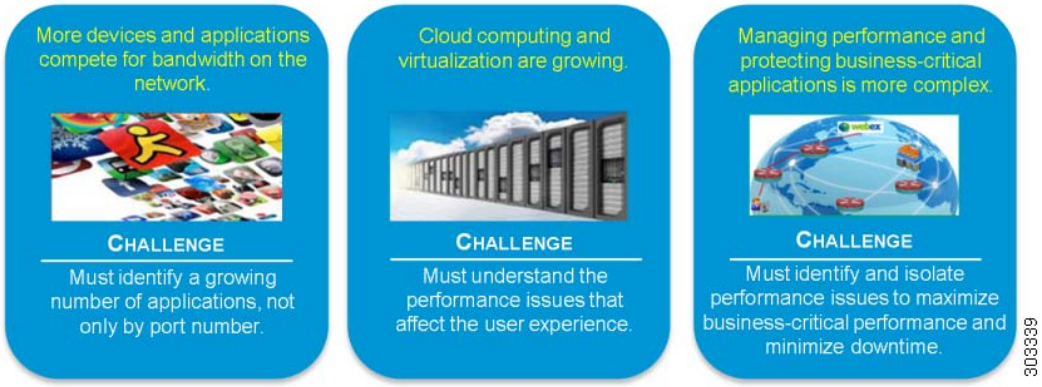

Business Use Case

The following use case illustrates how Cisco AVC can improve the user experience.

A user asks: “Why is Exchange running so slowly?”

IT engineers need answers to questions such as:

- Is Exchange actually running slowly? What are the users seeing?
- Where is the delay: branch LAN, WAN, data center LAN, or server?
- If the delay is in the network, why?
  - What is the mix of application traffic?
  - What are the key network performance metrics?

To solve the problem, IT engineers need to determine the best option. Cisco AVC offers tools to help find the best option.

- De-prioritize or block competing non-critical traffic.

  Cisco QoS tools can help.
- Send different applications over different routes.

  Cisco Performance Routing (PfR) can help.
- Squeeze more traffic over the same WAN links.

  Cisco Wide Area Application Services (WAAS) WAN optimization can help.
- Reduce apparent application latency over the WAN.

  Cisco Wide Area Application Services (WAAS) application acceleration can help.

  Or...
- Need to add more capacity?

  Cisco AVC integration with management and reporting tools, such as Cisco Prime Infrastructure, can help provide the data needed for planning new capacity.

Technology Overview

Revised: November 21, 2013, OL-30581-01

This overview of AVC technology includes the following topics:

Overview

The Cisco Application Visibility and Control (AVC) solution leverages multiple technologies to recognize, analyze, and control over 1000 applications, including voice and video, email, file sharing, gaming, peer-to-peer (P2P), and cloud-based applications. AVC combines several Cisco IOS/IOS XE components, as well as communicating with external tools, to integrate the following functions into a powerful solution.

- **Application Recognition**

  Operating on Cisco IOS and Cisco IOS XE, NBAR2 utilizes innovative deep packet inspection (DPI) technology to identify a wide variety of applications within the network traffic flow, using L3 to L7 data.

  NBAR2 can monitor over 1000 applications, and supports Protocol Pack updates for expanding application recognition, without requiring IOS upgrade or router reload.

- **Metrics Collection and Exporting**

  Metric providers, an embedded monitoring agent, and Flexible NetFlow combine to provide a wide variety of network metrics data. The monitoring agent collects:

  – TCP performance metrics such as bandwidth usage, response time, and latency.

  – RTP performance metrics such as packet loss and jitter.

  Performance metrics can be measured at multiple points within the router.

  Metrics are aggregated and exported in NetFlow v9 or IPFIX format to a management and reporting package. Metrics records are sent out directly from the data plane when possible, to maximize system performance. When more complex processing is required, such as when the router is maintaining a history of exported records, records may be exported by the route processor, which is slower than direct export from the data plane.

- **Management and Reporting Systems**

  Management and reporting systems, such as Cisco Prime Infrastructure or third-party tools, receive the network metrics data in Netflow v9 or IPFIX format, and provide a wide variety of system management and reporting functions. These functions include configuring metrics reporting, creating application and network performance reports, system provisioning, configuring alerts, and assisting in troubleshooting.

  Using the Cisco Prime Infrastructure management console, an administrator can configure each router in the network remotely using a GUI.

- **Control**

  Administrators can use industry-leading Quality of Service (QoS) capabilities to control application prioritization, manage application bandwidth, and so on. Cisco QoS employs the same deep packet inspection (DPI) technology used by NBAR2, to enable Cisco routers to reprioritize critical applications and enforce application bandwidth use.

Figure 2-1 provides a high level overview the functions of the Cisco AVC solution.

Figure 2-1 Functional Overview of the Cisco AVC Solution

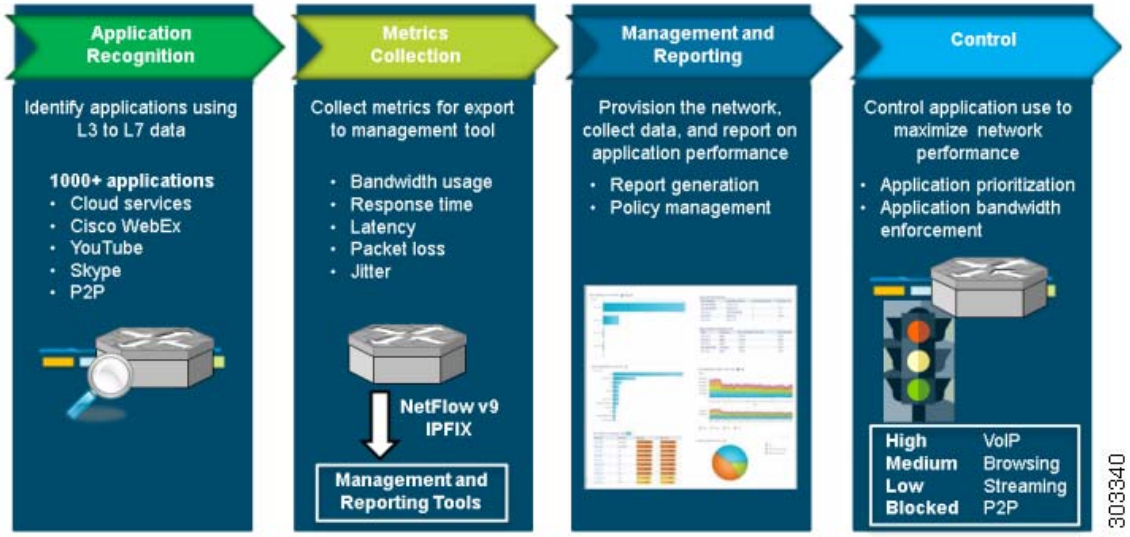

AVC Features and Capabilities

Table 2-1 describes individual Cisco AVC solution features and their availability on Cisco IOS and Cisco IOS XE platforms. For a release-by-release history of AVC features and enhancements, see Appendix B, "AVC Feature History".

Table 2-1 AVC Features
1.

Feature	Description	Available on IOS Platforms[1]	Available on IOS XE Platforms[2]
General			
Unified Solution	Cisco AVC combines application recognition, advanced metrics collection, sophisticated reporting, and network traffic control and optimization technologies into a unified solution.	Yes	Release 3.4S and later
IPv6 Support	Cisco AVC supports both IPv4 and IPv6.	Yes	Release 3.5S and later
Support on a wide range of Cisco routers operating with Cisco IOS and Cisco IOS XE	For details about supported platforms and feature activation, see: AVC Supported Platforms, page A-1 AVC Licensed Features, page 3-1	Yes	Yes
Interoperability with Cisco GET VPN	For information, see NBAR Interoperability with Cisco GET VPN, page 2-13.	—	Release 3.11S and later
Application Recognition			
Network Based Application Recognition 2 (NBAR2)	Provides application recognition. Uses an innovative deep packet inspection (DPI) technology to identify a wide variety of applications within the network traffic flow, using L3 to L7 data. NBAR2 can monitor over 1000 applications.	Yes	3.4S
Protocol Pack updates	Expands NBAR2 application recognition without requiring IOS upgrade or router reload.	Yes	3.4S
Metrics Collection			
Accounting	• Accounting of all metrics is performed by Flexible NetFlow (FNF) and the IPFIX exporter. • Multiple parallel monitors with overlapping data for the same traffic are permitted. • Flexible record keys provide different aggregation schemes for different traffic types.	Yes	3.4S
Media Monitoring	Media performance metrics are provided by Cisco Medianet technology. For more information, see: http://www.cisco.com/web/solutions/trends/medianet/index.html	Yes	3.4S

Feature	Description	Available on IOS Platforms[1]	Available on IOS XE Platforms[2]
Traffic Filtering	A policy-map defined in Cisco Common Classification Policy Language (C3PL) filters the traffic to be reported. Traffic filters operate separately from other types of policy-maps employed in the system.	Yes	3.4S
Interoperability with Cisco AppNav	Cisco AppNav is the Wide Area Application Services (WAAS) diversion mechanism. AVC provides statistics before and after the AppNav WAAS service controller (AppNav SC), as well as inspecting and reporting application information on optimized traffic. For more information about Cisco AppNav, see: http://www.cisco.com/en/US/prod/collateral/contnetw/ps5680/ps6474/white_paper_c11-705318.html	—	3.4S
Packet Capture	Cisco Embedded Packet Capture (EPC) technology performs packet capture. For more information about Cisco EPC, see: http://www.cisco.com/en/US/products/ps9913/products_ios_protocol_group_home.html	—	3.4S
Reporting on Individual Transactions	Flexible NetFlow (FNF) monitors can report on individual transactions within a flow. This enables greater resolution for traffic metrics. For more information, see: Connection/Transaction Metrics, page 4-24	—	3.9S
QoS Metrics	Cisco AVC provides monitors to collect metrics related to Quality of Service (QoS) policy. Monitors can indicate: • Packets dropped on an interface, per QoS queue, due to a QoS policy that limits resources available to a specific type of traffic. • Class hierarchy (indicating traffic priority) of a reported flow, as determined by the QoS policy map. For more information, see: QoS Metrics: Cisco IOS XE Platforms, page 4-18	Yes	3.4S
Easy Performance Monitor Configuration	The Easy Performance Monitor (“Easy perf-mon” or “ezPM”) feature provides an “express” method of provisioning monitors. Easy perf-mon provides “profiles” that represent typical deployment scenarios. After a user selects a profile and specifies a small number of parameters, Easy perf-mon provides the remaining provisioning details. In its initial release, it provided one profile, based on the exporting model improvements of the same release. Subsequent releases provide additional options. For more information, see: Easy Performance Monitor, page 4-27	15.4(1)T	3.10S
Customizing attribute values	See Customizing Attribute Values, page 4-11.	15.4(1)T	3.11S

Feature	Description	Available on IOS Platforms[1]	Available on IOS XE Platforms[2]
Management and Reporting			
Cisco Prime Infrastructure 2.0 or later	The Cisco Prime Infrastructure management and reporting system is an integral part of the Cisco AVC solution and provides extensive management and reporting features, including provisioning the system, storing exported data, and generating reports. For more information about Cisco Prime Infrastructure, see: http://www.cisco.com/en/US/products/ps12239/index.html	Yes	3.4S
Management and reporting products available from Cisco certified partners.	For information, see the Cisco Developer Network Solutions Catalog: http://marketplace.cisco.com/catalog 1. Select **Technology**. 2. In the **Technologies** list, select **Application Visibility and Control**. 3. Click **Find Solution**. A list of partner solutions appears. A **Cisco Compatible** logo indicates that the solution has passed compatibility tests with AVC. **Note** Operation of Solutions Catalog page is subject to change.	Yes	Yes
Control			
Cisco Quality of Service (QoS)	For information, see: http://www.cisco.com/go/qos	Yes	Yes

1. Applicable prior to Cisco IOS release 15.4(1)T where not specified.
2. Applicable prior to Cisco IOS XE release 3.11S where not specified.

AVC Architecture

The following Cisco AVC components are described in this section:

Figure 2-2 describes the components in the Cisco AVC architecture.

Figure 2-2 AVC Architecture for Cisco IOS and Cisco IOS XE

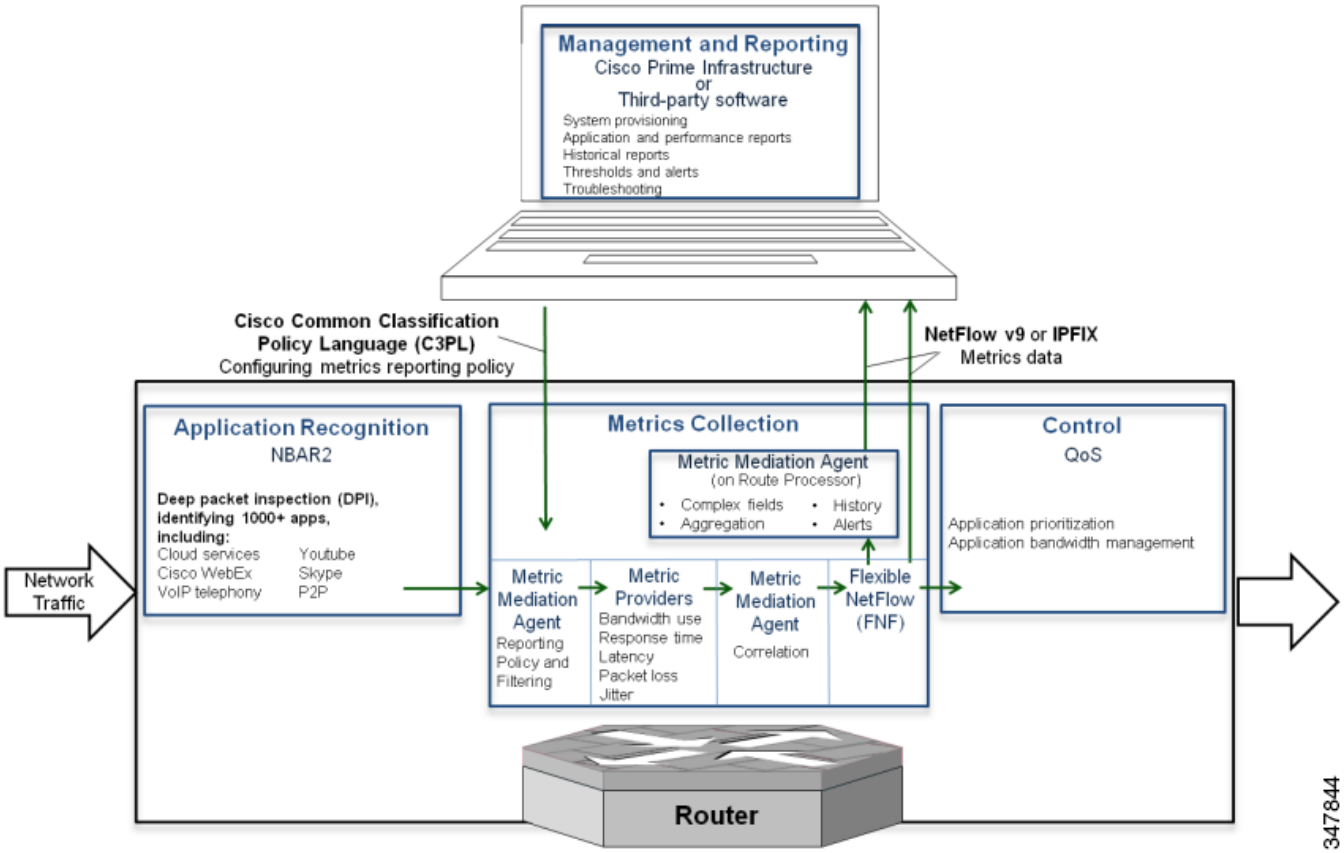

NBAR2

Network Based Application Recognition 2 (NBAR2) provides native stateful deep packet inspection (DPI) capabilities. NBAR2 is the next generation of NBAR, enhancing the application recognition engine to support more than 1000 applications.

Note Full NBAR2 functionality requires an advanced license, providing full AVC functionality. See AVC Licensed Features, page 3-1.

NBAR2 provides powerful capabilities, including:

- Categorizing applications into meaningful terms, such as category, sub-category, application group, and so on. This categorization simplifies report aggregation and control configuration.
- Field extraction of data such as HTTP URL, SIP domain, mail server, and so on. The extracted application information can be used for classification or can be exported by IPFIX to the collector for creating reports.
- Customized definition of applications, based on ports, payload values, or URL/Host of HTTP traffic.
- The set of attributes for each protocol can be customized.

Additional Application Protocol Definitions

With NBAR2 Protocol Packs, new and updated application signatures can be loaded into a router without upgrading the software image. Major protocol packs providing new and updated signatures are released periodically. Minor protocol packs are released between major releases; they provide updates and bug fixes. For information about protocol pack support, see: http://www.cisco.com/en/US/docs/ios-xml/ios/qos_nbar/prot_lib/config_library/nbar-prot-pack-library.html

In addition to the predefined application protocols, you can create customized application definitions based on ports, payload values, or URL/Host of the HTTP traffic. Protocol attributes, such as application categorization, sub-categorization, application group, and so on, can also be customized.

For more information, see: http://www.cisco.com/go/nbar

Metric Mediation Agent

Cisco IOS Platforms	Cisco IOS XE Platforms
Added in release 15.4(1)T. Prior to this release, on Cisco IOS platforms, Cisco AVC made use of the Measurement, Aggregation, and Correlation Engine (MACE). Beginning with the current release, MMA replaces MACE functionality. AVC continues to support MACE, but users are encouraged to migrate to MMA. For links to information about MACE configuration, see Appendix C, “References”.	Added in release 3.8S.

The Metric Mediation Agent (MMA) manages, correlates, and aggregates metrics from different metric providers. It provides the following functions:

- Controls traffic monitoring and filtering policy.
- Correlates data from multiple metric providers (see Metric Providers, page 2-7) into the same record.
- Aggregates metrics.
- Supports history and alert functions. This requires sending the metrics records to the route processor (RP) before exporting them to the management and reporting tools.

Metric Providers

Metric providers collect and calculate metrics and provide them to the Metric Mediation Agent (MMA) for correlation. There are a variety of metric providers: some collect simple, stateless metrics per packet, while other more complex metric providers track states and collect metrics per flow, transforming the metrics at the time of export and making sophisticated calculations. These transformations may require punting of records to the route processor (RP) before the metrics are exported to the management and reporting system.

The MMA compiles multiple metric providers of different types into the same record (see Metric Mediation Agent, page 2-7).

Flexible NetFlow

Netflow/IPFIX is the industry standard for acquiring operational data from IP networks to enable network planning, monitoring traffic analysis, and IP accounting. Flexible NetFlow (FNF) enables customizing traffic analysis parameters according to specific requirements. The AVC solution is compatible with NetFlow v9 (RFC-3954) and IPFIX (RFC-5101).

For more information, see: http://www.cisco.com/go/fnf

QoS

Cisco Quality of Service (QoS) provides prioritization, shaping, or rate-limiting of traffic. QoS can place designated applications into specific QoS classes/queues. This enables:

- Placing high priority, latency-sensitive traffic into a priority queue.
- Guaranteeing a minimum bandwidth for an individual application or for a group of applications within a QoS traffic class.

Similarly, QoS can also be used for "policing" or managing non-enterprise, recreational applications such as YouTube and Facebook.

The Cisco AVC solution integrates QoS functionality with NBAR2. QoS can use application information provided by NBAR2 in managing network traffic. The QoS class-map statements enable matching to NBAR2-supported applications and L7 application fields (such as HTTP URL or Host), as well as to NBAR2 attributes. Class-map statements can coexist with all other traditional QoS match attributes, such as IP, subnet, and DSCP.

For more information, see: http://www.cisco.com/go/qos

Embedded Packet Capture

Cisco IOS Platforms	Cisco IOS XE Platforms
Not available	Added in release 3.8S

Embedded Packet Capture (EPC) enables capturing the entire traffic for a given traffic class. The capture is limited only by available memory. The management and reporting system can read packets captured as a packet capture (pcap) file.

For more information, see: http://www.cisco.com/go/epc

Common Flow Table

The Common Flow Table (CFT) manages L4 connections and enables storing and retrieving states for each flow. Using a common flow table optimizes use of system memory and improves performance by storing and running data for each flow only once. The CFT standardizes flow management across the entire system.

Management and Reporting Systems

Cisco AVC operates with a variety of management and reporting systems.

- **Cisco Prime Infrastructure Management and Reporting**—For additional information, see Cisco Prime Infrastructure, page 2-9.
- **Third-Party Management and Reporting Solutions**—Cisco certifies solutions for AVC through the Cisco Developer Network. For a list of certified third-party management solutions, see the Cisco Developer Network Solutions Catalog:

  1. Navigate to http://marketplace.cisco.com/catalog
  2. Select **Technology**.
  3. In the **Technologies** list, select **Application Visibility and Control**.
  4. Click **Find Solution**. A list of partner solutions appears. A **Cisco Compatible** logo indicates that the solution has passed compatibility tests with AVC.

Note Operation of the Solutions Catalog page is subject to change.

Cisco Prime Infrastructure

Cisco Prime Infrastructure provides infrastructure lifecycle management and end-to-end visibility of services and applications for improved troubleshooting. It combines the solution lifecycle from design phase to monitor and troubleshooting phase.

For configuration, Cisco Prime Infrastructure has a provisioning GUI and built-in templates for enabling AVC capabilities on network devices.

For monitoring, Cisco Prime Infrastructure leverages the rich information provided by the network infrastructure, such as routers, and provides network administrators with a single tool for monitoring both network and application performance.

Network administrators can use Cisco Prime Infrastructure to drill down from an enterprise-wide network view to an individual user at a site, to proactively monitor and troubleshoot network and application performance problems.

For more information, see: http://www.cisco.com/go/primeinfrastructure

Interoperability of AVC with other Services

Cisco AVC is interoperable with many router features and services. This section provides additional information about AVC integration with AppNav WAAS, NAT, and VRF.

Interoperability with AppNav WAAS

Cisco IOS Platforms	Cisco IOS XE Platforms
Not available	Added in release 3.8S

Figure 2-3 shows a typical deployment scenario for Cisco AVC, demonstrating the integration with WAAS and the combination of optimized and pass-through traffic.

Figure 2-3 Typical AVC Deployment

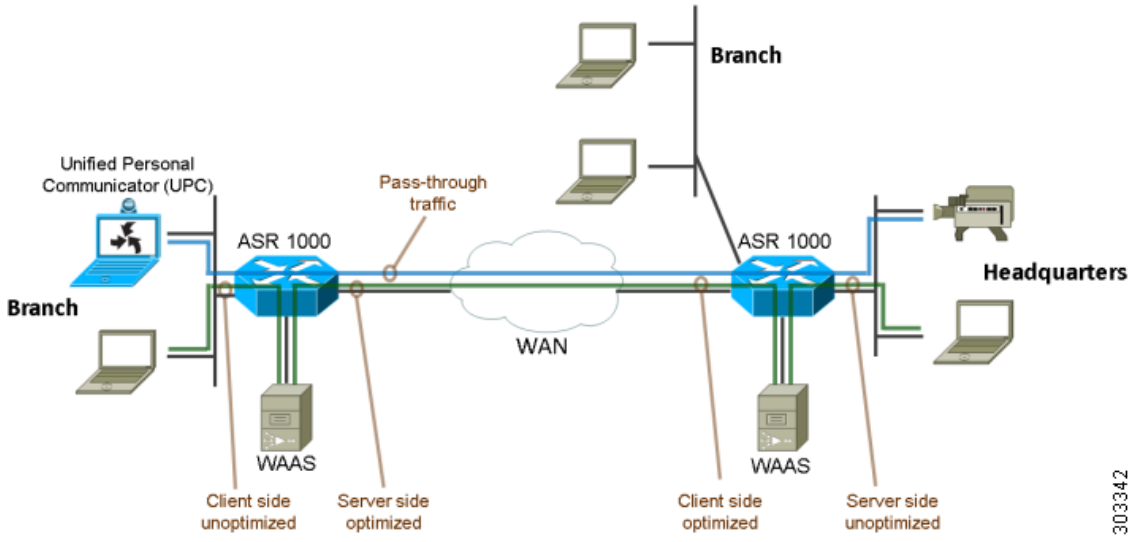

Attachment to a WAAS-Enabled Interface

Cisco Wide Area Application Services (WAAS) provides WAN optimization and application acceleration. The Cisco AVC solution operates closely with Cisco WAAS, reporting performance on both optimized and unoptimized traffic.

Figure 2-4 shows two recommended locations for metric collection. The monitoring location on the WAN interface collects metrics for optimized and unoptimized traffic. The monitoring location on the unoptimized virtual interface collects metrics for unoptimized traffic.

Figure 2-4 Recommended WAAS Monitoring Points

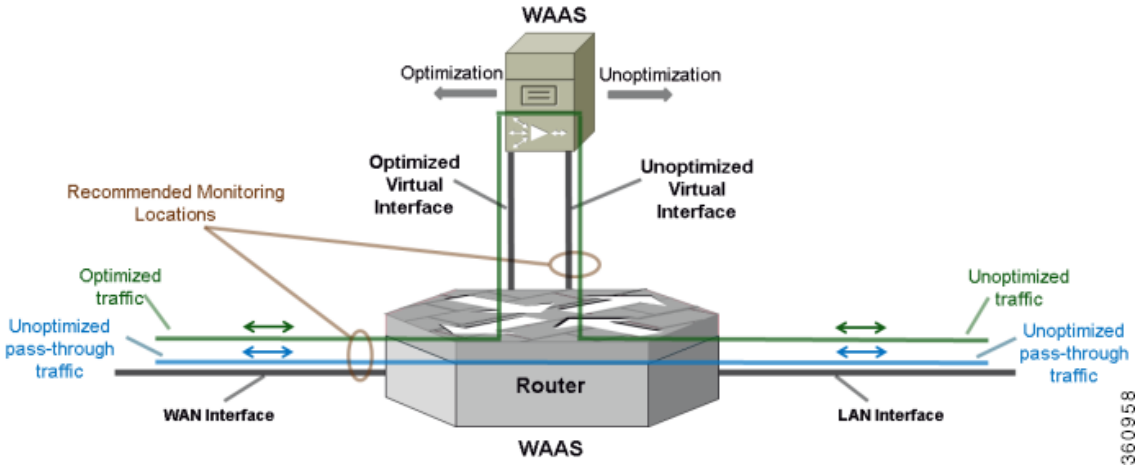

Because optimized traffic may be exported twice (pre/post WAAS), a new segment field, servicesWaasSegment, is exported within the record in order to describe the type of traffic at the monitoring location. Table 2-2 describes the segment definitions.

Table 2-2 AppNav "servicesWaasSegment" Field Values

Value	Description
0	Unknown
1	Client unoptimized
2	Server optimized
4	Client optimized
8	Server unoptimized
16	Pass-through

For pass-through traffic (bypassing WAAS), the servicesWaasPassThroughReason field indicates the reason for pass-through. See the *Cisco Application Visibility and Control Field Definition Guide for Third-Party Customers* for a description of this field.

Application Recognition on Optimized Traffic

The interoperability of Cisco AVC and WAAS enables executing traffic policies and monitoring on optimized traffic, utilizing NBAR2 application recognition.

Note When using WAAS, application L7 fields are only supported on unoptimized traffic. URL records must be attached on the unoptimized AppNav virtual interface.

Reported Input/Output Interfaces

Table 2-3 describes the input/output interface field values used by AppNav when a monitor is attached to the WAN, LAN, or an AppNav virtual interface.

Table 2-3 AppNav Exported Interfaces

Interface	Direction	Input interface value	Output interface value
WAN	Ingress	WAN	LAN
WAN	Egress	LAN	WAN
Optimized VI	Ingress	Optimized VI	LAN
Optimized VI	Egress	WAN	Optimized VI
UnOptimized VI	Ingress	UnOptimized VI	LAN
UnOptimized VI	Egress	LAN	UnOptimized VI
LAN	Ingress	LAN	WAN
LAN	Egress	WAN	LAN

AppNav Interoperability with NAT and VRF

Cisco IOS Platforms	Cisco IOS XE Platforms
Not available	Added in release 3.8S

When AppNav is enabled, it uses the virtual routing and forwarding (VRF) configuration of the LAN interface although it is installed on the WAN interface. AppNav uses the LAN VRF to divert traffic to WAAS, based on local addresses.

Up to three tuples can be used per flow. Figure 2-5 shows an example. Using more than one tuple can be necessary because of different VRF configurations and/or NAT translation. The NBAR/FNF/AppNav features in the path interact together using the same flow.

Figure 2-5 AppNav Interaction in VRF/NAT Cases

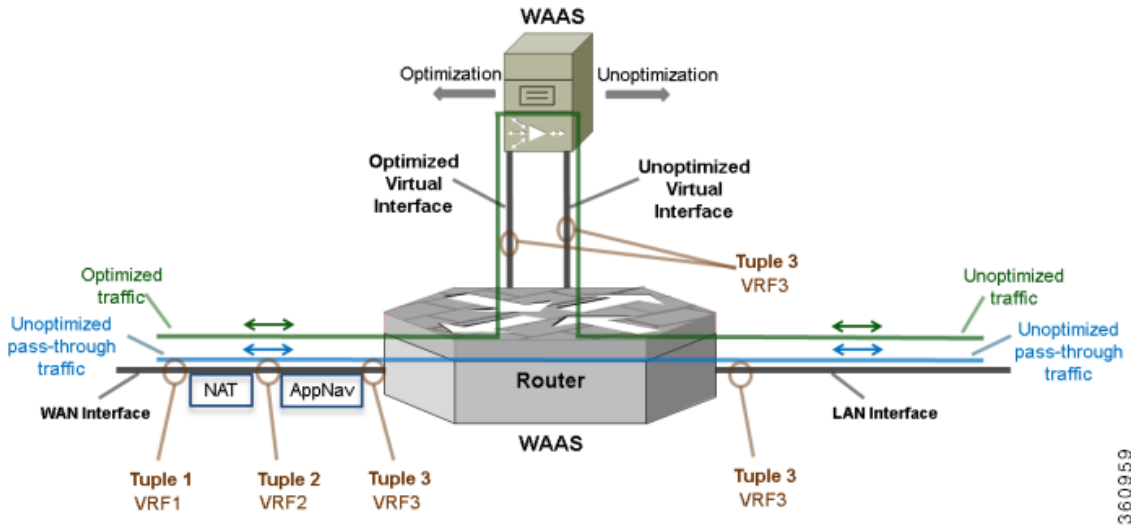

NBAR Interoperability with Cisco GET VPN

Cisco IOS Platforms	Cisco IOS XE Platforms
Not available	Added in release 3.11S

Background

Cisco Group Encrypted Transport VPN (GET VPN) is a tunnel-less VPN technology designed to provide the security of encrypted communication, with high media performance, such as lower audio/video latency, and advanced provisioning and management abilities. When using GET VPN, the router performs the encryption and decryption of the VPN traffic.

For more information about Cisco GET VPN, see: http://www.cisco.com/en/US/products/ps7180/

Encrypted Traffic and NBAR Functionality

Prior to IOS XE release 3.11S, for encrypted traffic, the NBAR component operated on the traffic in its encrypted form. As a result, NBAR was not able to provide deep packet inspection of GET VPN traffic.

Beginning with release 3.11S, NBAR operates on clear traffic (after decryption for ingress, or before encryption for egress). This enables running output QoS on inspected applications. In this release, input QoS and reporting in this release continue to operate on encrypted traffic.

To revert to the NBAR functionality that existed prior to release 3.11S, use the following command:

```
ip nbar disable classification encrypted-app
```

Note Enabling NBAR to operate on encrypted traffic requires additional processing, which may impact overall performance.

Limitations

The following limitations apply to NBAR interoperability with GET VPN:

- As in previous releases, QoS continues to operate on ingress traffic in its encrypted form, utilizing application identification information provided by the NBAR legacy component.
- In this release, only the operation of NBAR and QoS output have changed. AVC visibility functionality is not supported for GET VPN encrypted traffic.

AVC Licensing and Feature Activation

Revised: November 21, 2013, OL-30581-01

This chapter addresses Cisco AVC feature licensing and includes the following topic(s):

AVC Licensed Features

Cisco AVC software components are provided as part of each Cisco IOS and Cisco IOS XE release. No additional software packages are required for AVC functionality.

Activating full AVC functionality may require additional feature licensing and activation. License and activation details vary according to the platform. For information about supported platforms, see AVC Supported Platforms, page A-1. For platform-specific details about activating features, see AVC Feature Activation, page 3-3.

Licensing AVC

A very limited subset of AVC capabilities (recognizing approximately 150 applications based on port numbers) is available without a software license. For most platforms, the Application Experience, or AX, license is the recommended way to procure AVC functionality. Application Experience licenses provide a cost-effective combination of AVC, WAN Optimization, and Medianet capabilities.

Combined hardware-software offerings simplify the procurement of ISR G2, ISR 4451-X, ASR 1001, and ASR 1002-X routers with the Application Experience licenses. Software licenses may be used to add Application Experience capabilities to previously purchased routers.

Table 3-1 and Table 3-2 describe recommended OS images and feature licenses for a variety of platforms. For detailed, up-to-date information about images and licenses for a particular platform, refer to the documentation associated with the platform.

Table 3-1 Recommended OS Image and License—IOS Platforms

Platform	OS Image	Recommended AVC License[1]
Cisco ISR G2 (880 series)	Universal - Data For information, see: http://www.cisco.com/en/US/prod/collateral/routers/ps380/data_sheet_c78_459542_ps380_Products_Data_Sheet.html	Application Experience (AX) For information about purchasing and installing the AX license, see: *Cisco Software Activation on Integrated Services Routers*
Cisco ISR G2 (890 series)	Universal For information, see: http://www.cisco.com/en/US/prod/collateral/routers/ps380/data_sheet_c78-519930.html	Application Experience (AX) For information about purchasing and installing the AX license, see: *Cisco Software Activation on Integrated Services Routers*
Cisco ISR G2 (1900, 2900, 3900 series)	Universal	Application Experience (AX) For information about purchasing and installing the AX license, see: *Cisco Software Activation on Integrated Services Routers*

1. For the ISR G2 family, the Data license also enables the Right to Use the AVC feature set.

Table 3-2 Recommended OS Image and License—IOS XE Platforms

Platform	OS Image	Recommended AVC License
Cisco ASR 1001-5G and 1002X	AIS[1] or AES[2]	Application Experience (AX)
Cisco ASR 1000 routers not described above	AIS, AES, or Universal (depending on router model)	AVC Feature License
Cisco CSR1000v	Premium	Premium OS image includes Right to Use AVC
Cisco ISR 4451-X	Universal	Application Experience (AX) For information about purchasing and installing the AX license, see: *Cisco Software Activation on Integrated Services Routers*

1. AIS = Advanced IP Services
2. AES = Advanced Enterprise Services

AVC Feature Activation

The following sections describe the Cisco IOS/IOS XE image and license to use for full AVC feature activation, and the activation process for different platforms:

- AVC Feature Activation: Cisco ISR G2 Series, page 3-3
- AVC Feature Activation: Cisco ASR 1000 Series Routers, page 3-4
- AVC Feature Activation: Cisco ISR 4400 Series, page 3-5
- AVC Feature Activation: Cisco CSR 1000V, page 3-5

AVC Feature Activation: Cisco ISR G2 Series

Image and License Required

OS Image and License	Temporary Licence Activation Supported
See Table 3-1 Cisco IOS 15.4(1)T or later	Yes

Temporary Activation/Deactivation of the Application Experience License

Cisco ISR G2 platforms support temporary 90-day activation of Application Experience (AX) features, for evaluation, before obtaining a full license, using the `license boot module` CLI command. Activating AX features provides full AVC functionality.

Activation

To temporarily activate AX features, load the AX package and reboot the router. Execute the following from the console:

Step 1 **configure terminal**

Step 2 **license boot module** *module-name* **technology-package appxk9**

Step 3 **end**

Step 4 **reboot**

To display the *module-name* for your router, use the following command:

```
module ?
```

To display the software packages and features supported by your router, enter the following command:

```
technology-package ?
```

For additional information about activating an evaluation license, see *Software Activation on Cisco Integrated Services Routers and Cisco Integrated Service Routers G2*: http://www.cisco.com/en/US/docs/routers/access/sw_activation/SA_on_ISR.html#wp1155619

Deactivation

To deactivate the AX features, unload the AX package and reboot the router. Execute the following from the console:

Step 1 **configure terminal**

Step 2 **no license boot module** *module-name* **technology-package appxk9**

Step 3 **end**

Step 4 **reboot**

AVC Feature Activation: Cisco ASR 1000 Series Routers

Image and License Required

OS Image and License	Temporary Licence Activation Supported
See Table 3-2	Yes

Licenses

For information about purchasing and installing the AES or AIS license for Cisco ASR 1000 series routers, see:

- *Software Activation Configuration Guide, Cisco IOS XE Release 3S*
- *Cisco ASR 1000 Series Aggregation Services Routers Ordering Guide*

Temporary Activation/Deactivation of the AES or AIS License

Cisco ASR 1001 and Cisco ASR 1002-X routers support temporary 90-day activation of AES or AIS features, for evaluation, before obtaining a full license, using the `license boot level` CLI command. Activating either of these feature sets provides full AVC functionality.

Note Cisco ASR 1000 series models other than Cisco ASR 1001 and Cisco ASR 1002-X do not support temporary license activation.

Activation

To temporarily activate AES or AIS features, load the AES or AIS image and reboot the router. Execute the following from the console (using `adventerprise` for the AES image or `advipservices` for the AIS image):

```
conf t
     license boot level [adventerprise | advipservices]
end
reboot
```

Deactivation

To deactivate the AES/AIS license features, load the IPbase image and reboot the router. Execute the following from the console:

```
conf t
     license boot level ipbase
end
reboot
```

AVC Feature Activation: Cisco ISR 4400 Series

Image and License Required

OS Image and License	Temporary Licence Activation Supported
See Table 3-2	Yes

Temporary Activation/Deactivation of the Application Experience License

The Cisco ISR 4400 supports temporary 90-day activation of Application Experience (AX) features, for evaluation, before obtaining a full license, using the `license boot level` CLI command. Activating AX features provides full AVC functionality.

Activation

To temporarily activate AX features, load the AX package and reboot the router. Execute the following from the console:

```
conf t
     license boot level appxk9
end
reboot
```

Deactivation

To deactivate the AX features, unload the AX package and reboot the router. Execute the following from the console:

```
conf t
     no license boot level appxk9
end
reboot
```

AVC Feature Activation: Cisco CSR 1000V

Image and License Required

OS Image and License	Temporary Licence Activation Supported
See Table 3-2	Yes

License

For information about purchasing and installing the Application Experience (AX) license, see Cisco IOS and IOS XE Licenses, page 3-6.

Temporary Activation/Deactivation of the Premium License

Cisco CSR 1000V Cloud Services Routers support temporary 90-day activation of Premium license features, for evaluation, before obtaining a full license, using the `license boot level` CLI command. Activating Premium features provides full AVC functionality.

Activation

To temporarily activate Premium features, execute the following from the console:

```
conf t
     license boot level premium
end
reboot
```

Deactivation

To deactivate the Premium features, reboot the router and execute one of the following from the console:

Option 1:

```
conf t
     license boot level standard
end
reboot
```

Option 2:

```
conf t
     license boot level advanced
end
reboot
```

For information about images and licenses for the Cisco CSR 1000V, see:
Cisco CSR 1000V Series Cloud Services Router Release Notes

Cisco IOS Images and Licensing

Cisco IOS and IOS XE Licenses

For information about Application Experience (AX) licensing, see:

- *Application Experience*

  http://www.cisco.com/en/US/netsol/ns1226/index.html
- *Cisco Software Activation on Integrated Services Routers*

  http://www.cisco.com/en/US/products/ps10616/index.html

Universal Image and Software Activation License

The feature activation section for each platform indicates whether it supports use of a universal IOS XE software image. The universal software image includes all IOS XE functionality. You can purchase new software capabilities at any time for a deployed router. With the purchase, you receive a Product Activation Key (PAK). To activate the purchased functionality, you enter the PAK into the Software Activation License (SAL), which is preinstalled on the device.

The software activation licensing system simplifies IOS XE software deployment. The IOS XE software image remains unchanged, regardless of which functionality has been activated, and only one archive image must be maintained per device.

AVC Configuration

Revised: November 21, 2013, OL-30581-01

This chapter addresses Cisco AVC configuration and includes the following topics:

Recent Configuration Enhancements and Limitations

Table 4-1 describes configuration features added in recent releases, and limitations.

Table 4-1 Configuration Features and Enhancements

Feature	IOS Platforms	IOS XE Platforms	Information/Limitations
Export Spreading	Added in IOS 15.4(1)T	Added in IOS XE 3.11S	For information, see NetFlow/IPFIX Flow Monitor, page 4-4
Easy Performance Monitor "express" method of provisioning monitors	Added in IOS 15.4(1)T	Added in IOS XE 3.10S	For information, see Easy Performance Monitor, page 4-27
Support for configuring 40 fields for each FNF record	Not applicable	Added in IOS XE 3.10S	For limitations, see: Downgrading to an IOS XE Version that Does Not Support More than 32 Fields, page 6-2
CLI field aliases	Added in IOS 15.4(1)T	Added in IOS XE 3.10S	For limitations, see: Removing Aliases before Downgrading from Cisco IOS 15.4(1)T / Cisco IOS XE 3.10 or Later, page 6-2

Unified Policy CLI

Cisco IOS Platforms	Cisco IOS XE Platforms
Added in release 15.4(1)T	Added in release 3.8S

Monitoring a configuration is done using performance-monitor unified monitor and policy.

Configuration Format

```
policy-map type performance-monitor <policy-name>
   [no] parameter default account-on-resolution
   class <class-map name>
      flow monitor <monitor-name> [sampler <sampler name>]
      [sampler <sampler name>]
      monitor metric rtp
```

Usage Guidelines

- Supports:
  - Multiple flow monitors under a class-map
  - Up to 5 monitors per attached class-map
  - Up to 256 classes per performance-monitor policy
- No support for:
  - Hierarchical policy
  - Inline policy
- Metric producer parameters are optional.

- Account-on-resolution (AOR) configuration causes all classes in the policy-map to work in AOR mode, which delays the action until the class-map results are finalized (the application is determined by NBAR2).

Attaching a Policy

Attach a policy to the interface using following command:

```
interface <interface-name>
    service-policy type performance-monitor <policy-name> {input|output}
```

Displaying Policy Map Performance Monitor Data

Display policy map performance monitor data using the command below. Example output is shown here.

- On Cisco IOS platforms, the data is reported once per flow, either for the first packet of the flow or for the packet of resolution if AOR is enabled.
- On Cisco IOS XE platforms, the data is reported for all packets that match the policy map.

```
Router# show policy-map type performance-monitor interface
 Ethernet1/0

  Service-policy performance-monitor input: policy

    Class-map: classmap (match-all)
      20 packets, 1280 bytes
      5 minute offered rate 0000 bps, drop rate 0000 bps
      Match: access-group name seawolf_acl_ipv4_tcp

    Class-map: class-default (match-any)
      0 packets, 0 bytes
      5 minute offered rate 0000 bps, drop rate 0000 bps
      Match: any

  Service-policy performance-monitor output: policy

    Class-map: classmap (match-all)
      20 packets, 1160 bytes
      5 minute offered rate 0000 bps, drop rate 0000 bps
      Match: access-group name seawolf_acl_ipv4_tcp

    Class-map: class-default (match-any)
      0 packets, 0 bytes
      5 minute offered rate 0000 bps, drop rate 0000 bps
      Match: any
```

Metric Producer Parameters

Metric producer-specific parameters are optional and can be defined for each metric producer for each class-map.

Configuration Format

```
monitor metric rtp
    clock-rate {type-number| type-name | default} rate
    max-dropout number
    max-reorder number
    min-sequential number
    ssrc maximum number
```

Reacts

The **react** CLI defines the alerts applied to a flow monitor. The **react** CLI has a performance impact on the router. When possible, send the monitor records directly to the Management and Reporting system and apply the network alerts in the Management and Reporting system.

Note **Cisco IOS XE Platforms**: Applying reacts on the device requires punting the monitor records to the route processor (RP) for alert processing. To avoid the performance reduction of punting the monitor records to the RP, send the monitor records directly to the Management and Reporting system, as described above.

Configuration Format

```
react <id> [media-stop|mrv|rtp-jitter-average|transport-packets-lost-rate]
```

NetFlow/IPFIX Flow Monitor

Cisco IOS Platforms	Cisco IOS XE Platforms
export-spread feature added in IOS 15.4(1)T	**export-spread** feature added in IOS XE 3.11S

Flow monitor defines monitor parameters, such as record, exporter, and other cache parameters.

Configuration Format: Cisco IOS Platforms

```
flow monitor type performance-monitor <monitor-name>
    record <name | default-rtp | default-tcp>
    exporter <exporter-name>
    history size <size> [timeout <interval>]
    cache entries <num>
    cache timeout {{active | inactive} <value> | synchronized <value> {export-spread
    <interval>}}
    cache type {permanent | normal | immediate}
    react-map <react-map-name>
```

Configuration Format: Cisco IOS XE Platforms

```
flow monitor type performance-monitor <monitor-name>
    record <name | default-rtp | default-tcp>
    exporter <exporter-name>
    history size <size> [timeout <interval>]
    cache entries <num>
    cache timeout {{active | inactive} <value> | synchronized <value>
    {export-spread <interval>} event transaction end}
    cache type {permanent | normal | immediate}
    react-map <react-map-name>
```

Usage Guidelines

- The **react-map** CLI is allowed under the class in the policy-map. In this case, the monitor must include the exporting of the class-id in the flow record. The route processor (RP) correlates the class-id in the monitor with the class-id where the react is configured.
- Applying history or a react requires punting the record to the RP.

- Export on the "event transaction end" is used to export the records when the connection or transaction is terminated. In this case, the records are not exported based on timeout. Exporting on the event transaction end should be used when detailed connection/transaction granularity is required, and has the following advantages:
  - Sends the record close to the time that it has ended.
  - Exports only one record on true termination.
  - Conserves memory in the cache and reduces the load on the Management and Reporting system.
  - Enables exporting multiple transactions of the same flow. (This requires a protocol pack that supports multi-transaction.)
- Export spreading—In a case of synchronized cache, all network devices export records from the monitor cache at the same time. If multiple network devices are configured with the same monitor interval and synchronized cache, the collector may receive all records from all devices at the same time, which can impact the collector performance. The export-spreading feature spreads out the export over a time interval, which is automatically set by MMA or specified by the user.

NetFlow/IPFIX Flow Record

The flow record defines the record fields. With each Cisco IOS release, the Cisco AVC solution supports a more extensive set of metrics.

The sections that follow list commonly used AVC-specific fields organized by functional groups. These sections do not provide detailed command reference information, but highlight important usage guidelines.

In addition to the fields described below, a record can include any NetFlow field supported by the platform.

A detailed description of NetFlow fields appears in the *Cisco IOS Flexible NetFlow Command Reference*.

Note On Cisco IOS XE platforms, the record size is limited to 40 fields (key and non-key fields or match and collect fields).

L3/L4 Fields

The following are L3/L4 fields commonly used by the Cisco AVC solution.

```
[collect | match] connection [client|server] [ipv4|ipv6] address
[collect | match] connection [client|server] transport port
[collect | match] [ipv4|ipv6] [source|destination] address
[collect | match] transport [source-port|destination-port]
[collect | match] [ipv4|ipv6] version
[collect | match] [ipv4|ipv6] protocol
[collect | match] routing vrf [input|output]
[collect | match] [ipv4|ipv6] dscp
[collect | match] ipv4 ttl
[collect | match] ipv6 hop-limit
collect            transport tcp option map
collect            transport tcp window-size [minimum|maximum|sum]
collect            transport tcp maximum-segment-size
```

Usage Guidelines

The client is determined according to the initiator of the connection.

The **client** and **server** fields are bi-directional. The **source** and **destination** fields are uni-directional.

L7 Fields

The following are L7 fields commonly used by the Cisco AVC solution.

```
[collect | match] application name [account-on-resolution]
collect application http url
collect application http uri statistics
collect application http host
collect application http user-agent
collect application http referer
collect application rtsp host-name
collect application smtp server
collect application smtp sender
collect application pop3 server
collect application nntp group-name
collect application sip source
collect application sip destination
```

Usage Guidelines

- The application ID is exported according to RFC-6759.
- Account-On-Resolution configures FNF to collect data in a temporary memory location until the record key fields are resolved. After resolution of the record key fields, FNF combines the temporary data collected with the standard FNF records. Use the **account-on-resolution** option when the field used as a key is not available at the time that FNF receives the first packet.

  The following limitations apply when using Account-On-Resolution:

  - Flows ended before resolution are not reported.
  - On Cisco IOS XE platforms, FNF packet/octet counters, timestamp, and TCP performance metrics are collected until resolution. All other field values are taken from the packet that provides resolution or the following packets.
- For information about extracted fields, including the formats in which they are exported, see: *Cisco Application Visibility and Control Field Definition Guide for Third-Party Customers*

Interfaces and Directions

The following are interface and direction fields commonly used by the Cisco AVC solution:

```
[collect | match] interface [input|output]
[collect | match] flow direction
collect connection initiator
```

Counters and Timers

The following are counter and timer fields commonly used by the Cisco AVC solution.

Note Two aliases provide backward compatibility for configurations created on earlier releases:

- **connection client bytes transport long** is an alias for **connection client bytes long**.
- **connection server bytes transport long** is an alias for **connection server bytes long**.

```
collect connection server counter bytes network long
collect connection server counter bytes transport long
collect connection server counter bytes long
collect connection server counter packets long

collect connection client counter bytes network long
collect connection client counter bytes transport long
collect connection client counter bytes long
collect connection client counter packets long

collect counter bytes rate
collect connection server counter responses
collect connection client counter packets retransmitted
collect connection transaction duration {sum, min, max}
collect connection transaction counter complete
collect connection new-connections
collect connection sum-duration
collect timestamp sys-uptime first
collect timestamp sys-uptime last
```

On Cisco IOS platforms:

```
collect counter packets long
collect counter bytes long
```

On Cisco IOS XE platforms:

```
collect counter packets [long]
collect counter bytes [long]
```

TCP Performance Metrics

The following are fields commonly used for TCP performance metrics by the Cisco AVC solution:

```
collect           connection delay network to-server         {sum, min, max}
collect           connection delay network to-client         {sum, min, max}
collect           connection delay network client-to-server  {sum, min, max}
collect           connection delay response to-server        {sum, min, max}
collect           connection delay response to-server histogram
                                                [bucket1 ... bucket7 | late]
collect           connection delay response client-to-server {sum, min, max}
collect           connection delay application               {sum, min, max}
```

Usage Guidelines

The following limitations apply to TCP performance metrics:

- All TCP performance metrics must observe bi-directional traffic.
- The policy-map must be applied in both directions.

Figure 4-1 provides an overview of network response time metrics.

Figure 4-1 Network Response Times

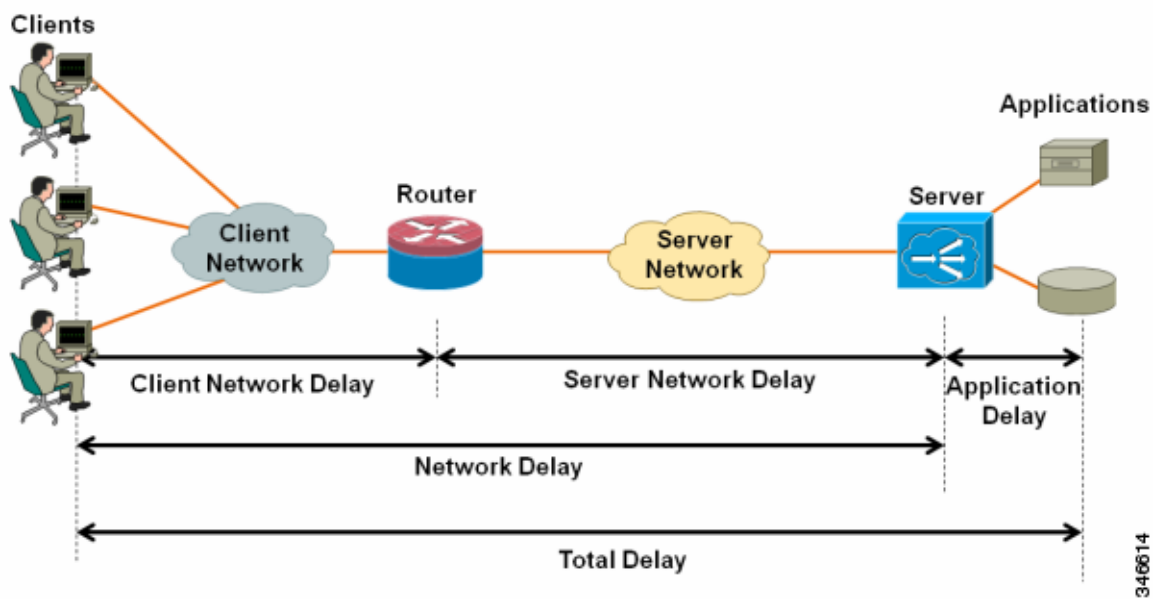

Figure 4-2 provides details of network response time metrics.

Figure 4-2 Network Response Time Metrics in Detail

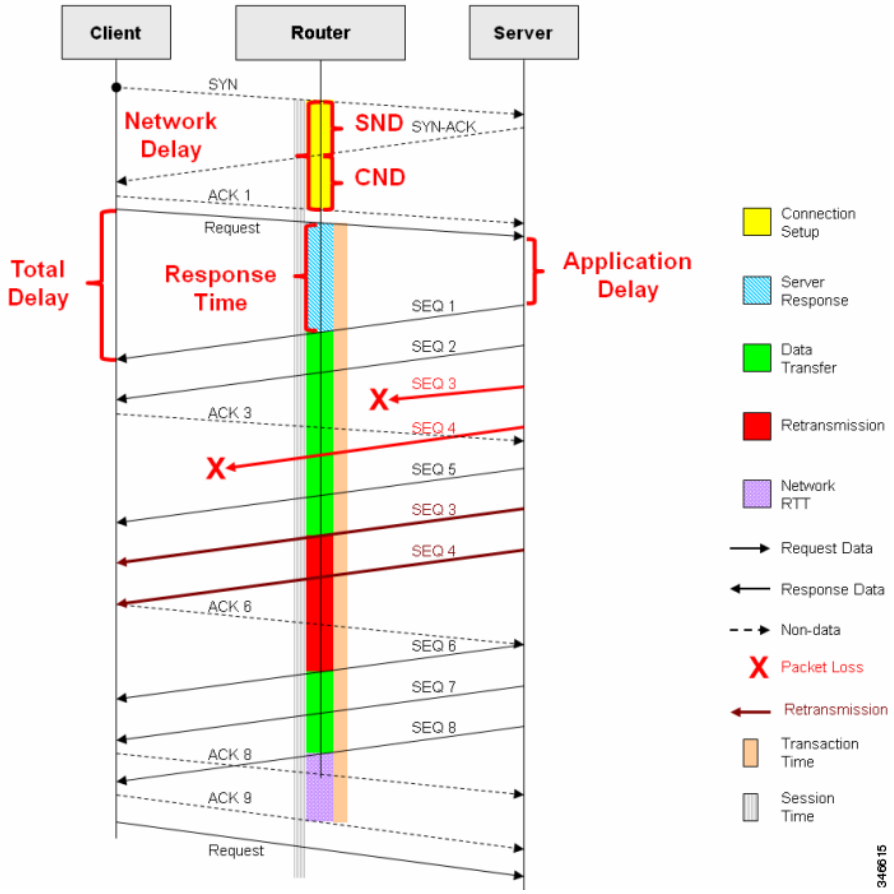

Media Performance Metrics

The following are fields commonly used for media performance metrics by the Cisco AVC solution:

```
[collect | match] match transport rtp ssrc
collect  transport rtp payload-type
collect  transport rtp jitter mean sum
collect  transport rtp jitter [minimum | maximum]
collect  transport packets lost counter
collect  transport packets expected counter
collect  transport packets lost counter
collect  transport packets lost rate
collect  transport event packet-loss counter
collect  counter packets dropped
collect  application media bytes counter
collect  application media bytes rate
collect  application media packets counter
collect  application media packets rate
collect  application media event
collect  monitor event
```

Usage Guidelines

Some of the media performance fields require punt to the route processor (RP). For more information, see *Cisco Application Visibility and Control Field Definition Guide for Third-Party Customers*.

L2 Information

The following are L2 fields commonly used by the Cisco AVC solution:

```
[collect | match] datalink [source-vlan-id | destination-vlan-id]
[collect | match] datalink mac [source | destination] address [input | output]
```

WAAS Interoperability

Cisco IOS Platforms	Cisco IOS XE Platforms
Not available	Available

The following are WAAS fields commonly used by the Cisco AVC solution:

```
[collect | match] services waas segment [account-on-resolution]
collect services waas passthrough-reason
```

Usage Guidelines

Account-On-Resolution configures FNF to collect data in a temporary memory location until the record key fields are resolved. After resolution of the record key fields, FNF combines the temporary data collected with the standard FNF records. Use this option (**account-on-resolution**) when the field used as a key is not available at the time that FNF receives the first packet.

The following limitations apply when using Account-On-Resolution:

- Flows ended before resolution are not reported.
- FNF packet/octet counters, timestamp and TCP performance metrics are collected until resolution. All other field values are taken from the packet that provides resolution or the following packets.

Classification

The following are classification fields commonly used by the Cisco AVC solution:

```
[collect | match]  policy performance-monitor classification hierarchy
```

Usage Guidelines

Use this field to report the matched class for the performance-monitor policy-map.

NetFlow/IPFIX Option Templates

NetFlow option templates map IDs to string names and descriptions:

```
flow exporter my-exporter
  export-protocol ipfix
  template data timeout <timeout>
```

```
  option interface-table timeout <timeout>
  option vrf-table timeout <timeout>
  option sampler-table timeout <timeout>
  option application-table timeout <timeout>
  option application-attributes timeout <timeout>
  option sub-application-table timeout <timeout>
  option c3pl-class-table timeout <timeout>
  option c3pl-policy-table timeout <timeout>
```

NetFlow/IPFIX Show commands

Use the following commands to show NetFlow/IPFIX information:

```
show flow monitor type performance-monitor [<name> [cache [raw]]]
show flow record type performance-monitor
show policy-map type performance-monitor [<name> | interface]
```

Customizing NBAR Attributes

Use the following commands to customize the NBAR attributes:

```
[no] ip nbar attribute-map <attribute-map-name>
   attribute category <category>
   attribute sub-category <sub-category>
   attribute application-group <application-group>
   attribute tunnel <tunnel-info>
   attribute encrypted <encrypted-info>
   attribute p2p-technology <p2p-technology-info>
[no] ip nbar attribute-set <protocol-name> <attribute-map-name>
```

Note These commands support all attributes defined by the NBAR2 Protocol Pack, including custom-category, custom-sub-category, and custom-group available in Protocol Pack 3.1 and later.

Customizing Attribute Values

Cisco IOS Platforms	Cisco IOS XE Platforms
Added in IOS 15.4(1)T	Added in IOS XE 3.11

Background

Attribute maps enable users to map various attribute values to protocols, changing the built-in grouping of protocols. The “custom attributes value” feature enables users to add new values to existing attributes.

For example, when using custom protocols to define enterprise specific protocols, it can be useful to classify the custom protocols as a new group (example: my-db-protocols-group). Beginning in the current release, new values can be defined for:

- category
- sub-category
- application-group

Customized attributes can be used for QoS matching, and the customized values appear in AVC reports.

Future Protocol Pack versions may enable defining additional attributes. For information about viewing which attributes can be customized and how many new groups can be defined, see Additional Usage Guidelines, page 4-12.

Basic Usage

CLI

```
[no] ip nbar attribute <attribute name> custom  <user-defined value>  [<user-defined help
string>]
```

Backward Compatibility

Previous releases of AVC included the following pre-defined attribute values, which could not be user-customized:

- For the category attribute: **custom-category**
- For the sub-category attribute: **custom-sub-category**
- For the application-group attribute: **custom-application-group**

To provide backward compatibility with existing configurations, the current release supports configurations that were created for earlier releases and that include one or more of these attributes.

Examples—Defining Values

The following examples define custom values for the category and sub-category attributes, and provide the optional explanatory help string:

```
ip nbar attribute category custom dc_backup_category "Data center backup traffic"
ip nbar attribute sub-category custom hr_sub_category "HR custom applications traffic"
ip nbar attribute application-group custom Home_grown_finance_group "our finance tools
network traffic"
```

Example—Removing Custom Values

The following example removes the custom value (“XYZ-app-group”) that had been assigned for the application-group attribute:

```
no ip nbar attribute application-group custom XYZ-app-group
```

Additional Usage Guidelines

Help

The following command provides help, indicating which attributes can have custom values.

```
ip nbar attribute ?
```

Displaying Customizable Attributes and Custom Values

The following command indicates which attributes can be defined with custom values (depends on the Protocol Pack version installed on the device), and displays the currently defined custom values.

```
show ip nbar attribute-custom
```

Customizing NBAR Protocols

Use the following commands to customize NBAR protocols and assign a protocol ID. A protocol can be matched based on HTTP URL/Host or other parameters:

```
ip nbar custom <protocol-name> [http {[url <urlregexp>]  [host <hostregexp>]}] [offset
[format value]] [variable field-name field-length] [source | destination] [tcp | udp ]
[range start end | port-number ] [id <id>]
```

Packet Capture Configuration

Cisco IOS Platforms	Cisco IOS XE Platforms
Not available	Available

Use the following commands to enable packet capture:

```
policy-map type packet-services <policy-name>
    class <class-name>
        capture limit packet-per-sec <pps> allow-nth-pak <np> duration <duration>
                    packets <packets>  packet-length <len>
        buffer size <size> type <type>

interface <interface-name>
   service-policy type packet-services <policy-name> [input|output]
```

QoS Metrics: Cisco IOS Platforms

This section applies to Cisco IOS platforms. (For information about QoS Metrics configuration for Cisco IOS XE platforms, see QoS Metrics: Cisco IOS XE Platforms, page 4-18.)

This section describes how to configure a performance monitor to include Quality of Service (QoS) metrics.

Background—QoS

QoS configuration is based on **class maps** and **policy maps**. Class maps categorize traffic; policy maps determine how to handle the traffic. Based on the policy identified for each packet, the packet is placed into a specific **QoS queue**, which determines the priority and pattern of transmission. Each queue is identified by a Queue ID field.

For additional information about QoS, see: http://www.cisco.com/go/qos

Exported Metrics

AVC enables configuration of QoS Packet Drop and QoS Class Hierarchy monitors on an interface, using one or more of the following QoS metrics, which can be included in exported performance monitor records:

- Queue ID—Identifies a QoS queue.
- Queue Packet Drops—Packets dropped (on the monitored interface) per QoS queue, due to a QoS policy that limits resources available to a specific type of traffic.
- Class Hierarchy—Class hierarchy of the reported flow. The class hierarchy is determined by the QoS policy map and determines the traffic priority.

QoS Packet Drop Monitor Output in Exported Record

When a QoS Packet Drop monitor is configured, the performance monitor record includes packet drop data per QoS queue in the following format:

Queue id	Queue packet drops
1	100
2	20

QoS Class Hierarchy Information Included in Exported Record

QoS class hierarchy information is exported using the following performance monitor fields:

- Hierarchy policy for each flow (defined by the policy map)
- Queue ID for each flow

This section provides an example of a QoS policy map configuration, followed by the information provided in a performance monitor record for three flows governed by this configuration.

The example includes two levels of policy map hierarchy. In the example, the `service-policy P11` statement in **bold** type creates a hierarchy with the P11 policy map as a child of the P1 policy map.

Note QoS class hierarchy reporting supports a hierarchy of five levels.

Based on the configuration, the following applies to a packet with, for example, a DSCP value of “ef” in the IP header:

1. The C1 class definition includes the packet by the `match any` statement.
2. The C11 class definition includes the packet by the `match ip dscp ef` statement.
3. Because the packet is included in class C1, policy map P1 defines the policy for the packet with the `shaping average` statement.

4. Policy map P1 invokes policy map P11 for class C1 with the `service-policy P11` statement.
5. Because the packet is included in class C11, policy map P11 assigns the packet to a queue which has been allocated 10% of remaining bandwidth.

```
class-map match-all C1
    match any
class-map match-all C11
    match ip dscp ef
class-map match-all C12
    match ip dscp cs2
!
policy-map P11
    class C11
        bandwidth remaining percent 10
    class C12
        bandwidth remaining percent 70
    class class-default
        bandwidth remaining percent 20

policy-map P1
    class C1
        shaping average 16000000
        service-policy P11
```

Table 4-2 shows an example of the information provided in an FNF record for three flows governed by this configuration.

Table 4-2 QoS Class Hierarchy Information in the Flow Record

Flow	Hierarchy	Queue id
Flow 1	P1, C1, C11	1
Flow 2	P1, C1, C11	1
Flow 3	P1, C1, C12	2

In Table 4-2, policy and class information is shown using the true policy and class names, such as P1 and C1. However, the record exports policy and class names using numerical identifiers in place of policy and class names. The monitor periodically outputs a “policy option template” and a “class option template” indicating the policy names and class names that correspond to the numbers used in the exported records. These option templates are defined in the exporter configuration, using statements such as the following, which create the option templates and indicate the time interval at which the monitor outputs the option template information:

```
option c3pl-class-table timeout <timeout>
option c3pl-policy-table timeout <timeout>
```

Configuration

Configuring a QoS Packet Drop Monitor

A QoS Packet Drop monitor can only export the Queue ID and Queue Packet Drop fields. It cannot be combined with other monitors to export additional fields. At the given reporting interval, the monitor reports only on queues that have dropped packets (does not report value of 0).

Step 1: Create the QoS Packet Drop Monitor

Use the following performance monitor configuration to create a QoS Packet Drop monitor. The process specifies a flow record of type performance monitor named "qos-record" and attaches the record to a monitor of type performance monitor named "qos-monitor." In the steps that follow, the qos-monitor is attached to the desired policy map.

```
flow record type performance monitor qos-record
    match policy qos queue index
    collect policy qos queue drops
flow monitor type performance monitor qos-monitor
    exporter my-exporter
    record qos-record
    cache timeout synchronized 60
```

Step 2: Configure the QoS Policy

The following example shows configuration of a QoS policy map. It includes a hierarchy of three policies: avc, avc-parent, and avc-gparent. Note that avc-gparent includes avc-parent, and avc-parent includes avc.

```
policy-map avc
    class prec4
        bandwidth remaining ratio 3
    class class-default
        bandwidth remaining ratio 1
policy-map avc-parent
    class class-default
        shape average 10000000
        service-policy avc
policy-map avc-gparent
    class class-default
        shape average 100000000
        service-policy avc-parent
```

Step 3: Create the QoS Class Hierarchy Record

To correlate the queue drops collected from the QoS Drops monitor, create a flow record that includes the class hierarchy and Queue id and flow key fields. The data exported by this monitor indicates which flows are assigned to which QoS Queue Id.

The following example configuration creates a QoS class record. The process specifies a record of type performance monitor named "qos-class-record."

```
flow record type performance-monitor qos-class-record
    match connection client ipv4 (or ipv6) address
    match connection server ipv4 (or ipv6) address
    match connection server transport port
    collect policy qos class hierarchy
    collect policy qos queue id
```

Step 4: Create the QoS Class Hierarchy Monitor

Use the following performance monitor configuration to create a QoS Class Hierarchy monitor. The process specifies a monitor of type "class-hier-monitor." In the steps that follow, the monitor is attached to the desired interface.

```
flow monitor type performance-monitor class-hier-monitor
    exporter my-exporter
    record qos-class-record
    cache timeout synchronized 60
```

Step 5: Create the Performance Monitor Policy

Use the following configuration to create a policy-map that will collect both monitors.

```
policy-map type performance monitor pm-qos
    class http
        flow monitor qos-monitor
        flow monitor qos-class-record
```

Step 6: Attach the Performance Monitor and QoS Policy to an Interface

Use the following to attach the monitor to the desired interface. For <*interface*>, specify the interface type—for example: GigabitEthernet0/2/1

Specify the IP address of the interface in IPv4 or IPv6 format.

```
interface <interface>
    ip address <interface_IP_address>
    service-policy type performance monitor output pm-qos
    service-policy output avc-gparent
```

Verifying the QoS Packet Drop Monitor Configuration

This section provides commands that are useful for verifying or troubleshooting a QoS Packet Drop Monitor configuration.

Verifying that the Monitor is Allocated

Use the following command to verify that the QoS monitor exists:

```
show flow monitor type performance monitor
```

Use the following commands to verify additional monitor details:

```
show flow monitor type performance monitor qos-monitor
show flow monitor type performance monitor qos-class-monitor
```

Verifying QoS Queue IDs, Queue Drops, and Class Hierarchies

The following show command displays the record collected:

```
show performance monitor history interval all
```

QoS Metrics: Cisco IOS XE Platforms

This section applies to Cisco IOS XE platforms. (For information about QoS Metrics configuration for Cisco IOS platforms, see QoS Metrics: Cisco IOS Platforms, page 4-13.)

This section describes how to configure Flexible NetFlow (FNF) monitors to include Quality of Service (QoS) metrics.

Background—FNF and QoS

FNF Monitors

Flexible NetFlow (FNF) enables monitoring traffic on router interfaces. FNF monitors are configured for a specific interface to monitor the traffic on that interface. At defined intervals, the monitor sends collected traffic data to a "collector," which can be a component within the router or an external component.

Beginning with Cisco AVC for IOS XE release 3.9, FNF records include new fields for QoS metrics.

QoS

QoS configuration is based on **class maps** and **policy maps**. Class maps categorize traffic; policy maps determine how to handle the traffic. Based on the policy identified for each packet, the packet is placed into a specific **QoS queue**, which determines the priority and pattern of transmission. Each queue is identified by a Queue ID field.

For additional information about QoS, see: http://www.cisco.com/go/qos

Exported Metrics

AVC enables configuration of QoS Packet Drop and QoS Class Hierarchy monitors on an interface, using one or more of the following QoS metrics, which can be included in exported FNF records:

- Queue ID—Identifies a QoS queue.
- Queue Packet Drops—Packets dropped (on the monitored interface) per QoS queue, due to a QoS policy that limits resources available to a specific type of traffic.
- Class Hierarchy—Class hierarchy of the reported flow. The class hierarchy is determined by the QoS policy map and determines the traffic priority.

QoS Packet Drop Monitor Output in Exported Record

When a QoS Packet Drop monitor is configured, the FNF record includes packet drop data per QoS queue in the following format:

Queue id	Queue packet drops
1	100
2	20

QoS Class Hierarchy Information Included in Exported Record

QoS class hierarchy information is exported using the following FNF fields:

- Hierarchy policy for each flow (defined by the policy map)
- Queue ID for each flow

This section provides an example of a QoS policy map configuration, followed by the information provided in an FNF record for three flows governed by this configuration.

The example includes two levels of policy map hierarchy. In the example, the `service-policy P11` statement in **bold** type creates a hierarchy with the P11 policy map as a child of the P1 policy map.

Note QoS class hierarchy reporting supports a hierarchy of five levels.

Based on the configuration, the following applies to a packet with, for example, a DSCP value of "ef" in the IP header:

1. The C1 class definition includes the packet by the `match any` statement.
2. The C11 class definition includes the packet by the `match ip dscp ef` statement.
3. Because the packet is included in class C1, policy map P1 defines the policy for the packet with the `shaping average` statement.
4. Policy map P1 invokes policy map P11 for class C1 with the `service-policy P11` statement.
5. Because the packet is included in class C11, policy map P11 assigns the packet to a queue which has been allocated 10% of remaining bandwidth.

```
class-map match-all C1
    match any
class-map match-all C11
    match ip dscp ef
class-map match-all C12
    match ip dscp cs2
!
policy-map P11
    class C11
        bandwidth remaining percent 10
    class C12
        bandwidth remaining percent 70
    class class-default
        bandwidth remaining percent 20

policy-map P1
    class C1
        shaping average 16000000
        service-policy P11
```

Table 4-3 shows an example of the information provided in an FNF record for three flows governed by this configuration.

Table 4-3 QoS Class Hierarchy Information in the FNF record

Flow	Hierarchy	Queue id
Flow 1	P1, C1, C11	1
Flow 2	P1, C1, C11	1
Flow 3	P1, C1, C12	2

In Table 4-3, policy and class information is shown using the true policy and class names, such as P1 and C1. However, the FNF record exports policy and class names using numerical identifiers in place of policy and class names. The monitor periodically outputs a "policy option template" and a "class option template" indicating the policy names and class names that correspond to the numbers used in the exported FNF records. These option templates are defined in the exporter configuration, using statements such as the following, which create the option templates and indicate the time interval at which the monitor outputs the option template information:

```
option c3pl-class-table timeout <timeout>
option c3pl-policy-table timeout <timeout>
```

Configuration

Enabling QoS Metric Collection

Enabling

To enable the QoS metrics collection feature for the platform, enter global configuration mode using `configure terminal`, then use the following QoS configuration command. The command causes QoS to begin collecting QoS metrics for FNF.

Note Enabling QoS metrics collection requires resetting all performance monitors on the device.

```
platform qos performance-monitor
```

Verifying

To verify that QoS metrics collection is enabled, use the following command:

```
show platform hardware qfp active feature qos config global
```

The following is an example of the output of the command:

```
Marker statistics are: disabled
Match per-filter statistics are: disabled
Match per-ace statistics are: disabled
Performance-Monitor statistics are: enabled
```

Configuring a QoS Packet Drop Monitor

A QoS Packet Drop monitor can only export the Queue ID and Queue Packet Drop fields. It cannot be combined with other monitors to export additional fields. At the given reporting interval, the monitor reports only on queues that have dropped packets (does not report value of 0).

Step 1: Create the QoS Packet Drop FNF Monitor

Use the following FNF configuration to create a QoS Packet Drop monitor. The process specifies a flow record of type "qos-record" and attaches the record to a monitor of type "qos-monitor." In the steps that follow, the qos-monitor is attached to the desired interface.

Note Ensure that QoS metrics collection is enabled. See Enabling QoS Metric Collection, page 4-20.

```
flow record qos-record
    match policy qos queue index
    collect policy qos queue drops
flow monitor qos-monitor
    exporter my-exporter
    record qos-record
```

Step 2: Configure the QoS Policy

The following example shows configuration of a QoS policy map. It includes a hierarchy of three policies: avc, avc-parent, and avc-gparent. Note that avc-gparent includes avc-parent, and avc-parent includes avc.

```
policy-map avc
    class prec4
        bandwidth remaining ratio 3
    class class-default
        bandwidth remaining ratio 1
policy-map avc-parent
    class class-default
        shape average 10000000
        service-policy avc
policy-map avc-gparent
    class class-default
        shape average 100000000
        service-policy avc-parent
```

Step 3: Attach the FNF Monitor and QoS Policy to an Interface

Use the following to attach the monitor to the desired interface. For <*interface*>, specify the interface type—for example: GigabitEthernet0/2/1

Specify the IP address of the interface in IPv4 or IPv6 format.

```
interface <interface>
    ip address <interface_IP_address>
    ip flow monitor qos-monitor output
    service-policy output avc-gparent
```

Verifying the QoS Packet Drop Monitor Configuration

This section provides commands that are useful for verifying or troubleshooting a QoS Packet Drop Monitor configuration.

Verifying that the Monitor is Allocated

Use the following command to verify that the QoS monitor exists:

```
show flow monitor
```

Use the following commands to verify additional monitor details:

```
show flow monitor qos-monitor
show flow monitor qos-monitor cache
show flow monitor qos-monitor statistics
show platform hardware qfp active feature fnf client flowdef name qos-record
show platform hardware qfp active feature fnf client monitor name qos-monitor
```

Verifying QoS queues and Class-Hierarchies

The following **show** commands display the statistics that QoS has collected. "gigX/X/X" refers to the interface for which the monitor has been configured.

```
show policy-map int gigX/X/X
show platform hardware qfp active feature qos queue output all
```

Verifying FNF-QOS FIA Activation

Use the following **show** command to verify that the FNF-QoS FIA (feature activation array) is enabled on the interface (GigabitEthernet0/2/1 in this example):

```
show platform hardware qfp active interface if-name GigabitEthernet0/2/1
```

Verifying the FNF Monitor and Record

Use the following **debug** commands to verify that the FNF monitor and record have been created:

```
debug platform software flow flow-def errors
debug platform software flow monitor errors
debug platform software flow interface errors


debug platform hardware qfp active feature fnf server trace
debug platform hardware qfp active feature fnf server info
debug platform hardware qfp active feature fnf server error
```

Configuring a QoS Class Hierarchy Monitor

In contrast to the QoS Packet Drop monitor, a QoS Class Hierarchy monitor can be combined with another monitor to export additional metrics.

Step 1: Create the QoS Class Record

The following example configuration creates a QoS class record. The process specifies a record of type "qos-class-record." The example specifies "ipv4 source" and "ipv4 destination" addresses, but you can configure the record to match according to other criteria.

Note Ensure that QoS metrics collection is enabled. See Enabling QoS Metric Collection, page 4-20.

```
flow record qos-class-record
    match ipv4 source address
    match ipv4 destination address
    collect counter bytes
    collect counter packets
    collect policy qos classification hierarchy
    collect policy qos queue index
```

Step 2: Create the QoS Class Hierarchy Monitor

Use the following FNF configuration to create a QoS Class Hierarchy monitor. The process specifies a monitor of type "class-hier-monitor." In the steps that follow, the monitor is attached to the desired interface.

```
flow monitor class-hier-monitor
    exporter my-exporter
    record qos-class-record
```

Step 3: Attach the QoS Class Hierarchy Monitor to an Interface

Use the following to attach the monitor to the desired interface. For <*interface*>, specify the interface type—for example: GigabitEthernet0/2/1

Specify the IP address of the interface in IPv4 or IPv6 format.

Note Attaching the service-policy to the interface, as indicated by the "service-policy" statement below, is a required step.

```
interface <interface>
    ip address <interface_IP_address>
    ip flow monitor class-hier-monitor output
    service-policy output avc-gparent
```

Verifying the QoS Class Hierarchy Monitor Configuration

This section provides commands that are useful for verifying or troubleshooting a QoS Class Hierarchy Monitor configuration.

Verifying that the Monitor is Allocated

Use the following command to verify that the QoS monitor exists:

```
show flow monitor
```

Use the following commands to verify additional details:

```
show flow monitor class-hier-monitor
show flow monitor class-hier-monitor cache
show flow monitor class-hier-monitor statistics

show platform hardware qfp active feature fnf client flowdef name qos-class-record
show platform hardware qfp active feature fnf client monitor name qos-monitor
```

Verifying FNF-QOS FIA Activation

In the following feature invocation array (FIA) verification example, the interface is GigabitEthernet0/2/1.

```
show platform hardware qfp active interface if-name GigabitEthernet0/2/1
```

Verifying the FNF Monitor and Record

Use the following **debug** commands to verify that the FNF monitor and record have been created:

```
debug platform software flow flow-def errors
debug platform software flow monitor errors
debug platform software flow interface errors

debug platform hardware qfp active feature fnf server trace
debug platform hardware qfp active feature fnf server info
debug platform hardware qfp active feature fnf server error
```

Connection/Transaction Metrics

Cisco IOS Platforms	Cisco IOS XE Platforms
Not available	Added in release 3.9S

Flexible NetFlow (FNF) monitors can report on individual transactions within a flow. This enables greater resolution for traffic metrics. This section describes how to configure connection and transaction metrics, including **transaction-id** and **connection id**, for FNF monitors. The connection/transaction monitoring feature is referred to as "Multi-transaction."

Note The Multi-transaction feature requires an NBAR protocol pack that supports the feature. The protocol pack provided with Cisco AVC for IOS XE release 3.9S and later protocol packs support this feature.

Introduction

Flexible NetFlow (FNF) monitors typically report traffic metrics per flow. (A flow is defined as a connection between a specific source address/port and destination address/port.) A single flow can include multiple HTTP transactions. Enabling the Multi Transaction feature for a monitor enables reporting metrics for each transaction individually.

You can configure the FNF record to identify the flow or the flow+transaction, using one of the following two metrics:

- connection id—A 4-byte metric identifying the flow.
- transaction-id—An 8-byte metric composed of two parts:
  - MSB—Identifies the flow and is equivalent to the connection id metric.
  - LSB—Identifies the transaction. The value is a sequential index of the transaction, beginning with 0.

Configuration

The following subsections describe the following for the Multi-transaction feature:

Requirements

The following requirements apply when using the Multi-transaction feature:

- The record configuration must use **match**, not **collect**.
- Specify only "connection id" or "transaction-id," but not both.
- Include "application name" in the record.
- Include "cache timeout event transaction-end" which specifies that the record is transmitted immediately and not stored in the monitor cache.

Configuring Exporter, Record, and Monitor in Performance Monitor Mode

Flexible Netflow (FNF) performance monitor (perf-monitor) mode enables configuring monitors with advanced filtering options that filter data before reporting it. Options for configuring filtering include IP access list, policy-map, and so on.

The following perf-monitor example configures a monitor and specifies the **transaction-id** metric for the FNF record, as shown in **bold**. Alternatively, you can specify the **connection id** metric.

Note See Configuring Exporter, Record, and Monitor in Performance Monitor Mode, page 4-25 for additional configuration information.

```
ip access-list extended mt_perf_acl
    permit ip any any

class-map match-all mt_perf_class
    match access-group name mt_perf_acl
    match protocol http

flow exporter mt_perf_exporter
    destination 64.128.128.128
    transport udp 2055

flow record type performance-monitor mt_perf_record
    match connection transaction-id
    collect counter packets
    collect application name
    collect application http url

flow monitor type performance-monitor mt_perf_monitor
    record mt_perf_record
    exporter mt_perf_exporter
    cache type normal
    cache timeout event transaction-end

policy-map type performance-monitor mt_perf_policy
    parameter default account-on-resolution
    class mt_perf_class
    flow monitor mt_perf_monitor

interface GigabitEthernet0/0/2
    service-policy type performance-monitor input mt_perf_policy
```

Verifying and Troubleshooting the Configuration

This section describes commands useful for verification and troubleshooting the FNF configuration. There are subsections for:

- Native or Performance Monitor Mode, page 4-26
- Native FNF Mode, page 4-26
- Performance Monitor Mode, page 4-26

Note For information about the **show** commands in the sections below, see the FNF command reference guide: http://www.cisco.com/en/US/docs/ios-xml/ios/fnetflow/command/fnf-cr-book.html

Native or Performance Monitor Mode

Verifying Multi-transaction Status

Display the Multi-transaction status:

```
show plat soft nbar statistics | inc is_multi_trs_enable
```

If Multi-transaction is enabled, the value is: `is_multi_trs_enable==1`

Native FNF Mode

Validating the Configuration

Use the following **show** commands to validate the configuration.

```
show flow exporter <exporter_name> templates
show flow monitor <monitor_name>
show platform hardware qfp active feature fnf client flowdef name <record_name>
show platform hardware qfp active feature fnf client monitor name <monitor_name>
```

Viewing Collected FNF Data and Statistics

Use the following **show** commands to view the collected FNF data and statistics.

```
show flow monitor <monitor_name> cache
show flow monitor <monitor_name> statistics
show flow exporter <exporter_name> statistics
show platform hardware qfp active feature fnf datapath aor
```

Performance Monitor Mode

Validating the Configuration

Use the following **show** commands to validate the configuration.

```
show flow exporter <exporter_name> templates
show flow record type performance-monitor <record_name>
show platform hardware qfp active feature fnf client monitor name <monitor_name>
```

Viewing Collected FNF Data and Statistics

Use the following **show** commands to view the FNF collected data and statistics.

```
show performance monitor cache monitor <monitor_name> detail
show flow exporter <exporter_name> statistics
show platform hardware qfp active feature fnf datapath aor
```

Easy Performance Monitor

Cisco IOS Platforms	Cisco IOS XE Platforms
Added in release 15.4(1)T	Added in release 3.10S

Overview

The Easy Performance Monitor ("Easy perf-mon" or "ezPM") feature provides an "express" method of provisioning monitors. This new mechanism adds functionality and does not affect the existing methods for provisioning monitors.

Easy perf-mon does not provide the full flexibility of the traditional perf-mon configuration model. Easy perf-mon provides "profiles" that represent typical deployment scenarios. After selecting a profile and specifying a small number of parameters, Easy perf-mon provides the remaining provisioning details.

For additional information about configuring Easy perf-mon, see:
Easy Performance Monitor

Application Experience Profile

In the current release, Easy perf-mon includes one profile, called "Application Experience," and five different traffic monitors, described in Table 4-4. Future releases will provide additional options.

Table 4-4 Application Experience Traffic Monitors

	Monitor Name	Default Traffic Classification
1	Application-Response-Time (ART)	All TCP
2	URL	HTTP applications
3	Media	RTP applications over UDP
4	Conversation-Traffic-Stats	Remaining traffic not matching other classifications
5	Application-Traffic-Stats	DNS and DHT

Users can override a small set of parameters in each of the traffic monitors, as described in Table 4-5. For an example of how to configure parameters in the Application Experience profile, see Easy Perf-Mon Configuration Example 2: Application Experience Profile, page 4-30.

Table 4-5 Application Experience Traffic Monitors: Parameters

		Configurable Parameters			
	Monitor Name	IPv4 / IPv6	In / Out	Traffic Class	Cache Size
1	Application-Response-Time (ART)	Y	N	Class and Application only	Y
2	URL	Y	N	Class and Application only	Y
3	Media	Y	Y	Class and Application only	Y

	Monitor Name	Configurable Parameters			
		IPv4 / IPv6	In / Out	Traffic Class	Cache Size
4	Conversation-Traffic-Stats	Y	N	N	Y
5	Application-Traffic-Stats	N	N	N	Y

The Application Experience profile implements the improved data exporting model, which is optimized for maximum performance, exporting the maximum possible amount of available information for monitored traffic. Based on the requirements of the reports that have been defined:

- For each type of traffic, the exported record contains all of the collected data required for the defined reports, with the required granularity.
- Exported records do not contain unnecessary data, such as data redundant with previously exported records or data that is not required for the defined reports.
- Exported records include server information.

Figure 4-3 illustrates how the "Application Experience" profile exports different types of traffic statistics.

Figure 4-3 Export Model—Application Experience Profile

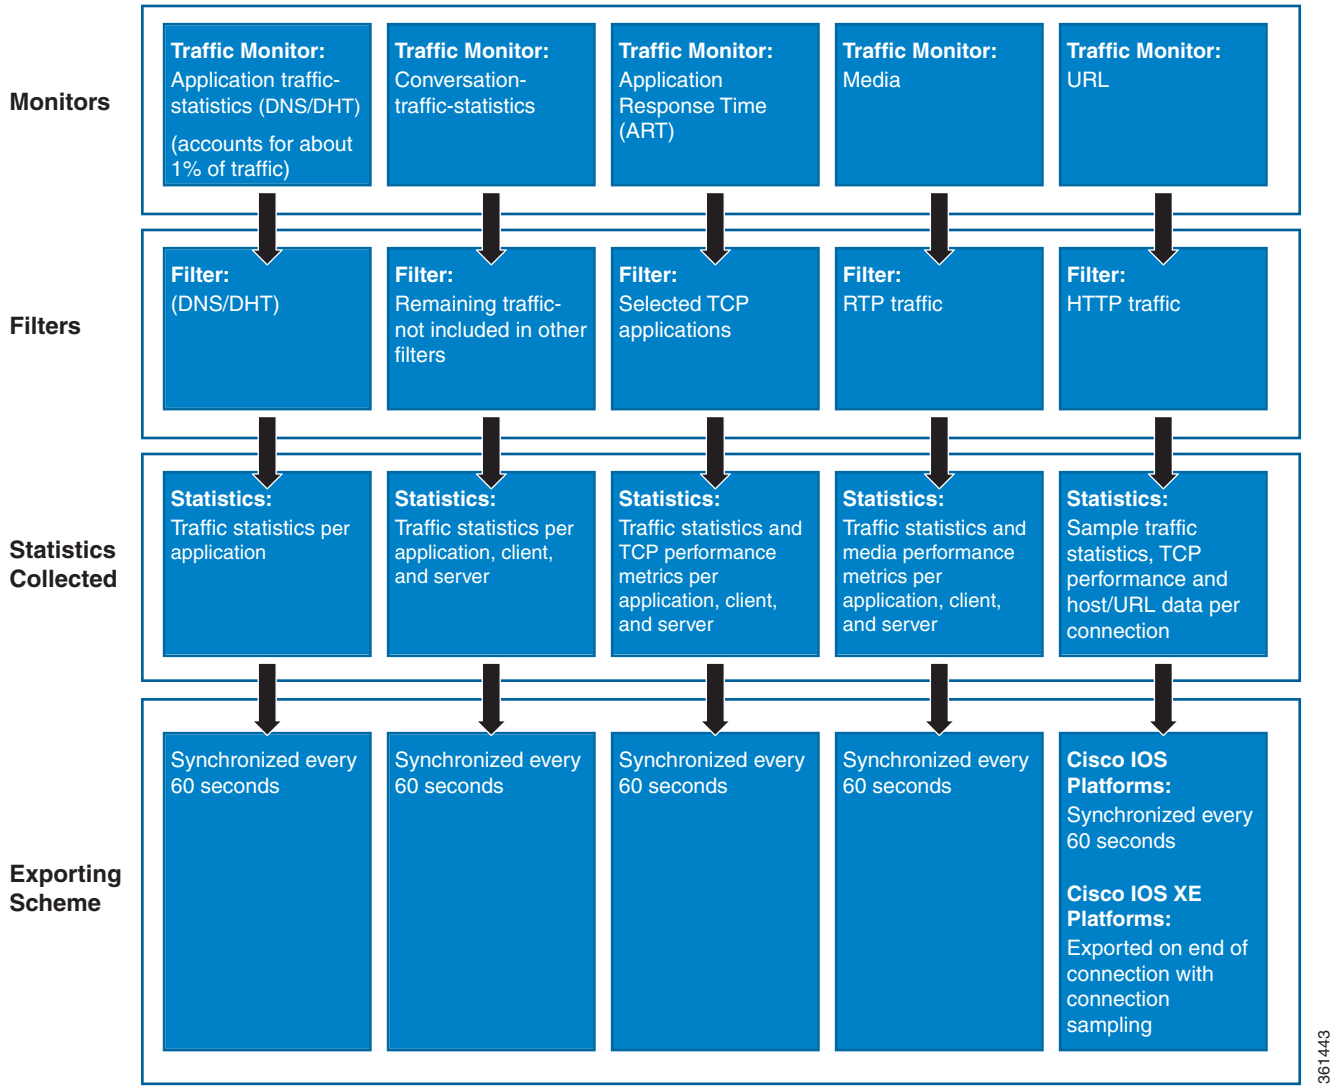

Configuration Examples

Easy Perf-Mon Configuration Example 1

The following Easy perf-mon configuration example activates all traffic monitors in the profile and attaches the policy-maps, both ingress and egress, to the GigabitEthernet0/0/1 interface:

```
!
! Easy performance monitor context
! --------------------------------
!
performance monitor context my-avc profile application-experience
   exporter destination 1.2.3.4 source GigabitEthernet0/0/1 port 4739
   traffic-monitor all
!
!
! Interface attachments
! ---------------------
interface GigabitEthernet0/0/1
  performance monitor context my-avc
```

Easy Perf-Mon Configuration Example 2: Application Experience Profile

The following Easy perf-mon “Application Experience” profile configuration example activates three traffic monitors, and specifies monitoring only IPv4 traffic. The context is then attached to two interfaces:

```
!
! Easy performance monitor context
! --------------------------------
!
performance monitor context my-visibility profile application-experience
   exporter destination 1.2.3.4 source GigabitEthernet0/0/1 port 4739

   traffic-monitor application-response-time ipv4
   traffic-monitor conversation-traffic-stats ipv4
   traffic-monitor media ipv4
!
! Interface attachments
! ---------------------
interface GigabitEthernet0/0/1
  performance monitor context my-visibility
interface GigabitEthernet0/0/2
  performance monitor context my-visibility
```

CLI Field Aliases

Cisco IOS Platforms	Cisco IOS XE Platforms
Added in release 15.4(1)T	Added in release 3.10S

Aliases provide a mechanism for simplifying configuration statements. The **all** alias refers to the set of all fields possible for a given statement. For example, "`collect connection delay` **`all`**" configures all fields that are possible to configure by the "`collect connection delay`" statement.

The following are examples:

```
collect connection delay all
collect connection transaction all
collect connection client all
collect connection server all
collect connection delay response to-server histogram all
```

Caution When using aliases, see Removing Aliases before Downgrading from Cisco IOS 15.4(1)T / Cisco IOS XE 3.10 or Later, page 6-2 before downgrading from Cisco IOS release 15.4(1)T or later, or from Cisco IOS XE release 3.10S or later.

Identifying the Monitored Interface

Cisco IOS Platforms	Cisco IOS XE Platforms
—	Added in release 3.11S

The "observation point id" metric identifies a monitored interface for traffic in both directions (ingress and egress). A single flow definition using this metric can be used in place of **match interface input** and **match interface output**, making configuration more compact and enabling a single record collected on an interface to include metrics for traffic in both directions.

Usage Guidelines

Configure the monitor on both the ingress and egress directions.

Example

In the following example configuration, a single monitor identifies the interface for traffic in both directions:

```
flow record my-application-record
    match application name account-on-resolution
    match flow observation point
    match flow direction
    collect counter packets
    collect counter bytes
```

Configuration Examples

This section contains AVC configuration examples. These examples provide a general view of a variety of configuration scenarios, combining multiple AVC features. Configuration is flexible and supports different types of record configurations.

Conversation Based Records—Omitting the Source Port

The monitor configured in the following examples sends traffic reports based on conversation aggregation. For performance and scale reasons, it is preferable to send TCP performance metrics only for traffic that requires TCP performance measurements. It is recommended to configure two similar monitors:

- One monitor includes the required TCP performance metrics. In place of the line shown in **bold** in the example below (collect <any TCP performance metric>), include a line for each TCP metric for the monitor to collect.
- One monitor does not include TCP performance metrics.

The configuration is for IPv4 traffic. Similar monitors should be configured for IPv6.

Example 1: For Cisco IOS Platforms

```
flow record type performance-monitor conversation-record
    match connection client ipv4 (or ipv6) address
    match connection server ipv4 (or ipv6) address
    match connection server transport port
    match ipv4 (or ipv6) protocol
    match application name account-on-resolution
    collect interface input
    collect interface output
    collect connection server counter bytes long
    collect connection client counter bytes long
    collect connection server counter packets long
    collect connection client counter packets long
    collect connection sum-duration
    collect connection new-connections
    collect policy qos class hierarchy
    collect policy qos queue id
    collect <any TCP performance metric>


flow monitor type performance-monitor conversation-monitor
    record conversation-record
    exporter my-exporter
    history size 0
    cache type synchronized
    cache timeout synchronized 60
    cache entries <cache size>

flow record qos-record
    match policy qos queue index
    collect policy qos queue drops
flow monitor qos-monitor
    exporter my-exporter
    record qos-record
```

Example 2: For Cisco IOS XE Platforms

```
flow record type performance-monitor conversation-record
    match services waas segment account-on-resolution
    match connection client ipv4 (or ipv6) address
    match connection server ipv4 (or ipv6) address
    match connection server transport port
    match ipv4 (or ipv6) protocol
    match application name account-on-resolution
    collect interface input
    collect interface output
    collect connection server counter bytes long
    collect connection client counter bytes long
    collect connection server counter packets long
    collect connection client counter packets long
    collect connection sum-duration
    collect connection new-connections
    collect policy qos class hierarchy
    collect policy qos queue id
    collect <any TCP performance metric>


flow monitor type performance-monitor conversation-monitor
    record conversation-record
    exporter my-exporter
    history size 0
    cache type synchronized
    cache timeout synchronized 60
    cache entries <cache size>
```

HTTP URL

The monitor configured in the following example sends the HTTP host and URL. If the URL is not required, the host can be sent as part of the conversation record (see Conversation Based Records—Omitting the Source Port, page 4-32).

```
flow record type performance-monitor url-record
    match transaction-id
    collect application name
    collect connection client ipv4 (or ipv6) address
    collect routing vrf input
    collect application http url
    collect application http host
    <other metrics could be added here if needed.
     For example bytes/packets to calculate BW per URL
     Or performance metrics per URL>

flow monitor type url-monitor
    record url-record
    exporter my-exporter
    history size 0
    cache type normal
    cache timeout event transaction-end
    cache entries <cache size>
```

HTTP URI

The **uri statistics** command enables exporting the first level of a parsed URI address. The command exports the value in the URI statistics field, which contains the depth 1 URI value, followed by a URI hit count value.

Note **Cisco IOS XE Platforms**: The URI hit count value is always 1 because the URI statistics field can only be configured per connection or transaction.

If no backslash exists at all after the URL, a zero length field is exported.

If the depth 1 value of the parsed URI exceeds a maximum number of characters, the value is truncated to the maximum length.

Note **Cisco IOS XE Platforms**: The **uri statistics** command must be configured with either the **connection id** or **transaction-id** commands.

Configuration Example

```
flow record er_uri_stat_record_1
    match connection transaction-id
    collect application name
    collect counter packets
    collect application http uri statistics
```

Example of Exported Value—Typical Address

Address: http://usr:pwd@www.test.com:81/dir/dir.2/index.htm?q1=0&&test1&test2=value#top

The **uri statistics** command exports: **/dir:1**

- **/dir** is the URI depth 1 level value.
- The "**:**" indicates a null character, followed by a URI hit count value of **1**.

Example of Exported Value—No Backslash after URL

Address: http://usr:pwd@www.test.com

The **uri statistics** command exports a zero length field.

Application Traffic Statistics

The monitor configured in the following example collects application traffic statistics:

```
flow record type performance-monitor application-traffic-stats
    match ipv4 protocol
    match application name account-on-resolution
    match ipv4 version
    match flow direction
    collect connection initiator
    collect counter packets
    collect counter bytes long
    collect connection new-connections
    collect connection sum-duration

flow monitor type application-traffic-stats
    record application-traffic-stats
```

```
    exporter my-exporter
    history size 0
    cache type synchronized
    cache timeout synchronized 60
    cache entries <cache size>
```

Media RTP Report

The monitor configured in the following example reports on media traffic:

```
flow record type performance-monitor media-record
    match ipv4(or ipv6) protocol
    match ipv4(or ipv6) source address
    match ipv4(or ipv6) destination address
    match transport source-port
    match transport destination-port
    match transport rtp ssrc
    match routing vrf input
    collect transport rtp payload-type
    collect application name
    collect counter packets long
    collect counter bytes long
    collect transport rtp jitter mean sum
    collect transport rtp payload-type
    collect <other media metrics>

flow monitor type media-monitor
    record media-record
    exporter my-exporter
    history size 10 // default history
    cache type synchronized
    cache timeout synchronized 60
    cache entries <cache size>
```

QoS Example 1: Control and Throttle Traffic

The following QoS configuration example illustrates how to control and throttle the peer-to-peer (P2P) traffic in the network to 1 megabit per second:

```
class-map match-all p2p-class-map
  match protocol attribute sub-category p2p-file-transfer

policy-map p2p-attribute-policy
  class p2p-class-map
    police 1000000
interface Gig0/0/3
    service-policy input p2p-attribute- policy
```

QoS Example 2: Assigning Priority and Allocating Bandwidth

The following QoS configuration example illustrates how to allocate available bandwidth on the eth0/0 interface to different types of traffic. The allocations are as follows:

- Business-critical Citrix application traffic for "access-group 101" users receives highest priority, with 50% of available bandwidth committed and traffic assigned to a priority queue. The `police` statement limits the bandwidth of business-critical traffic to 50% in the example.
- Web browsing receives a committed 30% of the remaining bandwidth after the business-critical traffic. This is a commitment of 15% of the total bandwidth available on the interface.
- Internal browsing, as defined by a specific domain (myserver.com in the example), receives a committed 60% of the browsing bandwidth.
- All remaining traffic uses the remaining 35% of the total bandwidth.

The policy statements commit minimum bandwidth in the percentages described for situations of congestion. When bandwidth is available, traffic can receive more than the "committed" amount. For example, if there is no business-critical traffic at a given time, more bandwidth is available to browsing and other traffic.

Figure 4-4 illustrates the priority and bandwidth allocation for each class. "Remaining traffic" refers to all traffic not specifically defined by the class mapping.

Figure 4-4 Bandwidth Allocation

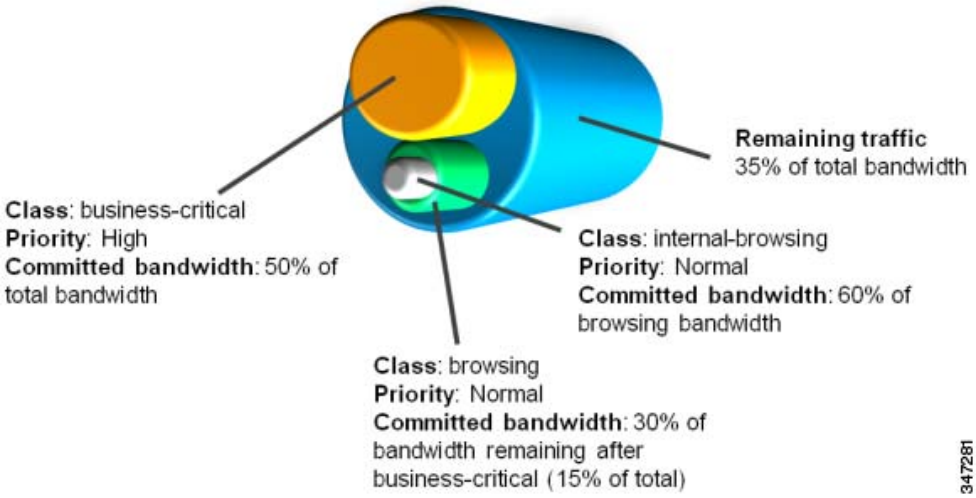

In class-map definition statements:

- **match-all** restricts the definition to traffic meeting all of the “match” conditions that follow. For example, the “business-critical” class only includes Citrix protocol traffic from IP addresses in “access-group 101.”
- **match-any** includes traffic meeting one or more of the “match” conditions that follow.

```
class-map match-all business-critical
    match protocol citrix
    match access-group 101
class-map match-any browsing
    match protocol attribute category browsing

class-map match-any internal-browsing
    match protocol http url “*myserver.com*”

policy-map internal-browsing-policy
    class internal-browsing
        bandwidth remaining percent 60

policy-map my-network-policy
    class business-critical
        priority
        police cir percent 50
    class browsing
        bandwidth remaining percent 30
        service-policy internal-browsing-policy

interface eth0/0
    service-policy output my-network-policy
```

Troubleshooting

Revised: November 21, 2013, OL-30581-01

This troubleshooting section includes the following topics:

Report Is Not Displayed Correctly

The following may be helpful for troubleshooting a report that is not displayed correctly:

- Verify that your flow exporter is configured with the correct destination IP.
- If you are using a VRF, ensure that it is added at the destination.

```
(config-flow-exporter)# destination 1.1.1.1 vrf myVrf
```

- Check whether samplers are configured correctly.
- Check the flow exporter statistics for errors.

```
# show flow exporter statistics
Flow Exporter my_exporter:
Packet send statistics (last cleared 4d00h ago):
  Successfully sent:         203808                (280136412 bytes)
Client send statistics:
  Client: Option options interface-table
Records added:           18528
   - sent:                18528
 Bytes added:             1852800
   - sent:                1852800
 Client: Option options vrf-id-name-table
   Records added:           3474
     - sent:                3474
 Bytes added:             125064
     - sent:                125064
```

```
     Client: Option options sampler-table
       Records added:           0
       Bytes added:             0
    Client: Option options application-name
       Records added:           1213584
```

- Check the cache output and verify that the specific monitor is not empty.

```
# show performance monitor cache detail [format record]
# show performance monitor history
```

- Verify policy and class-map hits (counters should increase).

```
# show policy-map type performance-monitor interface g0/0/2
 GigabitEthernet0/0/2
  Service-policy performance-monitor input: mymon_in
    Class-map: select_ipv4_tcpperf (match-all)
      354704 packets, 75729623 bytes
      30 second offered rate 1000 bps, drop rate 0000 bps
      Match: protocol ip
      Match: access-group name ipv4_tcpperf
    Class-map: class-default (match-any)
      0 packets, 0 bytes
      30 second offered rate 0000 bps, drop rate 0000 bps
      Match: any
```

- Review the running-config and verify that nothing is missing or misconfigured. The problem can be caused by even a single access-list missing.
- **Cisco IOS XE Platforms**: Verify that account-on-resolution (AOR) is active.
  - If AOR is active, handles will have a non-zero value, as shown in the following example:

```
# show platform hardware qfp active feature fnf datapath aor
 CFT:  ConfigAddress 0x8a1e16a0, Instance 0x8a1de760, Feat ID 1, FlowObj ID 1
 CVLA: handle 0x97f00000 epoch 0x4
```

  - If AOR is inactive, handles will have the value of zero, as shown in the following example:

```
# show platform hardware qfp active feature fnf datapath aor
 CFT:  ConfigAddress 0x8a1e16a0, Instance 0x00000000, Feat ID 0, FlowObj ID 0
 CVLA: handle 0x0 epoch 0x4
```

Incorrect TCP Performance Statistics

The following may be helpful for troubleshooting incorrect TCP performance statistics:

- Verify that the monitor that includes TCP performance metrics is applied to only one interface.
- For that interface, service-policy must be attached in both directions.
- Check for asymmetric routing.
- Verify that routes/route-maps are configured correctly.
- If filtering applications, ensure that the appropriate class-map has hits.
- Verify that account-on-resolution (AOR) is active. For details about verifying AOR, see Report Is Not Displayed Correctly, page 5-1.
- Enable IP NBAR Protocol Discovery on the interface to determine whether the protocol of interest is identified.

```
Router(config-if)# ip nbar protocol-discovery
Router# show ip nbar protocol-discovery interface g0/0/3
```

```
GigabitEthernet0/0/3
Last clearing of "show ip nbar protocol-discovery" counters 00:00:10
                            Input                     Output
                            -----                     ------
  Protocol                  Packet Count              Packet Count
                            Byte Count                Byte Count
                            30sec Bit Rate (bps)      30sec Bit Rate (bps)
                            30sec Max Bit Rate (bps)  30sec Max Bit Rate (bps)
  ------------------------- ------------------------- -------------------------
  http                      7                         8
                            3472                      1740
                            0                         0
                            0                         0
```

Memory/Cache Warning

An error message typically occurs if the total memory consumed by all monitors exceeds 25% of the total available memory. If the memory required for all enabled features exceeds the memory available, the following may be helpful for troubleshooting:

- Review the configuration. If there are mismatches, remove the configuration and reapply it.
- Reduce the FNF monitor cache size.

Also see Cache Size Recommendation, page 6-1.

Cache Warning on Cisco IOS Platforms

On Cisco IOS platforms, the following type of MMA warning can occur:

```
7310: 2013-09-17T00:32:02: %SCRIPT-6-DIAG: Sep 16 23:55:56.459 PDT: %MMA-3-CACHE_OVERFLOW:
The number of flows has exceeded 95% of the configured size, monitor testing-url_ipv4,
please increase cache size
```

Memory Warning on Cisco IOS XE Platforms

On Cisco IOS XE platforms, the following type of FNF warning can occur:

```
Oct 28 14:44:10.358 IST: %QFP_FNF-4-FNF_MEM_UPLIMIT_WARN: F0: cpp_cp:  Netflow and
Flexible Netflow configuration is using (140199440) bytes of data plane DRAM which exceeds
the recommended maximum of (134217728) bytes.
```

This warning message indicates that a large amount of memory is allocated to Flexible NetFlow (FNF) monitors. Allocating this amount of memory to FNF monitors is acceptable, but the total memory required by all other enabled features must not exceed the available memory.

More Than 32 Matches per Class

The following may be helpful for troubleshooting the following type of error message regarding configuring more than 32 matching statements:

```
cannot configure more than 32 matching statements per class-map for the interface
```

- Review your class-map configuration.

  # **show class-map**

- Make sure every class-map has no more than 32 match instructions, including hierarchical classes. Remove redundant match instructions

More Than Five Monitors per Class

The following may be helpful if you receive the following type of error message regarding the limit of five (5) monitors per policy per class:

```
%Only 5 monitors allowed per policy per class
```

- Review the class-map configuration.

  # **show class-map**

- Verify that every class-map has no more than five monitors, including FNF monitors which are applied directly on the interface. Remove any redundant monitors and retry.

AVC Notes, Limitations, and Caveats

Revised: November 21, 2013, OL-30581-01

This section includes the following topics:

Notes

Hidden Fields

Two hidden fields (first/last timestamp) are implicitly added to each record, even when these fields are not explicitly configured. When the fields are not explicitly configured, the fields are not exported and are not displayed using **show** commands. Because of these two hidden fields, the effective maximum number of supported fields is the upper limit defined for the release, minus two.

Cache Size Recommendation

The cache size to configure is determined by the traffic profile. The cache should be large enough to store all traffic records, but not excessively large. A warning message may appear if the configured cache exceeds 25% of DRAM. For troubleshooting information, see Memory/Cache Warning, page 5-3.

Limitations

ISSU Limitations

Cisco In-Service Software Upgrade (ISSU) provides transparent router software upgrade or downgrade. ISSU enables bug fixes, deployment of new features, and even complete upgrade of the Cisco IOS software image. For more information, see:
http://www.cisco.com/en/US/products/ps7149/products_ios_protocol_group_home.html

This section describes ISSU limitations for AVC.

Removing Aliases before Downgrading from Cisco IOS 15.4(1)T / Cisco IOS XE 3.10 or Later

Cisco IOS Platforms	Cisco IOS XE Platforms
Applicable to release 15.4(1)T and later	Applicable to release 3.10S and later

In Cisco IOS XE release 3.10S and Cisco IOS release 15.4(1)T, aliases were introduced to the AVC monitor configuration syntax. Using the **all** alias simplifies configuration statements and optimizes performance. (See CLI Field Aliases, page 4-31.)

Before downgrading from one of these releases, or a later release, to a version that does not support aliases, remove the aliases and manually expand the statements to specify each of the required fields explicitly. Failure to remove aliases before downgrading will result in undesired behavior, including possible system crash.

Downgrading to an IOS XE Version that Does Not Support More than 32 Fields

Cisco IOS Platforms	Cisco IOS XE Platforms
Not applicable	Applicable to release 3.10S and later

AVC for Cisco IOS XE 3.10 introduced support for configuring records containing 40 fields. If a record configuration includes more than 32 fields, downgrading to an IOS XE version that does not support more than 32 fields is not supported.

Before downgrading from Cisco IOS XE 3.10 or later, to a version, such as IOS XE 3.9, that does not support more than 32 fields, remove any record configuration of more than 32 fields.

Note Some record configurations include hidden fields. Hidden fields count toward the total supported number of fields. See Hidden Fields, page 6-1.

Note Upgrading from a version that does not support more than 32 fields to a version that does support more than 32 fields is supported.

Error Caused By Using a Performance Monitor With Default Cache Size

Cisco IOS Platforms	Cisco IOS XE Platforms
Not applicable	Applicable to release 3.11S and later

Symptom

Using a performance monitor when the cache size is set to its default value may cause an error during the Cisco In-Service Software Upgrade (ISSU) process. An error in the console log will indicate a failure to update the monitor cache size.

Conditions

1. Applicable to all Cisco IOS XE platforms.
2. Occurs when running ISSU, which provides transparent router software upgrade or downgrade.
3. May occur when doing either one of the following:
   - Upgrading from Cisco IOS XE 3.10 or earlier to IOS XE 3.11 or later version
   - Downgrading from IOS XE 3.11 (or later) to a version earlier than 3.11

Workaround

A preventive workaround and typical use case is to configure the cache size manually rather than using the default.

If using the default cache size, use the following workaround to avoid the error:

1. Remove the service policy.
2. Run the system upgrade or downgrade.
3. Re-attach the service policy.

Performance Monitor Limitations

Cisco IOS Platforms	Cisco IOS XE Platforms
Applicable	This limitation is not applicable.

Performance monitors operate in different modes, depending on the metrics that they are configured to collect. For maximum performance, any of the following metrics may be used. Including other metrics may impact performance.

- Match Fields
  - match application name [account-on-resolution]
  - match connection client ipv4 (or ipv6) address
  - match connection server ipv4 (or ipv6) address
  - match connection client transport port
  - match connection server transport port
  - match ipv4 protocol

  - match policy qos index
  - match routing vrf input
- Collect Fields
  - collect application http host
  - collect application http uri statistics
  - collect connection all
  - collect datalink mac source address
  - collect interface [input/output]
  - collect ip dscp
  - collect ipv4 ttl (or ipv6 hop-limit)
  - collect policy qos classification hierarchy
  - collect policy qos queue [drops/index]
  - collect timestamp sys-uptime first
  - collect timestamp sys-uptime last

Example of Record Including Metrics That Do Not Reduce Performance

```
flow record type performance-monitor Conversation-Traffic-Stats-IPv4(6)
    match ipv4 protocol
    match application name account-on-resolution
    match connection client ipv4 (or ipv6) address
    match connection server ipv4 (or ipv6) address
    match connection server transport port
    match routing vrf input
    collect interface input
    collect interface output
    collect ipv4 dscp
    collect connection client counter packets long
    collect connection server counter packets long
    collect connection client counter bytes long
    collect connection server counter bytes long
    collect connection new-connections
    collect connection sum-duration
    collect ipv4 ttl   (or ipv6 hop-limit)
    collect timestamp sys-uptime first
    collect timestamp sys-uptime last

flow record type performance-monitor Application-Response-Time-IPv4(6)
    match ipv4 protocol
    match application name account-on-resolution
    match connection client ipv4 (or ipv6) address
    match connection server ipv4 (or ipv6) address
    match connection server transport port
    match routing vrf input
    collect interface input
    collect interface output
    collect ipv4 dscp
    collect connection client counter packets long
    collect connection server counter packets long
    collect connection client counter bytes long
    collect connection server counter bytes long
    collect connection new-connections
    collect connection sum-duration
    collect ipv4 ttl   (or ipv6 hop-limit)
    collect connection delay application sum
```

```
    collect connection delay application max
    collect connection delay response to-server sum
    collect connection delay response client-to-server sum
    collect connection delay network client-to-server sum
    collect connection delay network to-client sum
    collect connection delay network to-server sum
    collect connection transaction duration sum
    collect connection transaction counter complete
    collect connection client counter packets retransmitted
    collect connection server counter responses
    collect connection delay response to-server histogram late
    collect timestamp sys-uptime first
    collect timestamp sys-uptime last

flow record type performance-monitor URL-IPv4(6)
    match ipv4 protocol
    match application name account-on-resolution
    match connection client ipv4 (or ipv6) address
    match connection server ipv4 (or ipv6) address
    match connection server transport port
    match routing vrf input
    collect interface input
    collect interface output
    collect ipv4 dscp
    collect connection client counter packets long
    collect connection server counter packets long
    collect connection client counter bytes long
    collect connection server counter bytes long
    collect connection new-connections
    collect connection sum-duration
    collect ipv4 ttl   (or ipv6 hop-limit)
    collect connection delay application sum
    collect connection delay application max
    collect connection delay response to-server sum
    collect connection delay response client-to-server sum
    collect connection delay network client-to-server sum
    collect connection delay network to-client sum
    collect connection delay network to-server sum
    collect connection transaction duration sum
    collect connection transaction counter complete
    collect connection client counter packets retransmitted
    collect connection server counter responses
    collect connection delay response to-server histogram late
    collect timestamp sys-uptime first
    collect timestamp sys-uptime last
    collect application http uri statistics
    collect application http host
```

Caveats

Caveats describe unexpected behavior. Severity 1 caveats are the most serious caveats. Severity 2 caveats are less serious. Severity 3 caveats are moderate caveats.

To view caveats related to the use of AVC, see the release notes for your platform.

If you have an account on Cisco.com, you can also use the Bug Search tool to find select caveats of any severity. See:
https://tools.cisco.com/bugsearch/search

(If the defect that you have requested is not displayed, it may be that the defect number does not exist, the defect does not have a customer-visible description, or the defect is for internal Cisco use.)

Derived Fields Caveat

Cisco IOS Platforms	Cisco IOS XE Platforms
Not applicable	Releases prior to 3.10S

Caveat **CSCue53207**, described in the *Cisco ASR 1000 Series Aggregation Services Routers Release Notes*, describes a bug in some earlier releases, in which a record that contains certain derived fields (listed below) may be punted incorrectly to the route processor (RP) and lost. When using any of the **connection delay** fields listed in the Workaround description below, downgrading to a release that contains this bug is not recommended.

The following is a description of the bug:

Symptom

A record that contains certain derived fields (listed below) may be punted incorrectly to the route processor (RP) and lost.

Conditions

Records can collect "derived" fields; calculating derived fields is dependent on the values of other fields. The fields listed below are incorrectly defined as derived and dependent on other fields. When a record contains one of these fields and does not include its dependent fields, the record is punted to the route processor (RP) to complete the record processing. Punting these records might lead to record loss.

Workaround

When configuring a monitor to collect one of the fields listed below, collect each of the dependent fields also. The list indicates the dependencies:

1. “connection delay application sum” is dependent on:

   connection delay response to-server sum

   connection delay network to-server sum

   connection server response sum

2. “connection delay application min” is dependent on:

   connection delay response to-server min

   connection delay network to-server sum

3. “connection delay application max” is dependent on:

   connection delay response to-server max

   connection delay network to-server sum

4. “connection delay response client-to-server sum” is dependent on:

   connection delay response to-server sum

   connection delay network to-server sum

   connection server response sum

5. “connection delay response client-to-server min” is dependent on:

   connection delay response to-server min

   connection delay network to-server sum

   connection server response sum

   connection delay response to-server sum

   connection delay network to-server min

6. “connection delay response client-to-server max” is dependent on:

   connection delay response to-server max

   connection delay network to-server sum

   connection server response sum

   connection delay response to-server sum

   connection delay network to-server max

Oversubscribed FNF Monitor Caveat

Cisco IOS Platforms	Cisco IOS XE Platforms
Not applicable	Releases prior to 3.10S

Caveat **CSCud15949**, described in the *Cisco ASR 1000 Series Aggregation Services Routers Release Notes*, describes a bug affecting releases prior to IOS XE 3.10S. For these releases, you can attach up to two policies per interface and direction. The total number of monitors included in the two policies should not exceed 10. In calculating the total number of monitors:

- Each policy is considered to include at least five monitors, even if fewer than five monitors are configured for the policy.
- An FNF static monitor is counted as 1 monitor.

The bug may occur (on the affected releases) if these limits are exceeded on any interface, either for ingress or egress traffic on the interface. This condition is called “oversubscribed.”

When a system is oversubscribed, downgrading to a release that contains this bug is not recommended. For oversubscribed systems, Cisco In-Service Software Upgrade (ISSU) does not enable downgrading to a release prior to 3.10S.

The following is a description of the bug:

Symptom

The CPP traceback notifying monitor cannot be reserved.

Conditions

The issue was seen when the MMA policy, mediatrace policy, and one FNF monitor were attached to an interface.

Workaround

Ensure that the total number of monitors does not exceed the limits outlined above, in the description of this bug.

Use Synchronized Cache for Optimized Monitors

Cisco IOS Platforms	Cisco IOS XE Platforms
Release 15.4(1)T	Not applicable

Caveat **CSCuh87789** describes a limitation affecting routers running Cisco IOS 15.4(1)T. On affected releases, use "synchronized cache" when configuring optimized monitors. Do not use, for example, the "normal cache" option. Synchronized cache is the default cache mode for the router.

Using a cache option other than synchronized may result in failure to export certain metrics, resulting in incomplete records.

Network Time Mismatch Between IOS and QFP Causing Dropped Records

Cisco IOS Platforms	Cisco IOS XE Platforms
Release 15.4(1)T	Not applicable

Caveat **CSCul27478** describes a problem that may occur when there is a clock mismatch between Cisco IOS and the router's QuantumFlow Processor (QFP). When this occurs, records punted from the QFP to IOS may be identified as late records, and incorrectly dropped instead of being exported.

The workaround for this issue is to configure an NTP server that allows the IOS clock to be synchronized with network time.

AVC Supported Platforms and Interfaces

Revised: November 21, 2013, OL-30581-01

This chapter addresses the following topics:

AVC Supported Platforms

Cisco AVC is supported on the following platforms:

- Cisco IOS Platforms (Cisco ISR G2 Routers)
  - Cisco C1921-AX/K9
  - Cisco C1941-AX/K9
  - Cisco C2901-AX/K9
  - Cisco C2911-AX/K9
  - Cisco C2921-AX/K9
  - Cisco C2951-AX/K9
  - Cisco C3925-AX/K9
  - Cisco C3925E-AX/K9
  - Cisco C3945-AX/K9
  - Cisco C3945E-AX/K9
  - Cisco 800 Series: C892FSP-K9, C896VA-K9, C897VA-K9, C897VAW-A-K9, C897VA-M-K9, C898EA-K9, C897VAW-E-K9, C897VAM-W-E-K9.
- Cisco IOS XE Platforms
  - Cisco ASR1000 Series Aggregation Services Routers
  - Cisco ISR4400 Series Integrated Services Routers
  - Cisco CSR 1000V Cloud Services Routers

For information about licensing and features for supported platforms, see:
AVC Licensed Features, page 3-1

Logical Interfaces Not Supported by AVC

Logical interfaces *not* supported by Cisco AVC in the current release:

- Dialer interfaces

  Supported on Cisco IOS platforms, not supported on Cisco IOS XE platforms
- Dynamic tunnels such as Dynamic Virtual Tunnel Interface (DVTI)
- IPv6 tunnels that terminate on the device
- Pass-through tunneled IPv6
- Multiprotocol Label Switching (MPLS)
- Overlay Transport Virtualization (OTV) overlay interfaces
- Virtual template interface

Logical interfaces *partially* supported by Cisco AVC in the current release:

- FLEXVPN

  Supported only in spoke-to-spoke topology

AVC Feature History

Revised: November 21, 2013, OL-30581-01

This chapter addresses the following topic:

- Feature History, page B-1

Feature History

The tables below describe highlights of new features and optimizations in recent AVC releases. They do not provide a full feature history of Cisco AVC.

- For Cisco IOS Releases, see Table B-1.
- For Cisco IOS XE Releases, see Table B-2.

Table B-1 AVC Feature History for Cisco IOS Releases

Feature	Description
Cisco IOS 15.4(1)T	
Convergence of Cisco AVC Architecture Across Platform Types	The convergence of AVC architecture brings together the strongest AVC features from IOS and IOS XE platforms, providing powerful features and greater standardization of configuration tasks across different Cisco platforms. Metrics, such as ART, HTTP, and QoS metrics, that were available in earlier releases can now be configured in the same way as on Cisco IOS XE platforms. Additional metrics are also newly available for Cisco IOS.
Metric Mediation Agent (MMA)	The Metric Mediation Agent (MMA) introduces an enhancement to Cisco AVC infrastructure, enabling addition of stateful and derived parameters with dynamic registration. The MMA provides aggregation of connections, history, and alarms from the route processor. The aggregated data is exported at a lower speed than the data path export. For more information about the MMA, see: Metric Mediation Agent, page 2-7

Feature	Description
QoS Metrics	This Cisco AVC release provides new monitors for collecting metrics related to Quality of Service (QoS) policy. Monitors can indicate: • Packets dropped on an interface, per QoS queue, due to a QoS policy that limits resources available to a specific type of traffic. • Class hierarchy (indicating traffic priority) of a reported flow, as determined by the QoS policy map. For more information, see: QoS Metrics: Cisco IOS Platforms, page 4-13
Easy Performance Monitor (ezPM) Configuration	Easy Performance Monitor "express" method of provisioning monitors. Easy perf-mon provides "profiles" that represent typical deployment scenarios. After a user selects a profile and specifies a small number of parameters, Easy perf-mon provides the remaining provisioning details. This release provides one profile, which includes five different traffic monitors. Future releases will provide additional options. For more information, see: Easy Performance Monitor, page 4-27
Customizing attribute values	See Customizing Attribute Values, page 4-11.
Export Spreading	The export-spreading feature spreads out the export of records from the monitor cache over a time interval, to improve collector performance. For more information, see: NetFlow/IPFIX Flow Monitor, page 4-4
IPv6 Support	The Cisco AVC solution supports both IPv4 and IPv6.

Features Available Prior to Cisco IOS 15.4(1)T

Unified Solution	Unifies the technologies of several reporting/control solutions. AVC technologies include the configuration mechanism, metrics, and reports of such components as TCP performance, Medianet, and so on.
Media Metrics	For an overview of the metrics collected by Cisco routers, both for Cisco IOS and for Cisco IOS XE, see: *Cisco Application Visibility and Control Field Definition Guide for Third-Party Customers*
Cisco Performance Agent (MACE) Metrics, including: • Application response (ART) • FNF • HTTP • QoS	For information about using these metrics, see: http://www.cisco.com/en/US/docs/ios-xml/ios/wan_waas/configuration/15-mt/cfg-avc-mace.html
TCP Performance Metrics	AVC includes several TCP performance measurements for traffic performance reporting.
Cisco Prime Infrastructure	The Cisco Prime Infrastructure management and reporting system is an integral part of the Cisco AVC solution and provides extensive management and reporting features, including provisioning the system, storing exported data, and generating reports.

Table B-2 AVC Feature History for Cisco IOS XE Releases

Feature	Description
Cisco IOS XE 3.11S	
New metric added to track information about the interface being monitored	The "observation point id" metric provides the physical port number of the interface to which the monitor is attached.
Customizing attribute values	See Customizing Attribute Values, page 4-11.
Interoperability with Cisco GET VPN	See NBAR Interoperability with Cisco GET VPN, page 2-13.
Cisco IOS XE 3.10S	
Improved Exporting Model	An improved and optimized exporting configuration model includes: • Exporting only a single record per packet, reducing duplicate data. • Optimizing monitor assignment. • Filtering low-bandwidth traffic. • Per server reports. The improved exporting model is used as part of the Easy Performance Monitor profile included in this release.
Easy Performance Monitor Configuration	Easy Performance Monitor "express" method of provisioning monitors. Easy perf-mon provides "profiles" that represent typical deployment scenarios. After a user selects a profile and specifies a small number of parameters, Easy perf-mon provides the remaining provisioning details. This release provides one profile, which includes five different traffic monitors. Future releases will provide additional options. For more information, see: Easy Performance Monitor, page 4-27
Performance Improvements	This release includes changes, such as an improved exporting model, predefined monitors, and MMA optimization, that improve performance by up to 30%.
Parsing URI Address	This release introduces the ability to parse URI addresses, enabling AVC to report depth 1 of the URI and filter traffic according to that value. For more information, see: HTTP URI, page 4-34
Support for Records with 40 Fields	This release introduces support for configuring records containing 40 fields.
Cisco IOS XE 3.9S	
Enhanced Connection/Transaction Metrics	Beginning with IOS XE release 3.9S, Flexible NetFlow (FNF) monitors can report on individual transactions within a flow. This enables greater resolution for traffic metrics. For more information, see: Connection/Transaction Metrics, page 4-24

Feature	Description
QoS Metrics	This Cisco AVC release provides new monitors for collecting metrics related to Quality of Service (QoS) policy. Monitors can indicate: • Packets dropped on an interface, per QoS queue, due to a QoS policy that limits resources available to a specific type of traffic. • Class hierarchy (indicating traffic priority) of a reported flow, as determined by the QoS policy map. For more information, see: QoS Metrics: Cisco IOS XE Platforms, page 4-18
Cisco IOS XE 3.8S	
Interoperability with Cisco AppNav	Cisco AppNav is the Wide Area Application Services (WAAS) diversion mechanism. Beginning with IOS XE release 3.8S, AVC provides statistics before and after the AppNav WAAS service controller (AppNav SC), as well as inspecting and reporting application information on optimized traffic.
Unified Solution	Unifies the technologies of several reporting/control solutions. AVC technologies include the configuration mechanism, metrics, and reports of such components as TCP performance, Medianet, and so on.
Metric Mediation Agent (MMA)	The Metric Mediation Agent (MMA) is a new infrastructure element developed in the IOS XE 3.8 release to manage, correlate, and aggregate metrics from different metric providers. MMA provides the following functions: • Controls traffic monitoring and filtering policy. • Correlates data from multiple metric providers (see Metric Providers, page 2-7) into the same record. • Aggregates metrics. • Supports history and alert functions. This requires sending the metrics records to the route processor (RP) before exporting them to the management and reporting tools.
TCP Performance Metrics	This release adds several TCP performance measurements for traffic performance reporting.
Interoperability with AppNav	AppNav is the Wide Area Application Services (WAAS) diversion mechanism. AVC for IOS XE 3.8 provides statistics before and after the AppNav WAAS service controller (AppNav SC), as well as inspecting and reporting application information on optimized traffic.
Packet Capture	Cisco Embedded Packet Capture (EPC) technology performs packet capture.
Cisco Prime Infrastructure	The Cisco Prime Infrastructure management and reporting system is an integral part of the Cisco AVC solution and provides extensive management and reporting features, including provisioning the system, storing exported data, and generating reports.
IPv6 Support	The Cisco AVC solution supports both IPv4 and IPv6.

References

Revised: November 21, 2013, OL-30581-01

The following table provides additional reference material.

Document	Description
AVC	
Application Visibility and Control (AVC)	Cisco Application Visibility and Control (AVC) home page (www.cisco.com/go/avc).
Cisco Application Visibility and Control Field Definition Guide for Third-Party Customers	Overview of the metrics and Flexible NetFlow (FNF) IDs exported by Cisco routers, both for Cisco IOS and for Cisco IOS XE.
AVC Configuration on Cisco IOS Platforms	
Configuring AVC to Monitor MACE Metrics	Information about AVC configuration using MACE on Cisco ISR G2 routers, prior to the Cisco IOS release 15.4(1)T. AVC continues to support MACE configuration but users are encouraged to migrate to MMA.
Licensing	
Software Activation Configuration Guide, Cisco IOS XE Release 3S	Activating licensed features in Cisco IOS XE.
Cisco ASR 1000 Series Aggregation Services Routers Ordering Guide	License ordering guide for Cisco ASR 1000 Series routers.
Cisco CSR 1000V Series Cloud Services Router Release Notes	Image and license information for CSR 1000V
Cisco Software Activation on Integrated Services Routers	Activating licensed features for Cisco Integrated Services Routers (ISR)
Related Components	
Applying QoS Features Using the MQC	Defining traffic policy using the Modular Quality of Service CLI (MQC).
Cisco IOS Quality of Service Solutions Configuration Guide	Configuring Cisco QoS.
Classifying Network Traffic Using NBAR in Cisco IOS XE Software	Configuring Cisco NBAR.
NBAR Protocol Pack Library	NBAR protocol library and NBAR2 protocol packs.

Document	Description
Cisco Performance Monitor and Mediatrace QuickStart Guide	Cisco Performance Monitor and Mediatrace.
Cisco Prime Infrastructure	Cisco Prime Infrastructure home page, with links to product documentation.
Cisco IOS Embedded Packet Capture	Cisco IOS Embedded Packet Capture (EPC) documentation.
Configuration	
Cisco IOS Flexible NetFlow Command Reference	Flexible NetFlow commands.
Getting Started with Configuring NetFlow and NetFlow Data Export	Configuring NetFlow and NetFlow Data Export.
Configuring NetFlow and NetFlow Data Export	Configuring NetFlow network traffic data export.
Easy Performance Monitor	Configuring Easy Performance Monitor (ezPM).

Revised: November 21, 2013, OL-30581-01

A

AVC Application Visibility and Control

C

CFT Common Flow Table

CP Control Plane

CSR Cloud Services Router

D

DP Data plane

E

ESP Embedded Services Processor

F

FNF Flexible NetFlow

FW Firewall

I

IP Internet Protocol – Layer 3 Datagram Protocol. IP provides features for addressing, type-of-service specification, fragmentation and reassembly, and security. Defined in RFC 791 (IPv4) and RFC 2460 (IPv6).

IPC Inter Process Communication

IPFIX Internet Protocol Flow Information Export

L

L2 Datalink Layer (layer 2) of the ISO reference model

L3 Network Layer (layer 3) of the ISO reference model

L4 Transport Layer (layer 4) of the ISO reference model

L7 Application Layer (layer 7) of the ISO reference model

M

MACE Measurement, Aggregation, and Correlation Engine

MMA Metric Mediation Agent

MMON Media Monitoring

N

NAT Network Address Translation

NBAR/NBAR2 Network Based Application Recognition

P

PA Performance Agent

R

RP Route Processor

RSVP Resource Reservation Protocol

S

SNMP Simple Network Management Protocol

SSRC Synchronization Source

T

TCP Transmission Control Protocol—L4 Reliable Transport Mechanism. Connection-oriented transport layer protocol that provides reliable full-duplex data transmission. TCP is part of the TCP/IP protocol stack.

TSS TCP Session State

U

UDP User Datagram Protocol—L4 Transport Mechanism. Connectionless transport layer protocol in the TCP/IP protocol stack. UDP is a simple protocol that exchanges datagrams without acknowledgments or guaranteed delivery, requiring that error processing and retransmission be handled by other protocols. UDP is defined in RFC 768.

V

VRF Virtual Routing and Forwarding

W

WAAS Wide Area Application Services

WAN Wide Area Network

WCM WAAS Central Manager

Z

ZBFW Zone Based Firewall