



AVC Notes, Limitations, and Caveats

Revised: March 2014

This section includes the following topics:

- [Notes, page 6-1](#)
- [Limitations, page 6-2](#)
- [Caveats, page 6-7](#)

Notes

- [Hidden Fields, page 6-1](#)
- [Cache Size Recommendation, page 6-1](#)
- [Fragmented Packets, page 6-1](#)

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](#).

Fragmented Packets

AVC handles fragmented packets as follows:

The first fragment packet is treated normally. AVC treats and reports subsequent fragments as non-TCP/UDP packets.

Limitations

- [General Limitations, page 6-2](#)
- [ISSU Limitations, page 6-4](#)
- [Performance Monitor Limitations, page 6-5](#)

General Limitations

- [Multicast, page 6-2](#)
- [Delay Before New NBAR Configuration Is Activated, page 6-2](#)
- [Do Not Use the Management Interface for Exporting Records, page 6-3](#)
- [Minimum Interval Between Assigning and Removing a Performance Monitor, page 6-3](#)

Multicast

AVC support for multicast is as follows:

- Supported:
 - MediaMonitoring (Calculates and reports media (RTP) performance metrics)
- Not supported:
 - Account on Resolution (AOR)
 - Application Response Time (ART) metric collection
 - NBAR

Delay Before New NBAR Configuration Is Activated

| Cisco IOS Platforms | Cisco IOS XE Platforms |
|---------------------|------------------------|
| 15.4(2)T | 3.12S |

After updating an NBAR configuration, there is a delay before the new configuration is active on the data path.

When you update a configuration, a parallel configuration (the previous) operates during the changeover of configuration to prevent any impact on classification during the transition from one configuration to the next.

After changing a configuration, it takes some time before NBAR classifies traffic according to the new configuration.

The following query indicates the status of the new configuration:

```
Device# show platform software nbar statistics | i NBAR
NBAR state is ACTIVATED
NBAR config send mode is ASYNC
NBAR config state is READY
```

Possible output:

- NBAR State—Status of NBAR component.
 - ACTIVATED
 - DEACTIVATED
- NBAR Configuration Send Mode
 - ASYNC—In normal operation, the NBAR configuration send mode is asynchronous.
- NBAR Configuration State
 - Ready—The new configuration is active.
 - Pending—The new configuration is not yet active. The previous configuration remains active until the new configuration becomes active.
 - Error—The new configuration is not active in the data path due to an error. In this error state, the previous configuration may or may not be active.

Do Not Use the Management Interface for Exporting Records

| Cisco IOS Platforms | Cisco IOS XE Platforms |
|---------------------|--|
| Not applicable | Applicable to all Cisco IOS XE platforms operating with AVC. |

Applicable To

All Cisco IOS XE platforms operating with AVC.

Description

Do not use the management interface (typically GigabitEthernet0) as the source interface of the exporter. The management interface can be identified by the “MGMT ETHERNET” or “GigE” labeling of the physical port. Not following this guideline may cause unexpected behavior, including system crash.

Minimum Interval Between Assigning and Removing a Performance Monitor

A minimum interval of approximately 5 seconds is required between assigning a performance monitor and removing it, or between removing and assigning again.

Specifically:

- After assigning an AVC performance monitor to an interface, wait approximately 5 seconds before removing the performance monitor from the interface.
- After removing an AVC performance monitor from an interface, wait approximately 5 seconds before re-assigning the same performance monitor to the interface.

Not waiting the required time interval may cause some AVC functionality to fail; waiting the required 5 seconds and attempting the configuration again typically resolves the issue. In extreme cases, AVC may stop functioning; to resolve this, restart the router.

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: [In-Service Software Upgrade](#).

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, page 6-4](#)
- [Downgrading to an IOS XE Version that Does Not Support More than 32 Fields, page 6-4](#)
- [Error Caused By Using a Performance Monitor With Default Cache Size, page 6-5](#)

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 4-33](#).)

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

- [Affect of Specific Metrics on Performance, page 6-5](#)

Affect of Specific Metrics on Performance

| 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, page 6-8](#)
- [Oversubscribed FNF Monitor Caveat, page 6-9](#)
- [Use Synchronized Cache for Optimized Monitors, page 6-10](#)
- [Incorrect Record Metric Values When FNF Cache Is Full, page 6-10](#)
- [Clock Mismatch Between QFP and Operating System Causes Records To Be Dropped, page 6-11](#)
- [CSR1000V Platform: Large Jitter Value Reported for Voice/Video Flow, page 6-12](#)
- [Incorrect Jitter Value Reported for RTP Streams, page 6-13](#)
- [After Route Processor Switchover, ezPM Record Export May Fail, page 6-13](#)

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.

Incorrect Record Metric Values When FNF Cache Is Full

| Cisco IOS Platforms | Cisco IOS XE Platforms |
|---------------------|--------------------------|
| Not applicable | 3.11.0 3.11.1 3.12 |

Caveat **CSCum52041** describes a problem that may occur when the FNF cache reaches a full state.

Symptom

Updating of some records in the FNF cache may fail intermittently. Metrics in these records may not reflect complete router traffic.

Conditions

1. A large number of match keys are defined in the configuration: total length of all key fields is more than 32 bytes.

2. The FNF record cache is full.

Workaround

None.

Further Problem Description

To determine if the FNF cache was full at some time during record collection, use one of the following commands. A value greater than 0 for the flows-not-added counter indicates that the cache reached the full state at some point.

For native FNF:

```
show flow monitor MONITOR-NAME cache
```

For a performance monitor:

```
show performance monitor cache MONITOR-NAME
```

Clock Mismatch Between QFP and Operating System Causes Records To Be Dropped

| Cisco IOS Platforms | Cisco IOS XE Platforms |
|---------------------|--|
| Not affected | <p>Affects the following releases:</p> <ul style="list-style-type: none"> • 3.10S: all releases • 3.11.0S <p>Issue resolved in:</p> <ul style="list-style-type: none"> • 3.11.1S and later • 3.12S: all releases |

Caveat **CSCu127478** describes a problem that may occur due to a clock mismatch between the IOS XE operating system 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.

Symptom

Records are dropped (not exported).

Conditions

The problem may occur when there is a clock mismatch between QFP and the IOS XE operating system.

Workaround

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

Alternatively, upgrade to a release that resolves this issue.

Further Problem Description

If the following CLI shows that there are late records, this problem may be occurring:

```
Device# show performance monitor statistics <monitor name>
MMA Internal Stats:
Agg Record Stats:
=====
Record total recv : 41
Record dropped Gen      : 0
Record dropped late : 0
Record total processed : 0
Malloc failed (low memory) : 0
Others : 1
Per Monitor Record Stats:
=====
```

It is also possible to compare timestamps between QFP and IOS XE to determine whether there is a clock mismatch. This may be done by comparing timestamps in an RP platform debug log.

Related Topics

- Caveat [CSCul00248](#). If you have an account on Cisco.com, you can use the [Bug Search](#) tool to view this caveat.
- Caveat [CSCum07636](#). If you have an account on Cisco.com, you can use the [Bug Search](#) tool to view this caveat.

CSR1000V Platform: Large Jitter Value Reported for Voice/Video Flow

| Cisco IOS Platforms | Cisco IOS XE Platforms |
|---------------------|--|
| Not affected | Affects CSR1000V platforms only. Affects the following releases: <ul style="list-style-type: none"> • 3.9S: all releases • 3.10S: all releases • 3.11.0S, 3.11.1S • 3.12S |

Caveat [CSCun33822](#) describes a problem affecting jitter values reported on CSR1000V platforms.

Symptom

Jitter values for voice/video flows are reported inaccurately, often in the hundreds of milliseconds.

Conditions

Relevant for a voice/video RTP flow on a CSR1000V platform.

A Medianet performance monitor is configured to monitor and report RTP statistics, such as jitter and packet-loss.

Workaround

None.

Incorrect Jitter Value Reported for RTP Streams

| Cisco IOS Platforms | Cisco IOS XE Platforms |
|---------------------|-----------------------------------|
| Not affected | All releases, beginning with 3.8S |

Symptom

The jitter measurement for RTP streams with a dynamic payload type (96-127) may be incorrect.

There is no dynamically learned mapping between the payload type and the clock frequency used in the specific RTP stream. The frequency is always set to 90 KHz.

Conditions

Affects RTP streams with a dynamic payload type.

Workaround

None.

After Route Processor Switchover, ezPM Record Export May Fail

| Cisco IOS Platforms | Cisco IOS XE Platforms |
|---------------------|------------------------|
| Not affected | 3.11.0 3.11.1 |

Caveat **CSCun24943** describes a problem affecting ezPM record export after a route processor switchover.

Symptom

After route processor (RP) switchover, ezPM does not operate on the newly active RP. Records are not exported.

Conditions

Stateful switchover (SSO) is configured. Switchover occurs.

Workaround

Re-apply the ezPM configuration or switchover to the original RP after it recovers from failure.

