



# Release Notes for Cisco IOS Release 15.2(4)GC2

---

**Current Release:**  
**15.2(4)GC2—July 2014**  
**Previous Release:**  
**15.2(4)GC1—April 2014**

These releases support the Cisco 5900 Embedded Services Routers (ESR) and select ISR G2 platforms. These release notes are updated to describe new features, limitations, troubleshooting, recommended configurations, caveats, and how to obtain support and documentation.

## Contents

This publication consists of the following sections:

- [Image Information and Supported Platforms, page 2](#)
- [Related Documentation, page 2](#)
- [Limitations, page 3](#)
- [Recommended Configuration Settings, page 5](#)
- [Caveats, page 9](#)
  - [Cisco IOS Release 15.2\(4\)GC2, page 9](#)
  - [Cisco IOS Release 15.2\(4\)GC1, page 10](#)
- [Obtaining Documentation and Submitting a Service Request, page 11](#)



# Image Information and Supported Platforms

These images are bug compatible with Cisco IOS Release 15.2(4)M6.

**Note**

---

You must have a Cisco.com account to download the software.

---

Cisco IOS Release 15.2(4)GC2 includes the following Cisco IOS image:

- c2900-universalk9-mz.SPA
- c2951-universalk9-mz.SPA
- c3900-universalk9-mz.SPA
- c3900e-universalk9-mz.SPA
- c5915-adventerprisek9-mz.SPA
- c5915-entbase-mz.SPA
- c5921i86-universalk9-ms.SPA
- c5930-adventerprisek9-mz.SPA
- c5940-adventerprisek9-mz.SPA

Cisco IOS Release 15.2(4)GC2 is supported on the following platforms:

- Cisco 2901 Integrated Services Router
- Cisco 2911 Integrated Services Router
- Cisco 2921 Integrated Services Router
- Cisco 2951 Integrated Services Router
- Cisco 3925 Integrated Services Router
- Cisco 3925e Integrated Services Router
- Cisco 3945 Integrated Services Router
- Cisco 3945e Integrated Services Router
- Cisco 5915 Embedded Services Router
- Cisco 5921 Embedded Services Router
- Cisco 5930 Embedded Services Router
- Cisco 5940 Embedded Services Router

## Related Documentation

The following documentation is available:

- *Software Configuration Guide, Cisco IOS Release 15.2(4)GC*  
[http://www.cisco.com/en/US/docs/solutions/GGSG-Engineering/15\\_2\\_4\\_GC/config\\_guide/Configuration\\_Guide.pdf](http://www.cisco.com/en/US/docs/solutions/GGSG-Engineering/15_2_4_GC/config_guide/Configuration_Guide.pdf)
- IOS Bulletins—You can find bulletins at  
<http://www.cisco.com/c/en/us/products/ios-nx-os-software/ios-software-releases-listing.html>

# Limitations

This release has the following limitations:

- The Quality of Service (QoS) policy can only be applied to one outgoing interface the Point-to-Point Protocol over Ethernet (PPPoE) session is traversing.

A QoS output policy can be applied to the Virtual Template or the Virtual Multipoint Interface (VMI), but not at the same time. If a policy is attached, the outgoing physical interfaces (i.e., physical-interface FastEthernet0/0) should not have output policy applied. It is recommended that the policy be attached to the Virtual Template. The other option is to apply the policy to the VMI, but not to the Virtual Template or Ethernet interface.

- When a service policy is applied to the VMI and packets are dropped on the VA due to credit starvation, the **show policy-map int VMI** command will not show these dropped packets. There is no back pressure between the interfaces in this configuration. The VMI does not know that packets were dropped by the Virtual Access (VA) or the Ethernet physical interfaces.
- RFC 5578 credits do not tie into QoS formulas. Credits only indicate to QoS the ability to transmit a packet or not. If there are enough credits a packet will be transmitted from the highest priority queue. When there are not enough credits, packets will be queued.
- QoS policy may drop on an interface when the interface receives an invalid CDR value.
- To prevent dropping QoS policy, the current data rate (CDR) is range checked to a floor value of 10kbps and a ceiling value of the physical interface speed. If CDR is below the floor value, the QoS rate shaping is set to the floor value. If CDR is above the ceiling value, the QoS rate shaping is set to the ceiling value. The floor value is 10kbps. The ceiling value is limited to the physical interface speed.
- Software Release 15.2(4)GC2 may have OSPFv3 incompatibility issues while in Mobile Ad Hoc Networks (MANET) mode with software versions 12.4(24)GC or previous versions. To avoid any issues, ensure that all routers run the same version of Cisco IOS.
- CSCtx62099

If the size of the IP mux MTU is larger than the MTU of the outbound interface, a superframe may not be transmitted, or it might be fragmented properly but not re-assembled at the receive side. A format error occurs on the demux interface. This is only an IPv4 issue; IPv6 does not support router packet fragmentation.

**Workaround** Reduce the MTU size in the profile to the size of the interface MTU.

## The Cisco 5921 ESR has the following limitations:

- With the 5921 ESR, if the user wishes to configure dot1q trunking, it is suggested that the Linux MTU be set to 1508. This is due to the fact that a frame of length 1500 (maximum, by default) is going to leave the 5921 application with a length of 1504 bytes (if using dot1q and for non-native VLAN) or of 1508 bytes (if using q-in-q and for non-native VLAN). Normally, Linux would either drop or fragment these frames. Configuring the MTU to 1508 in Linux will prevent the fragmentation or dropping of frames in the event that dot1q or q-in-q is configured on the 5921. Even if the Linux MTU is set to greater than 1500, the 5921 is going to "cap" its MTU at 1500 (to the 5921 application). This will mimic the behavior of a typical Cisco device.

- With the 5921 ESR, several situations have been encountered in which the e1000e Ethernet driver strips VLAN tags before a frame reaches the 5921. This will result in dot1q trunking not performing properly (the 5921 will receive frames with no VLAN tag, even though it is configured to expect VLAN tags). In IOS, you will notice this by seeing ARP or ping failures. If you see such behavior, please issue the following debug command from IOS:

```
#debug arp
```

Now, try the ping again. If the VLAN tag stripping issue is present, you will see a "wrong cable" message similar to the following:

```
*Jan 14 21:49:50.874: IP ARP rep filtered src 192.168.110.2 e05f.b986.5500, dst
192.168.110.1 0022.4d7b.e424 wrong cable, interface Ethernet0/0.130
```

Now switch to the Linux command line and see if the e1000e driver is being used by issuing the following command (using eth0 as an example):

```
[root@router ~]# ethtool -i eth0
driver: e1000e <== LOOK FOR THIS
version: 2.3.2-k
firmware-version: 2.1-2
bus-info: 0000:02:00.0
supports-statistics: yes
supports-test: yes
supports-eprom-access: yes
supports-register-dump: yes
supports-priv-flags: no
```

From the Linux command line, verify the VLAN mode of the device:

```
[root@router ~]# ethtool -d eth0
MAC Registers
-----
0x00000: CTRL (Device control register) 0x58100248
      Endian mode (buffers):          little
      Link reset:                      reset
      Set link up:                     1
      Invert Loss-Of-Signal:           no
      Receive flow control:            enabled
      Transmit flow control:           enabled
      VLAN mode:                       enabled <== LOOK FOR THIS
      Auto speed detect:                disabled
      Speed select:                    1000Mb/s
      Force speed:                      no
      Force duplex:                     no
```

If the VLAN mode is enabled, this indicates that the driver is stripping the VLAN tags.

To remedy this using CentOS, please upgrade to the latest e1000e driver by following these steps from the Linux command line:

```
#yum install kernel-devel gcc gcc-c++ make wget
#reboot

#mkdir /usr/local/src/e1000e
#cd /usr/local/src/e1000e
#wget http://downloadmirror.intel.com/15817/eng/e1000e-3.0.4.tar.gz
#tar xzf e1000e-3.0.4.tar.gz
#cd e1000e-3.0.4
#cd src
#make install
```

If the process ends with the following message, ignore it:

```
/bin/sh: man: command not found
```

```
#rmmod e1000e
#modprobe e1000e
```

From the Linux command line, verify that the new driver has been activated:

```
[root@centos src]# ethtool -i eth0
driver: e1000e
version: 3.0.4-NAPI    <== LOOK FOR THIS
firmware-version: 2.1-2
bus-info: 0000:02:00.0
supports-statistics: yes
supports-test: yes
supports-EEPROM-access: yes
supports-register-dump: yes
supports-priv-flags: no
-----
```

## The Cisco 5915 ESR has the following limitations:

- The Virtual Local Area Network (VLAN) database is maintained in the flash as a file (*vlan.dat*). This file is created or updated only while adding, deleting, or modifying a VLAN to a switchport. Do not delete the *vlan.dat* file.

Copying a configuration, with switch ports configured with VLANs, through Trivial File Transfer Protocol (TFTP) to the Cisco 5915 ESR running configuration, the *vlan.dat* file may not be created automatically in flash. As a result, the VLANs configured in the switch ports do not exist, even though the running config shows the switch ports configured with VLANs. If a VLAN appears as 'Line protocol down' the VLAN association needs to be updated by configuring the VLAN database explicitly by entering the **vlan** *vlan\_number* command.

When you copy any configuration file with VLAN configurations to TFTP server, back up the corresponding *vlan.dat* file to TFTP. And when you copy a configuration file from a TFTP server to the Cisco 5915 ESR, copy the corresponding *vlan.dat* file to C5915 router flash (CSCtq04969).

- When you copy files to the flash, the copy process may become very slow when flash has limited free space. Delete unused files to free up space. If the issue persists, enter the **format flash** command to reclaim space from the deleted files (CSCt42754).
- The Output drops counter may not be accurate under stress. You can verify the number of output drops in output for the **show controller** command (CSCt170552).

## Recommended Configuration Settings

Use the following configuration guidelines when enabling class-based weighted fair-queuing:

- Enter the following command to turn off creation of virtual-template subinterfaces:

```
no virtual-template subinterface
```

- Enter the following commands to create a policy map with class-based weighted fair-queuing and apply the newly created policy-map to the virtual template:

```
class-map match-any chat
  match dscp af11
class-map match-any voice
  match dscp ef
```

```

policy-map mypolicy
  class chat
    bandwidth percent 40
  class voice
    bandwidth percent 40

interface virtual-template number
  service-policy output mypolicy
    
```

- No additional configuration is supported on the policy-map.

Use the following configuration guidelines when disabling Point to Point Protocol (PPP) keepalives:

- You can turn off the PPP keepalive messages to decrease overhead when the radio alerts the router with a PPPoE Active Discovery Terminate (PADT) message that the layer-2 RF connection is no longer available. Turning off the PPP keepalive messages may also avoid the potential for the router to terminate the connection based on missed PPP keepalives over a poor radio frequency (RF) link.
- To turn off the PPP keepalive messages, enter the following command for the virtual-template:

```

interface virtual-template number
  no keepalive
    
```

Use the following configuration guidelines for setting the recommended Open Shortest Path First (OSPF) values of radio link metrics:

- You may have to dampen the amount of changes in order to reduce network-wide churn because cost components may change rapidly.
- The following recommended values are intended as a starting point for optimizing an OSPFv3 network and are based on network simulations that may reduce the rate of network changes. Each network may have unique characteristics that require different settings to optimize actual network performance.

You must configure these values for both OSPFv3 IPv4 and IPv6:

```

S1 = ipv6 ospf dynamic weight throughput
    Recommended value = 0
    Default=100
S2 = ipv6 ospf dynamic weight resources
    Recommended value = 29
    Default = 100
S3 = ipv6 ospf dynamic weight latency
    Recommended value = 29
    Default = 100
S4 = ipv6 ospf dynamic weight L2 factor
    Recommended value = 29
    Default = 100
    
```

The following is an example configuration for a VMI interface on the virtual template when running bypass mode:

```

interface vmi1
...
ipv6 ospf cost dynamic weight throughput 0
ipv6 ospf cost dynamic weight resources 29
ipv6 ospf cost dynamic weight latency 29
ipv6 ospf cost dynamic weight L2-factor 29
...
ospfv3 cost dynamic weight throughput 0
ospfv3 cost dynamic weight resources 29
ospfv3 cost dynamic weight latency 29
ospfv3 cost dynamic weight L2-factor 29
    
```

For more information on OSPF commands, refer to the following URL:

[http://www.cisco.com/en/US/docs/ios/iproute\\_ospf/command/reference/iro\\_book.html](http://www.cisco.com/en/US/docs/ios/iproute_ospf/command/reference/iro_book.html)

Use the following configuration guidelines for disabling split horizon in EIGRP:

- By default split horizon is enabled in EIGRP. You can disable split horizon by entering the **no ip split-horizon eigrp** command for the respective autonomous system number.

```
interface vmi number
no ip split-horizon eigrp as-number
```

- Enter the following command to disable the ip redirects on the vmi interface when you are configuring the vmi interface for EIGRP:

```
interface vmi number
no ip redirects
```

Use the following configuration guidelines for setting EIGRP values of radio link metrics:

- EIGRP monitors the following metrics on an interface allowing the tuning of the EIGRP metric calculations; use the metric weights router configuration command:

```
metric weights tos k1 k2 k3 k4 k5
```

where `tos` denotes type of service (currently, it must always be zero) and use the following default values for weights:

k1 - 1

k2 - 0

k3 - 1

k4 - 0

k5 - 0




---

**Note** The **no metric weights** command restores the K-values to the defaults listed above.

---

- Most configurations use the Delay and Bandwidth metrics with Bandwidth taking precedence.
- You must set the weights identically on all routers in an autonomous system.




---

**Note** If you wish to use the default K-values you do not need to enter the **metric weights** command.

---

- To set the metric dampening value for EIGRP, enter the following commands for either change-based or interval-based dampening of metric updates received through VMI:

- Change Based Dampening:

```
ip50-1(config)#int vmi 4
ip50-1(config-if)#eigrp 100 interface dampening-change 40
```

Default Value for Change Based Dampening:

50%

To enable change-based dampening using the default Value, enter the following command:

```
eigrp 100 interface dampening-change
```

To disable change-based dampening, enter the following command:

```
no eigrp 100 interface dampening-change
```

– Interval-based Dampening:

```
ip50-1(config)#int vmi 4
ip50-1(config-if)#eigrp 100 interface dampening-interval 20
```

Default Timer value for Interval-based Dampening:  
30 seconds

To enable interval-based dampening using the default Value, enter the following command:

```
eigrp 100 interface dampening-interval
```

To disable interval-based dampening, enter the following command:

```
no eigrp 100 interface dampening-interval
```

- The following exceptions will result in an immediate update:
  - a down interface
  - a down route
  - any change in a metric triggered outside the scope of the VMI metric update



**Note**

---

No recommended values other than default are currently available.

---

For more information on EIGRP commands, refer to:

[http://www.cisco.com/en/US/docs/ios/iproute\\_eigrp/command/reference/ire\\_book.html](http://www.cisco.com/en/US/docs/ios/iproute_eigrp/command/reference/ire_book.html)

Use the following configuration guidelines to enable Waas Express:

- Enter the following command to turn Waas Express on the Wan interface:

```
waas enable
```

- Enter the following commands to enable full optimizations:

```
parameter-map type waas waas_global
tfo optimize full
```

- Enter the following commands to enable application accelerators:

```
parameter-map type waas waas_global
accelerator http-express
enable
accelerator cifs-express
enable
accelerator ssl-express
enable
```

For more information on Waas Express configuration, refer to the following URLs:

[http://www.cisco.com/en/US/docs/ios/wan/configuration/guide/cfg\\_waas\\_exp.pdf](http://www.cisco.com/en/US/docs/ios/wan/configuration/guide/cfg_waas_exp.pdf)

[http://www.cisco.com/en/US/docs/ios-xml/ios/wan\\_waas/configuration/15-2mt/wan-cfg-waas-exp.html](http://www.cisco.com/en/US/docs/ios-xml/ios/wan_waas/configuration/15-2mt/wan-cfg-waas-exp.html)



# Caveats

Caveats describe unexpected behavior in Cisco IOS releases. Caveats listed as open in a prior release are carried forward to the next release as either open or closed (resolved).

## Cisco IOS Release 15.2(4)GC2

The following sections list caveats for Cisco IOS Release 15.2(4)GC2:

### Open Caveats

- **CSCup61110**

If the 5940 does not have a start up configuration, it displays a message that is not informative:

```
Original NVCONFIG doesnt have correct MAGIC number  
Backup NVCONFIG also doesnt have correct MAGIC number
```

**Workaround:** Ignore the message.

- **CSCup68078**

On the 5921, if there is a time mismatch between the IOS time and the Linux host time, a license invalidation error may occur.

**Workaround:** If your HW platform supports a real-time clock circuit, please ensure that the clock is set to the actual time and date.

If your HW platform does not support a real-time clock circuit, please manually set the time to 1 July 2014 before starting the 5921 application. Please note, this clock setting will not be persistent across reboots or power removals, so it needs to be done each time the platform is booted (before the 5921 is loaded). In CentOS, this can be done by adding the following command in the c5921-swr-init.sh script: `date --set="1 Jul 2014 00:00:00"`

In addition, if your platform does not support a real-time clock, do not use NTP inside IOS.

- **CSCup75978**

On the 5921, if the Linux MTU is set to below 100 for the TAP interface, ping will not work.

**Workaround:** There is no workaround.

### Closed Caveats

- **CSCul97065**

On the 5940, the running configuration would get cleared by randomly power cycling the system.

- **CSCun69568**

A router may occasionally reload when updating the configuration using nested TCL scripts.

- **CSCum13603**

The 5940 would occasionally reload when making changes to interface IP addresses. This has only been observed when NAT is configured on a VRF.

- **CSCum46854**

Cisco 5915 router with 'ip mobile router-service' crashes when Ethernet cable is unplugged.

- **CSCuo13216**

When the 5921 is starting up the user would see the following message:

```
WARNING: No HOME environment => No default SWROptions File
```

As long as the user has specified a specific location for the SWROPTIONS file, the warning message can be ignored.

- **CSCue62144**

A 5921 running under a load could see a jump in CPU usage when running an active console.

- **CSCum79890**

On the 5921, EIGRP update packets were getting lost between the 5921 IOS and linux interface, when DLEP session is established between peers.

- **CSCun11206**

On the 5921, in some cases you cannot ping or arp across a dot1q interface. Please see the [Limitations](#) section for suggested workaround.

- **CSCuj35636**

Performing a shutdown on an interface, in some scenarios, can result in very high CPU utilization. This CPU utilization can be determined by using the Linux application, top. The problem only occurred on the 5921.

## Cisco IOS Release 15.2(4)GC1

The following sections list caveats carried forward from Cisco IOS Release 15.2(4)GC1:

### Open Caveats

- **CSCun69568**

A system crash may happen when the configuration is updated by using TCL scripts.

**Workaround:** None

- **CSCum13603**

On the 5940, with NAT enabled on a VRF, you may experience a crash and reload when making changes to an interface IP address.

**Condition:** This has only been seen when NAT is enabled on a VRF.

**Workaround:** None

### Closed Caveats

- **CSCua04651**

On the 5915 and 5940 ESRs, there is a small memory leak when cache is quickly filled. The Router stops forwarding traffic due to memory exhaustion. This memory leak was reported on both IPv4 & IPv6.

- **CSCty54465**

When an ACL is placed on an inbound interface, and IP multiplexing is enabled on the interface, traffic that was "not filtered" was received after being de-muxed by IP multiplexing.

- **CSCuj17818**

When PPPoE is configured on radio interfaces between two routers, and a shut and no shut are issued on the remote interface of router2, 9 packets are stuck in the Router1 input queue.

- **CSCum77865**

A router configured to NAT an encrypted (IPsec) GRE tunnel source doesn't actually do the NAT-ing unless CEF is disabled.

**Previous Workaround:** Disable IP CEF switching ('no ip route-cache cef') on the physical interface, where the data packets arrive that are being forwarded out through the tunnel.

- **CSCua49735**

The WAAS-Express router crashes when HTTP-Express Accelerator is enabled and HTTP traffic is going through the WAAS-Express router. This condition is only seen if the response header is split across multiple receives on WAASX.

- **CSCun11619**

On the 5940 ESR, there were two conditions that sometimes cause the ESR to hang indefinitely rather than crash.

1. With NAT enabled on a VRF, you may have experienced a crash and reload when making changes to an interface IP address.
2. The WAAS-Express router might crash when HTTP-Express Accelerator was enabled and HTTP traffic is going through the WAAS-Express router. This condition was only seen if the response header is split across multiple receives on WAASX.

## Obtaining Documentation and Submitting a Service Request

For information on obtaining documentation, using the Cisco Bug Search Tool (BST), submitting a service request, and gathering additional information, see What's New in Cisco Product Documentation at: <http://www.cisco.com/en/US/docs/general/whatsnew/whatsnew.html>.

Subscribe to What's New in Cisco Product Documentation, which lists all new and revised Cisco technical documentation, as an RSS feed and deliver content directly to your desktop using a reader application. The RSS feeds are a free service.

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](http://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)

© 2012-2014 Cisco Systems, Inc. All rights reserved.