

Design Guide

# Packet Interception Mode (PIM) Design Guide 2.0

---



Juniper Networks, Inc.  
1194 North Mathilda Avenue  
Sunnyvale, CA 94089 USA  
408 745 2000 or 888 JUNIPER  
[www.juniper.net](http://www.juniper.net)

## Table of Contents

Table of Contents.....	2
Summary .....	3
Packet Interception Theory of Operation.....	3
Overview .....	3
How Packet Interception works .....	4
Deciding Which Method to Use .....	5
When to use Packet Interception Mode.....	5
When NOT to use Packet Interception Mode.....	6
Deploying Packet Interception Mode in a High Availability Network .....	6
Packet Interception Matrix .....	7
External Mode.....	8
Operational Details for External mode.....	8
WCCP.....	18
Operational Details for WCCP mode .....	18
High Availability using WCCP .....	24
Route Injection .....	26
Operational Details for Route Injection.....	26
Additional Network Considerations for Route Injection.....	32

## Summary

This document expands on the existing 5.0 documentation to provide more information about advanced topics for off-path deployments

## Packet Interception Theory of Operation

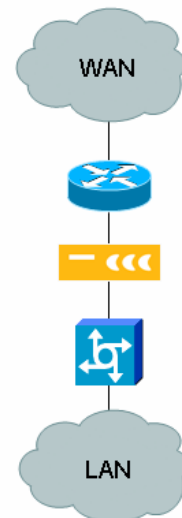
### Overview

In the simplest network deployment, WX and WXC application acceleration platforms (“WX/WXC devices”) are inserted into the network in “inline” mode (see Figure 1, the orange device is the WX/WXC device). The WX/WXC device is Ethernet attached to a switch on the LAN side and a router on the WAN side. All packets that flow between the LAN and WAN pass through the WX/WXC device and can receive WAN optimization services. The WX/WXC device is completely transparent to the routed network; it appears as a transparent bridge. Inline mode also offers hardware bypass support so that if anything goes wrong with it, including total power failure, the WX/WXC device “switches-to-wire,” in effect becoming an Ethernet cable, so that network operation can continue without interruption.

Inline mode is ideal for many environments and is generally the recommended mode where possible due to its simplicity, but there are network topologies that do not lend themselves to an inline deployment. For example, if the router is an integrated switch/router with direct connections to the LAN it may not be possible or convenient to insert the WX/WXC device inline. In cases where inline deployment is not possible or desired, the WX/WXC device can be inserted “off-path” by attaching it to the switch or router in a one-armed fashion.

When a WX/WXC device is installed ‘off-path’, this is called Packet Interception mode. Packet Interception Mode has been available in the WX operating system (WXOS™) software since release 4.0.

The WXOS 5.0 manual has a good overview of Packet Interception Mode and basic configuration examples. This document supplements the manual with more in-depth information about deployment issues and options.



**Figure 1 - Inline**

## How Packet Interception works

There are a variety of ways that a WX/WXC device can be inserted into the network.

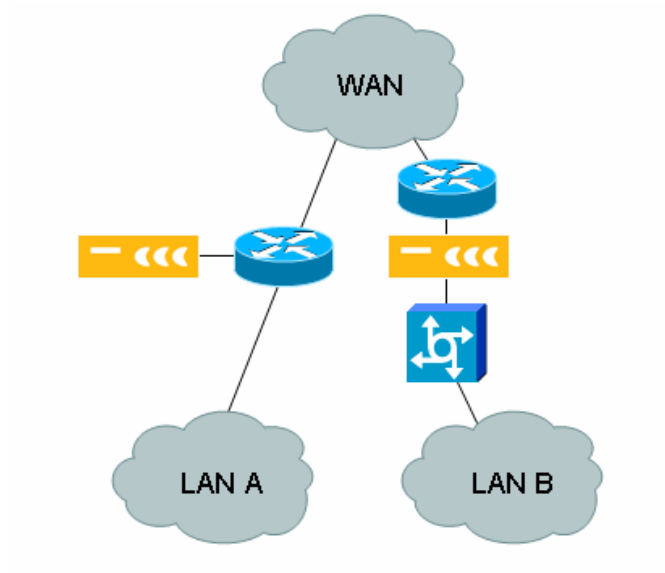


Figure 2 shows a simple example. In this case, the WX/WXC device on the left is directly attached to the WAN router in off-path mode while the WX/WXC device on the right is in inline mode between the switch and router.

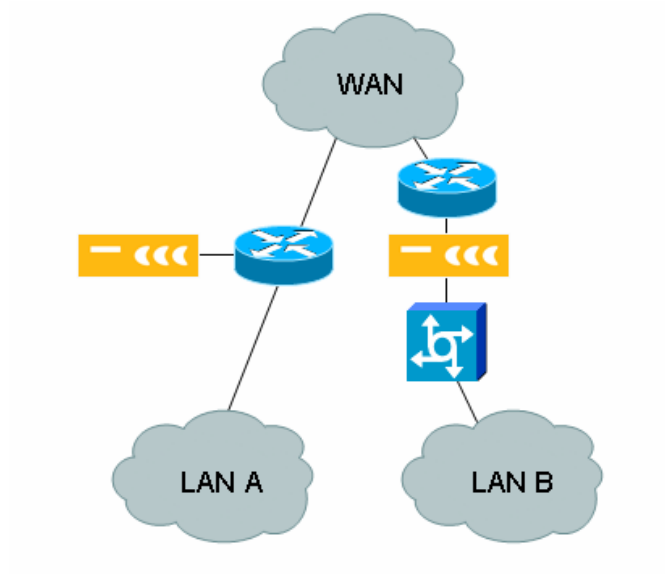


Figure 2 - Off-Path and Inline modes

## Off-Path versus Inline modes

In an "inline" deployment, traffic flows from the LAN to the WAN through the WX/WXC device. The WX/WXC device then performs WAN Optimization functions (acceleration, reduction, QoS, etc.). Optimized traffic flowing from the LAN to WAN is addressed to the remote WX/WXC device where all decompression and de-encapsulation is performed.

When deployed in 'off-path' mode the WX/WXC device is installed in a one-arm fashion thus we will need to direct traffic to it in order for to "flow through" the WX/WXC device. In an off-path deployment, one of the first considerations is to determine what traffic will be forwarded to the WX/WXC device. The goal is to only forward traffic that needs WAN Optimization services. This is not a strict requirement since the WX/WXC device will "pass-through" traffic that has no WAN Optimization policy configured.

The WX/WXC devices form tunnels and advertise the subnets for which they provide WAN optimization services to all other WX/WXC devices. Regardless of the redirection method selected, the goal is to redirect traffic to receive WAN optimization to the off-path WX/WXC device. Once the traffic has been optimized, it is forwarded via the WX/WXC tunnel to the appropriate remote WX/WXC device where it will be re-assembled and directed to the appropriate subnet.

There are three different ways traffic can be directed to an off-path WX/WXC device: Route Injection Mode via RIPv2, WCCP via WCCPv2 and External Mode via Policy Routing. The following sections discuss each of these options.

## Deciding Which Method to Use

### When to use Packet Interception Mode

Packet Interception should be used when it is not possible to deploy WX/WXC devices inline. One example of this is if traffic is originated or destined to incompatible media types (non-Ethernet). Another example is if there are multiple paths traffic can take from the LAN to the WAN and more control of traffic flow may be required, for example, in a High Availability (HA) environment, data may traverse multiple parallel paths to access the WAN. This may require multiple inline WX/WXC devices even if the traffic levels are low to get maximum benefit. In off-path mode traffic flow from several routers can be directed to a single WX/WXC device for WAN optimization.

Understanding the network topology, traffic paths, traffic types and load is very important when planning the deployment of any new network equipment. The preferred location for the WX/WXC device when installed in off-path mode is directly connected to a port on the WAN router or a port on the L3 switch in its own subnet. This implementation provides both ease of installation in both redundant and non-redundant deployments, fault isolation and in case of failure, fast recovery.

## When NOT to use Packet Interception Mode

Due to the different Packet Interception modes, an off-path WX/WXC device can usually be installed in most networks however; care should be taken if the WX/WXC device will be connected to a Layer 2 device. The reason is failure recovery. Using WCCP or Route Injection modes allows for failure recovery via the dynamic nature of these protocols thus recovery will occur when the protocol timers expire. Recovery may be done via dynamically switching to a secondary or backup WX/WXC device or back to native routing. External mode, however, is not dynamic in nature and relies on policy based routing for traffic redirection. Policy based routing will invalidate a route when the route goes away (i.e. the link to the 'next-hop' goes down, or ARP expires). If the WX/WXC device is connected to a L2 device and a failure occurs the L3 device will not immediately detect the failure because it is not directly connected.

## Deploying Packet Interception Mode in a High Availability Network

There are many different considerations when planning the deployment of equipment in any network. These are even more important in High Availability networks. What protocols are being used? How do they interact with each other? What part does each play in providing stability and uptime in the network? Commonly used techniques/protocols include OSPF, BGP, EIGRP, access lists, policy routing, route redistribution, HSRP, VRRP . The unique role each plays in the network should be evaluated. In many cases seemingly simple changes may have unexpected results. Due to the complex nature of these techniques and protocols, not every case can be examined here in detail. The goal of this document is to identify key points to take into account when planning the installation of WX/WXC devices in off-path mode.

## Packet Interception Matrix

The following matrix in Table 1 provides a quick overview of when and how each packet intercept mode could be deployed.

	Route Injection	External	WCCP
Can be installed without disruption of traffic flow	Yes	Yes	Yes
Non-Cisco Routers and Switches	Yes	Yes	No
Internal Routing Protocol used is EIGRP	Yes*	Yes	Yes
Internal Routing Protocol used is RIPv2	Yes*	Yes	Yes
Traffic from multiple routers can be directed to a single WX/WXC.	Yes	Yes*	No
WX/WXC connected to an L2 Switch	Yes*	Yes*	Yes
WX/WXC connected to an L3 Switch/ Router	Yes*	Yes*	Yes
High Availability	Yes	Yes*	Yes*
Pass-Through traffic handling	Tunneled to Remote	Redirected back to router by default or Tunneled via configuration option*	Redirected back to router via GRE tunnel for native routing.

**Table 1 - Packet Interception Matrix**

**\*Some restrictions or specific requirements may apply. See design considerations for the details.**

**Note:** WX/WXC devices can be installed using different deployment modes, off-path and inline, within the same network based on the requirements at each location.

## External Mode

### Operational Details for External mode

#### Basic Operation:

External mode is typically a non-dynamic method for off-path deployment. When deploying the WX/WXC device using this mode, traffic redirection to the WX/WXC device is achieved on a L3 device using Policy Based Routing.

The flexibility of access lists used in Policy Routing provides the capability to be as specific as necessary to redirect only traffic to be optimized by the WX/WXC device. When applying policies on a L3 device, apply the policies only to the interfaces that will receive traffic that is candidate for redirection (i.e. ports connected to user segments) – usually you do not apply these policies to the interface where the WX/WXC device will connect, as this may cause a route loop.

#### Failover and Recovery:

Failover recovery when using external mode will depend on the amount of time it takes to determine if the WX/WXC device is still able to forward traffic.

##### ■ Link Failure:

If the WX/WXC device is directly connected to the L3 device that is redirecting the traffic using Policy Based Routing then failover will occur once the L3 device is aware the 'link' has failed.

##### ■ Protocol Failure:

If the WX/WXC device is not directly connected to the L3 device redirecting traffic, failure detection will require another mechanism to notify the L3 device that the policy route next hop is unavailable. A common example is the router and the WX/WXC device are both connected to the same VLAN on an L2 switch. While the ARP and/or MAC address are still present in the L3 device, traffic will continue to be directed to the WX/WXC device. Aging of the ARP and/or MAC address for the WX/WXC device could take several minutes. During that time traffic will be lost.

Since this issue is not unique to WX/WXC devices and a common issue with Policy Routing, several routing vendors support tracking mechanisms to allow policy routing to have some dynamic capability to detect next hop reachability. Providing the ability to automatically change the route policy or ignore it completely. One major vendor Cisco has introduced a feature, PBR Support for Multiple Tracking Options, provides a mechanism to track a device using ICMP. This feature will track an object to confirm a device is reachable.

**Note:** The PBR Support for Multiple Tracking Options feature was introduced in IOS version 12.3(4)T.



### Pass-Through Traffic Handling:

In External mode pass-through traffic, filtered traffic or traffic destined for a non-WX/WXC site will be returned to the L3 device for forwarding. To avoid confusing the forwarding database, if the WX/WXC device sends a packet back to the router for normal forwarding the source MAC will be changed to that of the WX/WXC device. This is the default behavior for pass-through traffic in External mode.

External mode also has the option to disallow returning pass-through traffic to the router, which means the WX/WXC device will tunnel all network traffic including pass-through, if it is destined for a WX/WXC-advertised network configured for optimization. If the WX/WXC device receives traffic destined for a subnet that is not advertised by any WX/WXC site for optimization, it will drop this traffic, since it has no match. This should not occur if the route policies and access lists are configured correctly. The following CLI command will tunnel pass-through traffic in External mode: `config packet-interception external set pass-through [ off | on ]`.

The default setting is ON (which means send pass-through traffic back to the router).

Setting this to OFF will tunnel pass-through traffic as described above.

### Before making this change, consider the following:

If the remote WX/WXC device fails it will stop advertising these subnets as reachable then all traffic becomes pass-through which results in all traffic redirected to the WX/WXC device being dropped.

### Design Considerations:

- Place WX/WXC device directly on dedicated L3 interface in its own subnet for best fault isolation and simplicity..
- If WX/WXC device is placed on an L2 switch fault isolation is more difficult. In this the router should have a method for dynamically detecting the status of the next hop (like PBR Tracking Objects).
- Carefully plan out what traffic should be redirected. These are manually created overrides for the natural traffic flow of the network. If there is a problem it is usually with the way the policy is written.
- Route Policies are applied only on the interface receiving (input) the packets not the sending interface (output).
- Whenever possible try implementing routing policies in a test lab first.
- The interface connected to the WX/WXC device and router should always be full duplex. This is not a hard requirement, but the amount of data on this connection is doubled since it goes in and comes out the same interface.
- Disable reverse path forwarding on the interface attached to the WX/WXC device.

**Policy Routing Behavior on Cisco routers:**

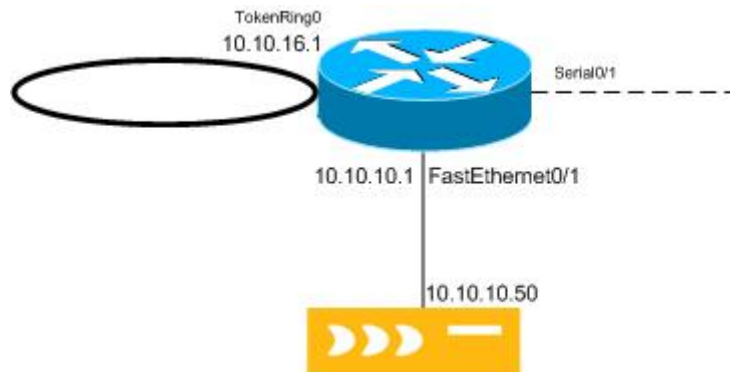
- If the packets do not match any route map statements, then all the set clauses are applied.
- If a statement is marked as deny, the packets meeting the match criteria are sent back through the normal forwarding channels and destination-based routing is performed.
- If the statement is marked as permit and the packets do not match any route map statements, the packets are sent back through the normal forwarding channels and destination-based routing is performed.

**Additional Information:**

Static Routes should not be used in place of Policy Routing, because static routes do not have the same granularity that Policy Routing does. Static Routes are entered on a “global” basis and do not have the options to make decisions on protocols, source/destination address, subnet and interface(s). This global nature has the potential to easily cause routing loops. A very common example is that pass-through traffic by default will be sent back to the router. Since you are redirecting all traffic for those subnets, there will be some pass-through traffic.

When the WX/WXC device sends the pass-through traffic back to the router, the router will send this back to the WX/WXC device since that is where its routing table says it should go. This will continue until the packet TTL expires. If you enable tunneling of pass-through traffic as described above, any traffic not destined for a WX/WXC-advertised subnet that is directed to the WX/WXC device will be dropped. What is the harm in doing this as long as routes are chosen carefully? If the remote WX/WXC device fails, it will stop advertising these subnets as reachable; then all traffic becomes pass-through, resulting in all traffic being redirected to the WX/WXC device being dropped. Due to these issues, static routing is not a recommended deployment configuration with External Mode.

## Basic External Mode Configuration



The diagram above provides a simple example of Packet Interception- using external mode. In this case the traffic we want to optimize is on a Token Ring network. Since the WX/WXC device only has Ethernet interfaces, off-path mode allows us to interoperate in this environment. Obviously if this was all Ethernet the configuration would be the same. As you can see the basic configuration requirements are fairly simple on both the WX/WXC device and the router.

**Note:** WCCP or Route Injection could also work here.

Configuration for WX/WXC device and router below:

### WX/WXC Device

```
Packet Intercept Mode = External
IP address : 10.10.10.50
Reduction subnet(s) : 192.168.0.0/16
Default gateway : 10.10.10.1
Static route :
192.168.0.0/16 gateway 10.10.10.1
```

### Router

```
interface tokenring 0
 ip address 10.10.16.1 255.255.255.0
 ! policy applied to token ring interface
 ip policy route-map Peribit

 ! access-list matches and permits traffic from any source IP
 destined for 192.168.0.0/16
```

```
access-list 120 permit ip any 192.168.0.0 0.0.255.255

! route map Peribit permits traffic that matches access list
120 and redirects to 10.10.10.50

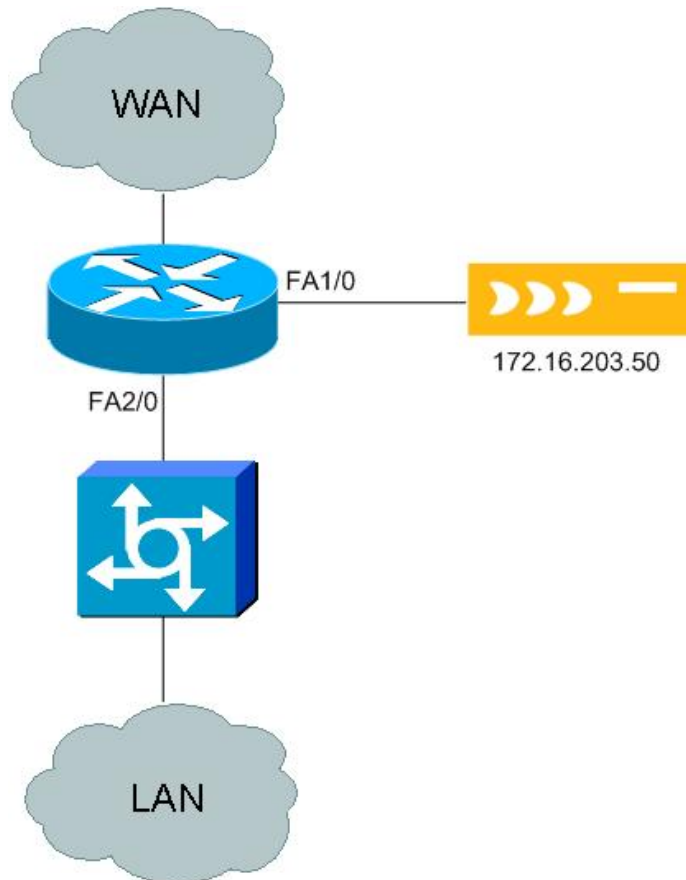
route-map Peribit permit 120

match ip address 120

set next hop 10.10.10.50
```

### External Mode Configuration Example in Detail

Below is another simple network example similar to the previous one. In this example we will look closer at some of the key configuration items that are used.



### Configuration Information:

As you can see, the WX/WXC device configuration is fairly straightforward when enabling External mode. All the traffic redirecting is done at the router and typically we just need a default route and our reduction subnet routes. So after a few clicks and reboot we are ready.

#### WX/WXC Device

```
Packet Intercept Mode = External
IP address 172.16.203.50
Reduction subnet(s) : 172.16.10.0/24, 172.16.100.0/24
Default gateway : 172.16.203.1
Static routes :
172.16.10.0/24 gateway 172.16.203.1
172.16.100.0/24 gateway 172.16.203.1
```

On the router side, we need to identify a way to direct traffic to the WX/WXC device. The method used is Policy Routing. A policy route is composed a several items.

1. Route-maps These are used to classify the traffic using access-lists, length of packet, etc. For simplicity we will use the access-list example since it is one of the most commonly used.
  - A. Route-maps are processed in order top to bottom. The traffic is compared to the Match statements in the route-map and when the first match occurs the corresponding Set statements are applied to that traffic (like setting the next-hop).

#### Router

```
! interface to the Peribit
interface fa1/0
ip address 172.16.203.1 255.255.255.0

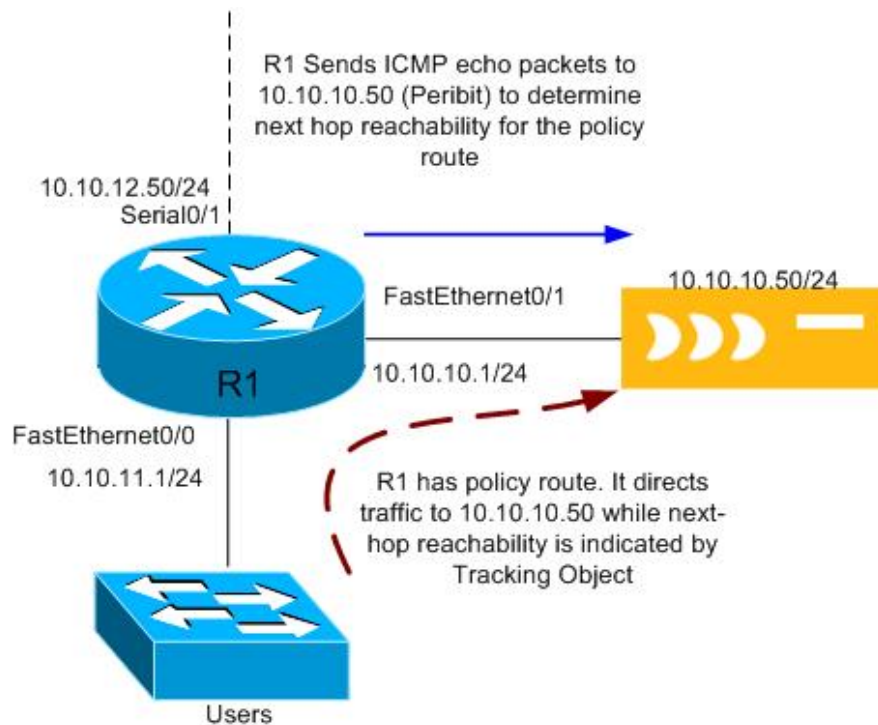
! interface to the LAN segment
interface fa2/0
ip address 172.16.1.1 255.255.255.0
! policy applied on this interface
ip policy route-map Peribit

! access list used for Policy matches traffic destined for
172.16.100.0/24 and 172.16.10.0/24
access-list 120 permit ip any 172.16.100.0 0.0.0.255
access-list 120 permit ip any 172.16.10.0 0.0.0.255
```

```

! Policy to forward traffic matching access list 120 to
specified next hop
route-map Peribit permit 120
  match ip address 120
  set ip next-hop 172.16.203.50
    
```

### Example 3: Dynamic Policy Based Routing



To provide a more dynamic recovery mechanism, we can use Cisco's PBR support for Multiple Tracking Objects. This configuration is very similar to the others we have seen except that this one allows for the tracking of next-hop reachability by the router, which will allow it to change the policy once the next hop goes down. There are too many options for us to cover here, so we will only cover the basics.

#### WX/WXC Device

```
Packet Intercept Mode = External
```

#### R1

```

!configure Response Time Reporter(RTR) rtr 1
rtr 1
  ! configure icmp echo parameters.
  ! and explicitly set the source ip just to be safe
    
```

```
type echo protocol ipIcmpEcho10.10.10.50 source-ipaddr
10.10.10.1

! we set timeout and frequency in this case timeout should
occur in 15 seconds.

timeout 5000
frequency 10

! starts Response Time Reporter instance rtr 1
rtr schedule 1 start-time now life forever

! starts tracking instance 123
track 123 rtr 1 reachability

interface fastethernet 0/0

ip address 10.10.11.1 255.255.255.0
! Applies policy routing to this interface
ip policy route-map Peribit

interface fastethernet 0/1
ip address 10.10.10.1 255.255.255.0

! Access-list matches and permits traffic from any source IP
destined for 192.168.0.0/16
access-list 120 permit ip any 192.168.0.0 0.0.255.255

! route map Peribit permits traffic that matches access-list
120
route-map Peribit permit 10

match ip address 120

set ip next-hop verify-availability 10.10.10.50 10 track
123

! above we set the next hop and its sequence number "10"
```

**Basic components of the tracking option used here.**

RTR Response Time Reporter.

**Creating the instance:**

*rtr <number>*

**Configuring the instance:**

Source IP address is optional, but we put it in just to be safe.

```
type echo protocol ipIcmpEcho <address> source-ipaddr <address>  
timeout <milliseconds>
```

We configured 5000 milliseconds which is the default.

```
Frequency <seconds>
```

We configured 10 seconds default is 60. This should provide a timeout of about 15 seconds if the gateway is lost 10 seconds plus 5000 milliseconds (5 seconds). This should also restore the service about 10 seconds after it becomes available.

Starting the Instance:

```
rtr schedule <rtr instance> start-time now life forever
```

**Creating the Tracking Object:**

```
track <tracking instance> rtr <rtr instance> reachability
```

This starts a tracking instance for the specified RTR and looks for reachability.

**The route-map statement:**

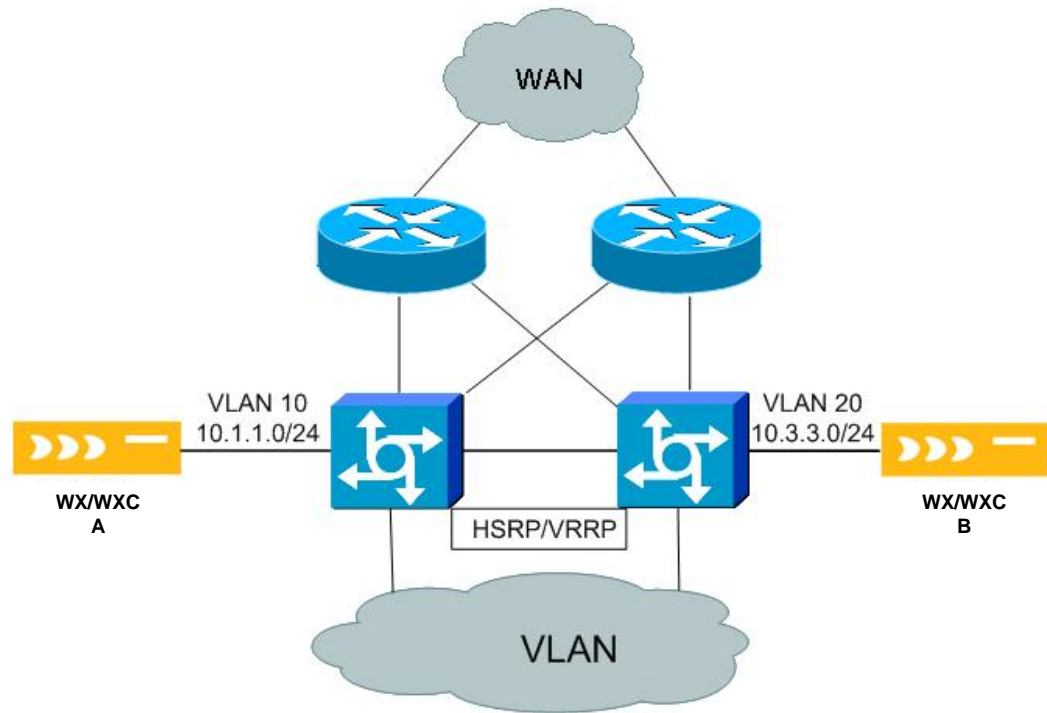
```
set ip next-hop verify-availability <ip address> <sequence number> track <tracking  
instance>
```

Since we covered route-maps previously we are just showing the one line. Multiple next hops can be tracked as well.

**Redundant Configuration using External Mode:**

Redundancy in external mode is normally accomplished by having each router connecting to a WX/WXC device operate independently. Each WX/WXC device would be connected to a single dedicated port or VLAN on the router so failure would be identified immediately. In this configuration the traffic that each router gets that is destined for the configured policy will be redirected to the WX/WXC device. If one WX/WXC device fails the router will continue to forward traffic using the native routing table. The other router will continue to forward traffic to the WX/WXC device for reduction and optimization.





Below is the configuration for both the WX/WXC device and routers in the diagram above. Note the WX/WXC devices are in their own dedicated VLANs.

**WX/WXC Device A:**

```
Packet Intercept Mode = External Mode
Router IP address : 10.1.1.1
Peribit IP Address : 10.1.1.50
```

**WX/WXC Device B:**

```
Packet Intercept Mode = External Mode
Router IP address : 10.3.3.1
Peribit IP Address: 10.3.3.50
```

**L3 Switch 1:**

```
Interface Vlan 1
ip address 10.2.2.1 255.255.255.0
ip policy route-map Peribit

Interface Vlan 10
ip address 10.1.1.1 255.255.255.0
```

```
Interface Vlan 30
ip address 10.4.4.253 255.255.255.0

! Access list matches any source traffic destined for the
10.1.14.0 network
access-list 120 permit ip any 10.1.14.0 0.0.0.255
route-map Peribit permit 120
match ip address 120
set ip next-hop 10.1.1.50
```

### L3 Switch 2

```
interface Vlan 1
ip address 10.2.2.2 255.255.255.0
ip policy route-map Peribit

Interface Vlan 20
ip address 10.3.3.1 255.255.255.0

Interface Vlan 30
ip address 10.4.4.254 255.255.255.0
! Access list matches any source traffic destined for the
10.1.14.0 network
access-list 120 permit ip any 10.1.14.0 0.0.0.255
route-map Peribit permit 120
match ip address 120
set ip next-hop 10.3.3.50
```

## WCCP

### Operational Details for WCCP mode

#### Basic Operation:

The WX/WXC device uses Cisco's WCCPv2 protocol as a dynamic method to direct traffic to the WX/WXC device in an off-path deployment. The WCCPv2 protocol redirects traffic based on any IP protocol number 0-255 that has an associated service group configured. You configure service groups on the WX/WXC device and associate them with an IP protocol number. The service group's WX/WXC device can be configured with are numbered 51-99. The WX/WXC device currently supports a maximum of five service groups. You must also configure the associated service group on the router as well.

WCCPv2 also provides a priority scheme 0-255 so that many web caches can be

configured on a single router. The priority of the cache will determine the order in which the router will redirect traffic; the higher the number, the higher the priority of the WX/WXC devices (default is 230.) Example: You want to use the WX/WXC device to redirect all TCP traffic, but you already have a web cache and want to make sure all http traffic goes there first. In this case you would configure the WX/WXC device with a lower priority so that the router does not bypass the web cache.

WCCPv2 allows for two different communication methods, unicast and multicast. As of the writing of this document, Juniper currently only supports the unicast method. This means a WX/WXC device can only redirect traffic from a single router. There are two different forwarding methods for router to web cache devices L2 rewrite and GRE tunnel. Juniper only supports the GRE tunnel method. All redirected traffic from the router is sent over this tunnel.

When configuring service groups on the router, configure them globally and then on the interface. Access lists can also be used in conjunction with service groups. This allows for fine-grained control of which traffic is redirected to the WX/WXC device.

As stated before, WCCP redirection is configured on a per interface basis. It can be applied either inbound or outbound depending on customer needs and the version of software.

#### **Inbound Redirection:**

Inbound redirection is typically the preferred method for deploying WCCP. Inbound redirection is less CPU intensive on the router since the route lookup and policy decisions are made on ingress of the packet. This also allows for more granular configuration since you only apply this to the interfaces where you want to redirect traffic from. Inbound redirection is newer and not all versions of code that support WCCP support applying it inbound.

#### **Outbound Redirection:**

Outbound redirection is applied on the outgoing interfaces where traffic would be exiting the router on its way to the WAN. This is more CPU intensive since the packets have already been processed on ingress so they must be processed again and then redirected or forwarded. Because traffic is processed in a specific order on the router, it may be that a customer needs to apply several different policies and services to their traffic and may need to process the packets for WCCP on the outgoing interface in order for the other policies/ services to function properly.

### **Failover and Recovery**

WCCP recovery is dependant on multiple factors:

- **Link Failure:**

If the WX/WXC device is directly connected to a router port and the link fails, failover will occur immediately. If the WX/WXC device is connected via an L2 switch, failure time will depend on the WCCPv2 protocol.

- **Protocol Failure:**

The WCCP protocol will age out a non-responding cache in approximately 30 seconds.

Here is a brief overview of the WCCPv2 protocol for a more detailed understanding.

The WX/WXC device sends “WCCP2\_HERE\_I\_AM” messages and the router sends “WCCP2\_I\_SEE\_YOU” messages in response. “WCCP2\_HERE\_I\_AM” messages are sent every 10 seconds. If the router does not see any “WCCP2\_HERE\_I\_AM” messages for 3 times the “WCCP2\_HERE\_I\_AM” interval, WCCP will consider the WX/WXC device as down and no longer redirect traffic.

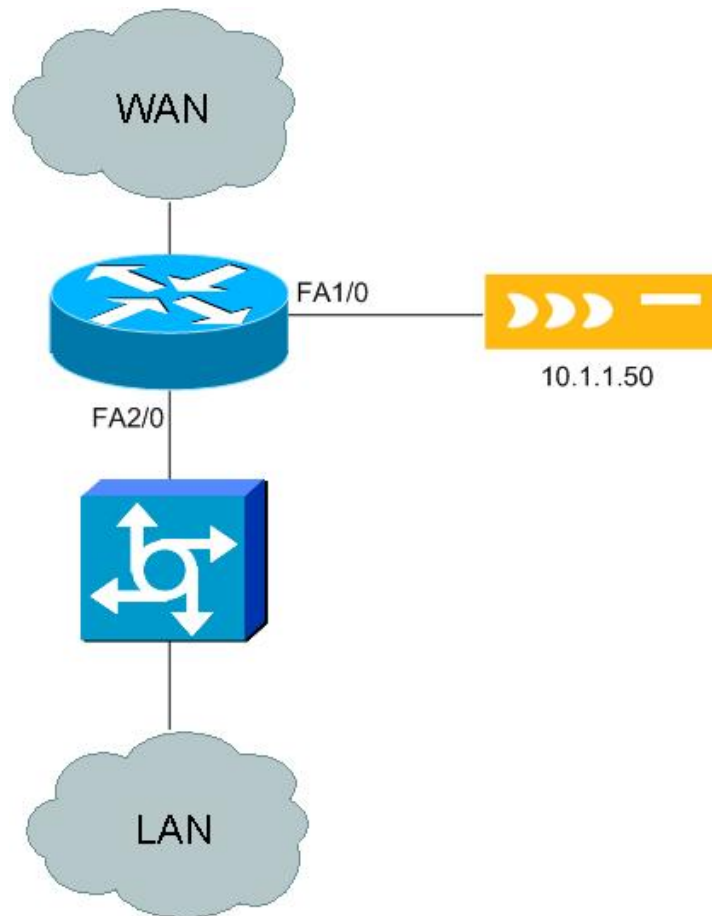
### **Pass-through Traffic Handling:**

In WCCP mode, if traffic redirected to the WX/WXC device does not match any of its criteria for processing it is considered to be pass-through traffic. Pass-through traffic is sent back to the router via the GRE tunnel where it will be forwarded normally by the router.

### **Design Considerations:**

- Place the WX/WXC device directly on a dedicated L3 interface in its own subnet for best fault isolation and simplicity.
- The WX/WXC only supports WCCPv2.
- The WX/WXC supports unicast WCCPv2. Traffic from only one router can be redirected to the WX/WXC device.
- WCCPv2 can dynamically identify a device as down after 30 seconds.
- Can be applied on inbound or outbound interfaces (see Basic Operation for details).
- Access-lists can be used for fine-grained control of traffic that will be redirected to the WX/WXC device.
- If other Web Caches are in use you may want to adjust the priority of the WX/WXC device so it does not interfere.
- Each WCCP service group used must be unique to the device advertising it. Example: you cannot have a WX/WXC device and a Web Cache advertising the same service group.
- While not required it is recommended that you have a loopback IP address on the router running WCCP.
- The interface connected to the WX/WXC device and router should always be full duplex. This is not a hard requirement, but the amount of data on this connection is doubled since it goes in and comes out the same interface.
- Disable reverse path forwarding on the interface attached to the WX/WXC device.

### Basic Inbound WCCP Configuration



The diagram above shows the WX/WXC device in a basic inbound WCCP deployment. Using WCCP the router will redirect all packets received on interface FastEthernet2/0, except those destined to 20.1.1.1 to the WX/WXC device. You could also use the access-list to further filter the type (TCP/ UDP, etc) of traffic redirected to the WX/WXC device. This feature allows for very granular control of the traffic.

**Note:** In most cases you will want to use inbound WCCP since it is more efficient.

#### WX/WXC Device

```
Packet Intercept Mode = WCCP
Router IP address = 10.1.1.1
Priority = 230
TCP = service-id 60
UDP = service-id 70
```

### Router

```

! create wccp service groups 60 and 70
ip wccp 60 redirect-list 101
ip wccp 70 redirect-list 101

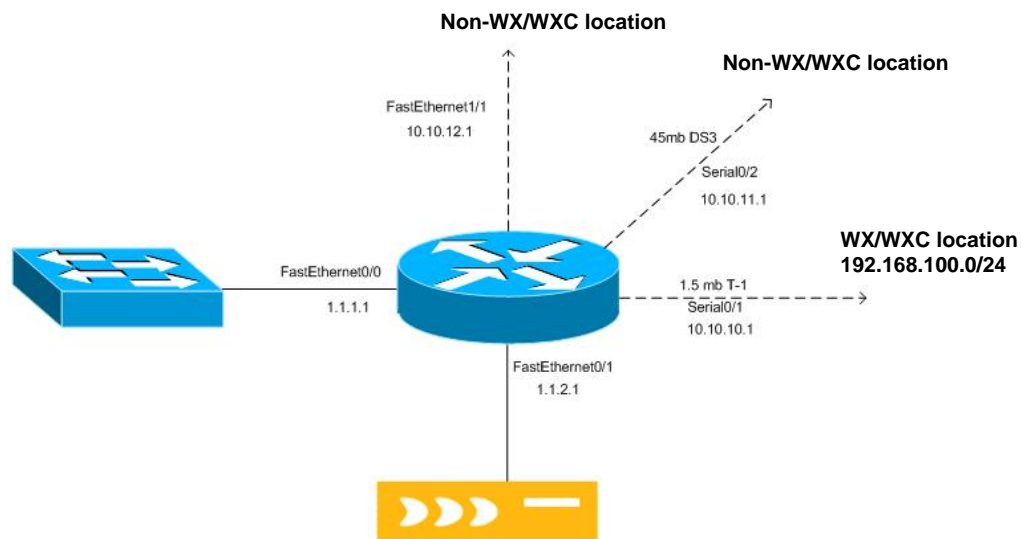
! configures wccp to match and redirect traffic based on
access-list 101

interface FastEthernet1/0
 ip address 10.1.1.1 255.255.255.0
interface FastEthernet2/0
 ip address 10.2.2.1 255.255.255.0

! applies inbound redirection on this interface
ip wccp 60 redirect in
ip wccp 70 redirect in

! access list created to permit all traffic except traffic
destined for 20.1.1.1
access-list 101 deny ip any host 20.1.1.1
access-list 101 permit ip any any
    
```

### Basic Outbound WCCP Configuration



Typical distribution site scenario using WCCP – traffic is received on all interfaces of the router with a percentage destined for the remote WX/WXC device location. This portion of traffic is to be redirected to the off-path WX/WXC device. Since there are multiple sites with high speed connections without WX/WXC devices and only one low speed connection, going to a WX/WXC device outbound redirection is used. If inbound redirection were used, all traffic received would have to be checked to see if

it should be forwarded to the WX/WXC device.

### WX/WXC Device

```
Packet Intercept Mode = WCCP
Router IP address = 1.1.2.1
TCP = service-id 60
UDP = service-id 70
```

### Router

```
ip wccp 60 redirect-list 101
ip wccp 70 redirect-list 101
interface FastEthernet0/0
  ip address 1.1.1.1 255.255.255.0

interface FastEthernet0/1
  ip address 1.1.2.1 255.255.255.0

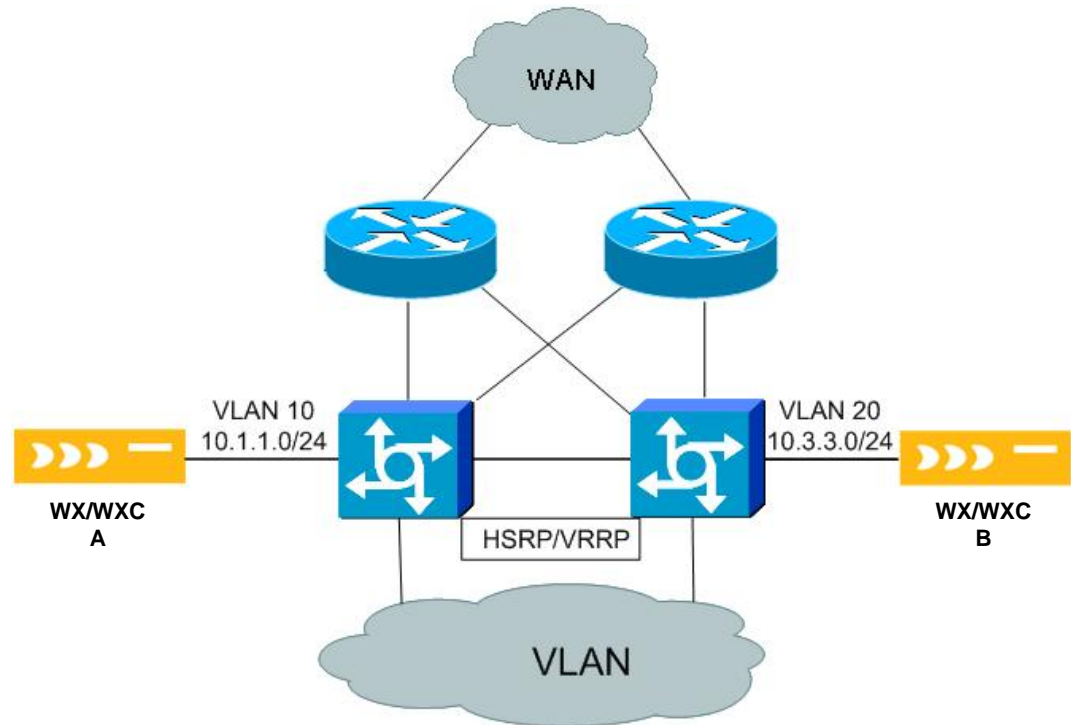
Interface Serial0/1
  ip address 10.10.10.1 255.255.255.0
! apply redirection on interface
  ip wccp 60 redirect out
  ip wccp 70 redirect out

interface Serial0/2
  ip address 10.10.11.1 255.255.255.0

interface FastEthernet1/1
  ip address 10.10.12.1 255.255.255.0

access-list 101 permit ip any 192.168.100.0 0.0.0.255
```

## High Availability using WCCP



Using WCCP in the above redundant topology is also a supported configuration; however, the WX/WXC device does not currently support multiple routers in a cache cluster so each WX/WXC device would function only with the router it is configured for

### WX/WXC Device A

```

Packet Intercept Mode = WCCP
Router IP address = 10.1.1.1
TCP = service-id 60
UDP = service-id 70
    
```

### WX/WXC Device B

```

Packet Intercept Mode = WCCP
Router IP address = 10.3.3.1
TCP = service-id 60
UDP = service-id 70
    
```



**L3 Switch 1**

```
ip wccp 60 redirect-list 101
ip wccp 70 redirect-list 101

interface Vlan 10
 ip address 10.1.1.1 255.255.255.0

interface Vlan 40
 ip address 10.4.4.253 255.255.255.252
interface Vlan 30
 ip address 10.2.2.1 255.255.255.0
 standby 1 priority 105
 standby 1 ip 10.2.2.10
 ip wccp 60 redirect in
 ip wccp 70 redirect in

access-list 101 permit ip any 192.168.100.0 0.0.0.255
```

**L3 Switch 2**

```
ip wccp 60 redirect-list 101
ip wccp 70 redirect-list 101

interface Vlan 20
 ip address 10.3.3.1 255.255.255.0

interface Vlan 40
 ip address 10.4.4.254 255.255.255.252

interface Vlan 30
 ip address 10.2.2.2 255.255.255.0
 standby 1 priority 100
 standby 1 ip 10.2.2.10
 ip wccp 60 redirect in
 ip wccp 70 redirect in
access-list 101 permit ip any 192.168.100.0 0.0.0.255
```

## Route Injection

### Operational Details for Route Injection

#### Basic Operation:

Route injection uses RIPv2 as a dynamic redirection method for Packet Interception. Before you stop reading because we mentioned RIP, you should know we are not recommending that you use RIPv2 as your interior routing protocol. We only use RIP to advertise reachability for the subnets being reduced. The attached router should only listen to RIPv2 and only on the interface attached to the WX/WXC device. We also suggest setting the timers lower so that in case of a failure the routes will age out in 30 seconds or less.

In order for the router to redirect traffic to the WX/WXC device, you must configure the RIP protocol to have a lower administrative distance so those routes advertised by RIP are preferred by the router. The router will then forward traffic for WX/WXC-advertised subnets to the WX/WXC device for reduction.

This probably still sounds a little strange, so we will explain how the WX/WXC device determines which routes to advertise via RIPv2. After tunnel establishment, the WX/WXC devices exchange reduction subnet information, notifying each other of the subnets configured for reduction. These are called remote routes to the receiving WX/WXC device. The WX/WXC device will use RIPv2 to advertise itself as the next hop gateway for all of its remote routes learned during this exchange to the attached router. Traffic for these subnets will now be directed to the WX/WXC device for reduction. . By default, the WX/WXC device will advertise a metric of two (2) for all networks learned during the exchange. The metric is a configurable option and is applied to all networks advertised by the WX/WXC device. The CLI command to configure the administrative distance: *config packet-interception rip set metric <1-15>*

When using Route Injection mode, the host address for each remote WX/WXC device will be “carved out” from the remote subnets learned (RIPv2 allows variable length subnets). This allows the WX/WXC devices to communicate with each other. If this was not done and the WX/WXC devices advertised their local subnets for reduction, the local router would send all traffic destined to the remote WX/WXC device directly to the local WX/WXC device which would then send it back to the router so communication would be impossible.

Because of this ‘carving’ the route table of the attached router will show multiple entries for advertised subnets using various subnet masks to exclude the host address for each remote WX/WXC device. Here is an example of what the route table might look like. Here 10.1.14.50 is the off-path WX/WXC device, and the IP address of the remote WX/WXC device (10.1.203.50) has been carved out.

```
10.1.0.0/16 is variably subnetted, 24 subnets, 9 masks
R 10.1.203.128/25 [30/2] via 10.1.14.50, 00:00:23, Ethernet0/1
R 10.1.203.51/32 [30/2] via 10.1.14.50, 00:00:23, Ethernet0/1
R 10.1.203.48/31 [30/2] via 10.1.14.50, 00:00:23, Ethernet0/1
```

```
R 10.1.203.52/30 [30/2] via 10.1.14.50, 00:00:23, Ethernet0/1
R 10.1.203.56/29 [30/2] via 10.1.14.50, 00:00:23, Ethernet0/1
R 10.1.203.32/28 [30/2] via 10.1.14.50, 00:00:23, Ethernet0/1
R 10.1.203.0/27 [30/2] via 10.1.14.50, 00:00:23, Ethernet0/1
R 10.1.203.64/26 [30/2] via 10.1.14.50, 00:00:23, Ethernet0/1
```

#### Advantages over WCCP or External Mode:

- **Ease of management:**

Route injection only injects routes for subnets that are in its remote route table. When new routes are added at a remote location no changes in local policies are required. In external mode or in WCCP mode if using access-lists, new access-lists may need to be created to allow this traffic.

- **High Availability networks:**

Using Route Injection it is simple to have several routers direct traffic to a single WX/WXC device. This simplifies design and minimizes any asymmetric routing concerns.

- **Simple backup scenario:**

By adding a second WX/WXC device that advertises the same routes via RIPv2 with a higher metric you have a backup solution using well understood protocols and behavior.

#### Failover and Recovery:

Route Injection recovery is dependant on multiple factors:

- **Link Failure:**

If the WX/WXC device is connected to a directly attached router port and the link fails, failover will occur immediately. If the WX/WXC device is connected via an L2 switch then failure time will depend on the RIPv2 timer's configuration.

- **Protocol Failure:**

Typically, recovery occurs in 15-30 seconds based on the setting of the following timers. Basic settings are usually 5 seconds for route updates and 15 seconds for invalid and hold-down timers (three (3) times the update timer). Flush is set for 30 seconds. Flush can be a smaller value than 30 seconds, but should always be larger than both the invalid and hold-down timers. You must make sure both the router and the WX/WXC device have the same update timer value or routes will flap.

To configure this on a Cisco router:

```
router rip
    timers basic 5 15 15 30
```

To configure the update timer through the WX/WXC CLI:

```
config packet-interception rip set update-timer 5
```

**Pass-through Traffic Handling:**

In Route Injection mode if traffic redirected to the WX/WXC device does not match any of its criteria for processing it is considered to be pass-through traffic. Pass-through traffic in Route Injection mode must still be tunneled to get to the remote side. If we did not do so it would be sent right back to the WX/WXC device because it is the next hop gateway for the route. Pass-through traffic is not reduced, just tunneled.

One side effect of this is traceroutes will not show the intermediate hops between the WX/WXC devices for subnets being reduced. Also, currently WX/WXC devices do not respond to traceroute, so they will appear invisible.

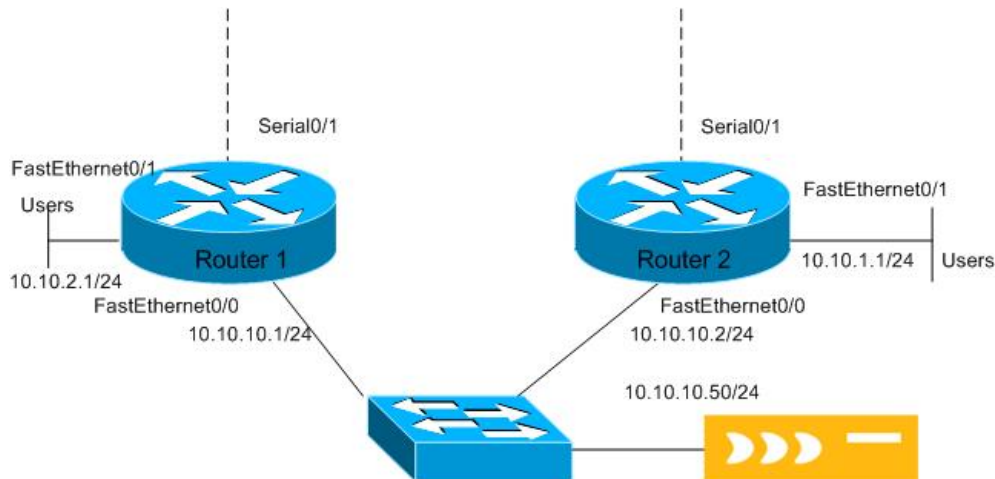
**Design Considerations:**

- Place the WX/WXC device directly on dedicated L3 interface in its own subnet for best fault isolation and simplicity.
- RIPv2 must have the same timers configured on the WX/WXC device and the Router or route instability may occur.
- Traceroute does not show intermediate hops between WX/WXC devices or the WX/WXC devices themselves.
- All Pass-through traffic must be tunneled.
- RIPv2 will dynamically identify a device as down after 30 seconds.
- When new routes are added for reduction at remote locations changes are automatically propagated via RIPv2. This prevents having to change local policies on all routers involved whenever a change is made.
- In High Availability environments it is simpler to direct traffic from multiple routers to the WX/WXC device for reduction.
- Simple backup configuration by having second WX/WXC device advertise subnets with a higher metric.
- If using EIGRP as the interior routing protocol. Some additional planning may be required for deployment. See the “Additional Network Considerations” section.
- The interface connected to the WX/WXC device and router should always be full duplex. This is not a hard requirement, but the amount of data on this connection is doubled since it goes in and comes out the same interface.
- Disable reverse path forwarding on the interface attached to the WX/WXC device.

**Additional Information:**

If you are using RIP as your interior routing protocol you can still use Route Injection, but careful planning should be done to ensure stability. Typically, a few route filters may be needed and/or offset lists to adjust the metrics for specific routes or route sources.

### Example 1: Basic Route Injection Configuration



This is a basic Route Injection configuration. Often Route Injection is used to direct traffic from multiple routers to a single WX/WXC device. In this case all traffic from Router 1 and Router 2 that is destined to a site with a WX/WXC advertising routes will be redirected to the local WX/WXC device.

#### WX/WXC Device

Packet Interception Mode = Route Injection

CLI Commands :

```
config packet-interception rip set update-timer 5
commit
```

#### Router1

```
router rip
version 2
! lower RIP flush timer to age routes faster- timers basic
5 15 15 30
network 10.0.0.0
! Specify RIP admin distance to be lower than all other
protocols (such as OSPF):
distance 30
```

```
! Disable auto-summarization of routes:
no auto-summary
! prevent RIP from broadcasting out all interfaces.
passive-interface default
interface fastethernet 0/0
ip address 10.10.10.1 255.255.255.0
no ip redirects
interface fastethernet 0/1
ip address 10.10.2.1 255.255.255.0
! Users on this interface.
```

### Router 2

```
router rip
version 2
! lower RIP flush timer to age routes faster- timer
parameter order: update invalid holddown ! flush
timers basic 5 15 15 30
network 10.0.0.0
! Specify RIP admin distance to be lower than all other
protocols (such as OSPF):
distance 30
! Disable auto-summarization of routes:
no auto-summary
passive-interface default
interface fastethernet 0/0
ip address 10.10.10.2 255.255.255.0
no ip redirects

interface fastethernet 0/1
ip address 10.10.1.1 255.255.255.0
! Users on this interface.
```

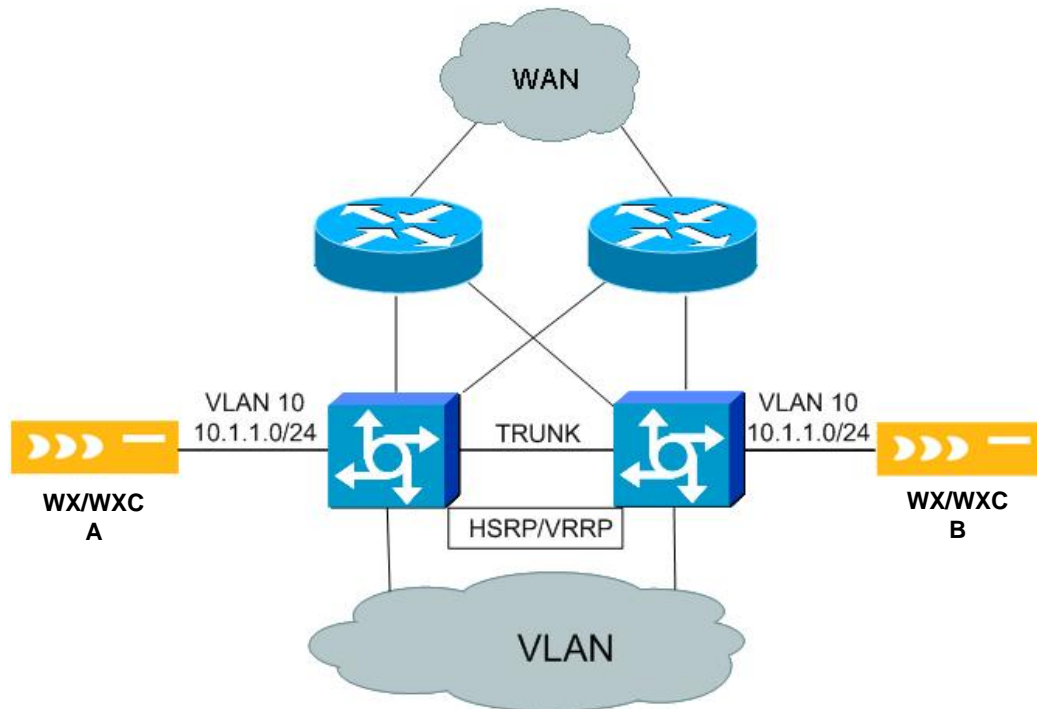
### Example: Two Typical Route Injection Redundant Configurations

Route Injection is commonly used to in networks that have many redundant connections or devices. By using RIPv2 to dynamically redirect traffic to the WX/WXC devices. Route Injection provides a simple method to direct traffic from multiple devices. Additionally it does not require reconfiguring all the routers with new policies as additional sites are added to the WX/WXC network.

In the example below the core routers/switches have a common VLAN and trunk port(s) between them. WX/WXC device A and B both advertise the remote subnets learned via RIPv2. Since we want the traffic to be focused to a single WX/WXC

device for better control. We have WX/WXC device A advertise the default distance of 2 for all routes learned and have WX/WXC device B advertise them with a distance of 5. In the event of a failure by WX/WXC device A, WX/WXC device B will become the best path via RIP and the traffic will be redirected automatically.

To provide faster convergence than the typical RIP timers, we can adjust the update timer interval to 5 seconds on the WX/WXC devices and routers. Then set the router to flush routes after 30 seconds.



Sample device configurations are below:

#### WX/WXC Device A

Packet Intercept Mode = Route Injection

IP Address: 10.1.1.50/24

CLI Commands:

```
config packet-interception rip set update-timer 5
commit
```

#### L3 Switch 1

```
interface Vlan 10
ip address 10.1.1.1 255.255.255.0
```

```
router rip
```

```
version 2
timers basic 5 15 15 30
network 10.0.0.0
no auto-summary
distance 30
passive-interface default
```

### **WX/WXC Device B**

Packet Intercept Mode = Route Injection

IP Address: 10.1.1.51/24

CLI Commands :

```
config packet-interception rip set update-timer 5
config packet-interception rip set metric 5
commit
```

### **L3 Switch 2**

```
interface Vlan Peribit B
ip address 10.1.1.2 255.255.255.0

router rip
version 2
! timer parameters in the following order: update invalid
holddown flush
timers basic 5 15 15 30
network 10.0.0.0
no auto-summary
distance 30
passive-interface default
```

## **Additional Network Considerations for Route Injection**

### **Route Validation**

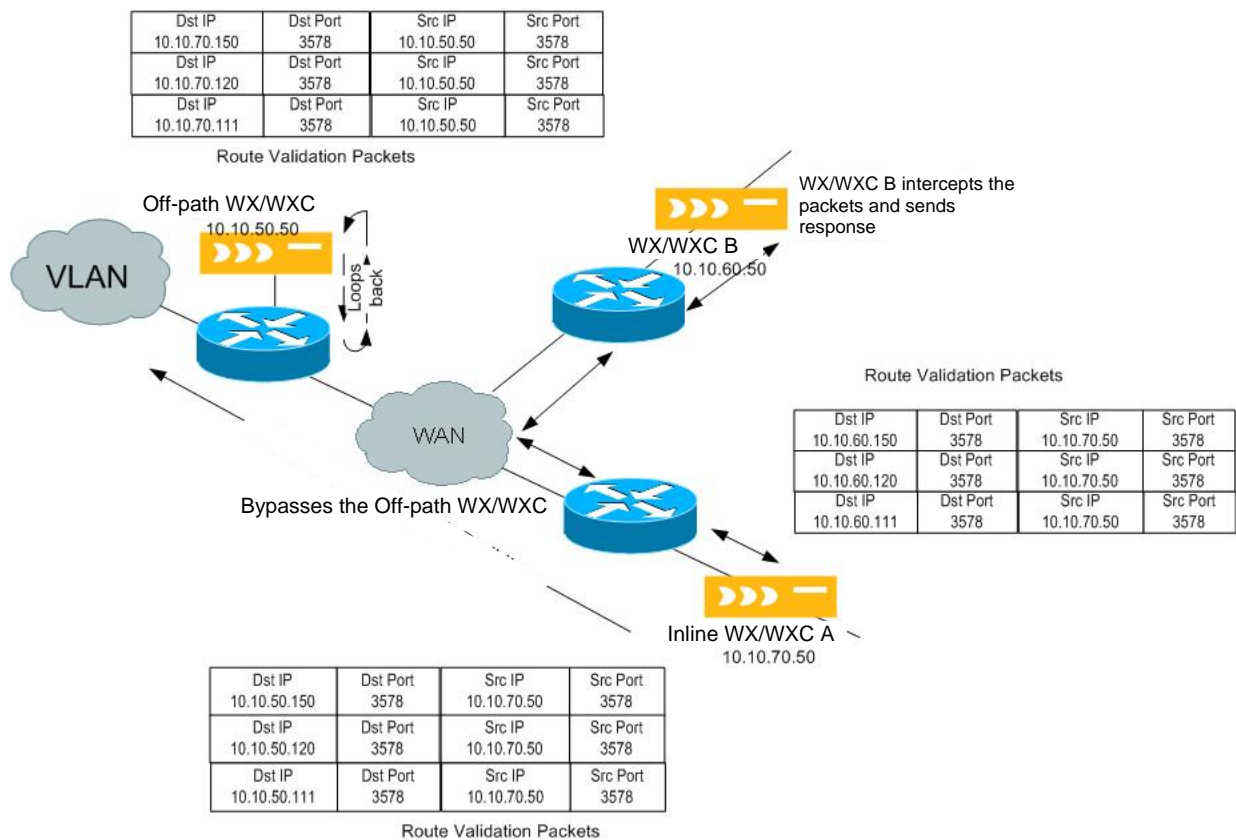
Route Validation is a feature that can be enabled to validate the path to the subnets advertised by remote WX/WXC devices. Each time remote routes are advertised or fetched, three probe packets are sent to three representative IP addresses in each advertised subnet. These packets are sent outside the WX/WXC tunnel so they can verify the natural routing path. If the remote WX/WXC device receives any of the probes, it discards the probes without forwarding them, and returns a report to the sending device. If a report is not received, the route is dropped from the remote routes.



If you have an existing inline deployment and are using Route Validation you will need to either disable Route Validation or not deploy off-path WX/WXC devices. Of course, if the WX/WXC devices doing Route Validation do not need to exchange traffic with the off-path WX/WXC device this does not matter. We will explain what happens when these two items are combined.

If an inline WX/WXC device is configured to use Route Validation and sends a validation packet to a subnet advertised by an off-path WX/WXC device. The validation packet would not be received by the WX/WXC device. Since validation packets are addressed to arbitrary host addresses on the remote subnet (not the WX/WXC device) the validation packets would bypass the off-path WX/WXC device and a response would not be returned to the sending WX/WXC device. This means this WX/WXC device will never be able to advertise its subnets for compression, because they are removed. See diagram below.

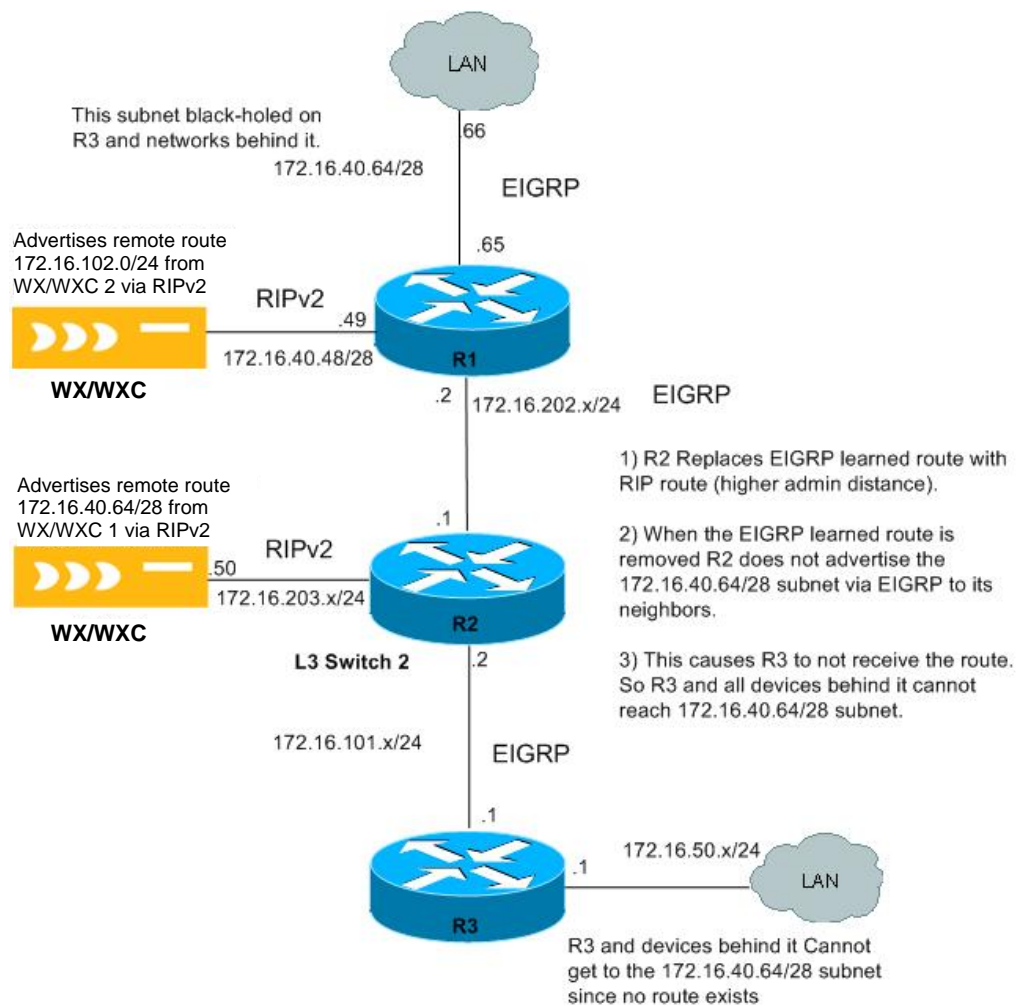
It is possible to enable outbound Route Validation on an off-path WX/WXC device in most cases, but this is generally not a good idea especially if running Route Injection. Since Route Injection uses RIPv2 advertisements to advertise itself as the next hop gateway to the remote subnets advertised by other WX/WXC devices, if it sends Route Validation packets outside the WX/WXC tunnel they will never reach the other side so all remote routes will be removed. For more details on why this is, see the Route Injection section.



### Route Injection Black Hole with EIGRP

In networks where EIGRP is in use and you are running WX/WXC devices in Route Injection mode some subnets may seem to have mysteriously gone into a 'black hole' and become unreachable. This can happen because EIGRP only exchanges local routing information with its neighbors. So if another route is learned with a better administrative distance it will replace the EIGRP route in the router. Since there is no EIGRP route, only a RIP route, the router will not advertise this route any longer to its neighbors. While the intent of Route Injection is to replace the route on the local router, preventing it from being advertised to other routers is not.

The diagram below and description of the issue following it will address how this happens and how it can be avoided.



### Potential Black Hole with EIGRP and Route Injection

In the diagram R1, R2 and R3 all use EIGRP to exchange routing information. WX/WXC device 1 and WX/WXC device 2 both transmit routing information using RIPv2 to R1 and R2 respectively. So R1 and R2 both run RIPv2 in listen only mode with a lower administrative distance than EIGRP, so RIPv2 routes will be preferred. By default they do not redistribute these routes to EIGRP and Auto Summary is disabled.

WX/WXC device 2 advertises the subnet that WX/WXC device 1 has identified for reduction 172.16.40.64/28 to R2 via RIPv2. Before WX/WXC device 2 advertised this subnet it was advertised by R1 via EIGRP. R2 then advertised this subnet to R3. Once R2 learns the route via RIPv2 it will use that as a preferred route to that subnet. When it does this it removes the EIGRP route and then stops advertising it to R3. This means R3 and all devices behind it will lose connectivity to the subnet.

To work around this issue there are a few different options.

1. Use static route for these subnets on R3 pointing to R2.
2. Use route redistribution from RIPv2 to EIGRP. This can be a bit tricky in some environments and requires some thought before deploying.
  - B. In this case you want to make sure only this route will be redistributed to EIGRP using a distribute-list. Distribute lists allow you to filter route advertisements between protocols and neighbors.
  - C. You want to make sure you only advertise it to R3 (again using distribute-lists).
  - D. If an alternate route to this subnet exists, you will need to ensure the metrics are correct to avoid any routing issues.
3. While this may sound daunting it is a very common practice in many networks.
4. Other Alternatives: Use WCCP or External Mode with Policy Based Routing.

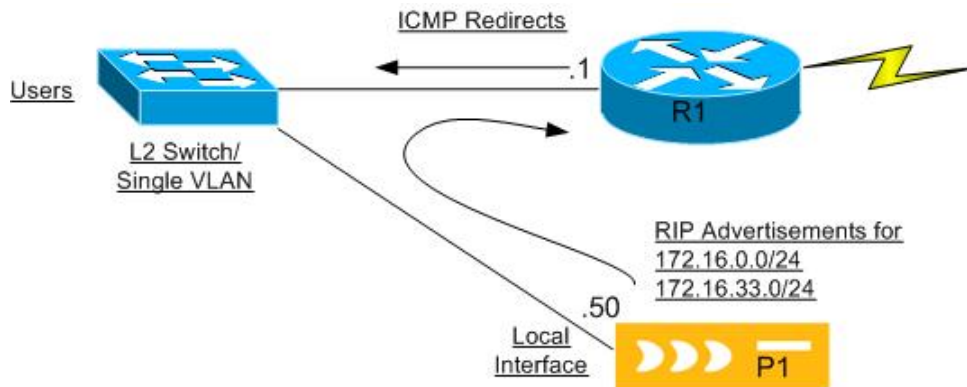
Okay, so why is this not a problem with OSPF? This issue does not occur with OSPF because in OSPF all routers in an area must know about all Link State Advertisements (LSA) so even though R2 will be using the RIPv2 route all routers below it would know that it was the correct gateway to get to the 172.16.40.64/28 subnet.

### ICMP Redirects and Route Injection

If users and the WX/WXC device performing Route Injection are on the same subnet/VLAN and the routers on that subnet have ICMP redirects enabled, then the following could happen.

When a user sends data destined for a subnet advertised by P1, the user sends that traffic to its default gateway R1. When that traffic is received by R1, it will correctly forward the data to P1 however; the router would also send an ICMP redirect to the user. This indicates that P1 is the better gateway for that destination. If the user replaces its gateway for this subnet with P1, this would cause connectivity issues for the user should the WX/WXC device fail. Connectivity to those networks for this user will be lost until the WX/WXC device is back up.

**Recommendation:** Place WX/WXC device on its own subnet or directly attach to the router. If not possible, disable icmp redirects on the router interface.



### Potential ICMP Redirect Issue with Layer 2

#### Packet Interception Mode and Reverse Path Forwarding

Packet Interception mode and routers that have reverse path lookup configured on them can cause data loss between sites running WX/WXC devices. Although the WX/WXC device can form tunnels across a Routed infrastructure where this feature is used, this is because the source IP of the packet(s) to establish a tunnel pass the "reverse path forwarding" criteria.

Because the WX/WXC device will re-inject packets into the network with a source IP addresses that matches an alternative interface. The problem can arise if you have Pass-Through traffic that is redirected back to the router, since it is now coming in on a different interface or when decompressing the traffic at the remote end of the connection.

#### So what is Reverse Path Forwarding?

When Reverse Path Forwarding is enabled on an interface, the router examines all packets received on that interface. The router checks to make sure that the source address appears in the routing table and matches the interface on which the packet

was received. The router must have a matching FIB entry for that source host on the interface where the packet was received and if it does not the packet will be dropped.

The key point is to ensure this feature is not enabled on the interface connected to a WX/WXC device. If this feature is enabled on a Cisco router you can use the following command on the interface connected to the WX/WXC device to disable it. *no ip verify unicast reverse-path.*

### **Packet Interception Mode - WCCPv2 and Cache Clusters**

WCCPv2 has the ability to allow multiple routers can use WCCPv2 to service a cache cluster, using multicasts. The WX/WXC device's Packet Intercept WCCP implementation does not currently support a "cache cluster" model.

### **Packet Interception WCCP and Multi-Path**

Multipath cannot be enabled when using WCCP mode. There is no support for the secondary IP address required for multi-path in this mode.

### **Using Default Assembler with Packet Interception Mode**

#### **WCCP and External Mode:**

Default assembler is fully supported with External or WCCP Packet Interception modes.

#### **Route Injection Mode:**

Default assembler will not have any effect if enabled on a WX/WXC device running Route Injection. The WX/WXC device will not advertise a default route which is essentially what a default assembler is.

Why not advertise a default route? Route Injection mode uses RIPv2 to direct traffic to itself. We can only advertise subnets that we know a remote WX/WXC device exists for. Since the router depends on the default route to get traffic to the other side of the network presumably and we use the router to get to the other side of the network. It cannot point to us for a default route when we need the router to forward packets to the remote destinations using this same route. When in Route Injection mode the WX/WXC device cannot send pass-through traffic back to the router to reroute (as is possible when using WCCP or External Mode). Doing so would cause the packets to loop between the router and WX/WXC device until the TTL becomes 0

**Note:** Remote WX/WXC devices can designate a WX/WXC device configured for Route Injection as their default assembler.

#### **Route Injection and Route Import or Route Polling:**

Route injection is currently not compatible with Route Import or Route Polling.

---

Copyright © 2006, Juniper Networks, Inc. All rights reserved. Juniper Networks and the Juniper Networks logo are registered trademarks of Juniper Networks, Inc. in the United States and other countries. All other trademarks, service marks, registered trademarks, or registered service marks in this document are the property of Juniper Networks or their respective owners. All specifications are subject to change without notice. Juniper Networks assumes no responsibility for any inaccuracies in this document or for any obligation to update information in this document. Juniper Networks reserves the right to change, modify, transfer, or otherwise revise this publication without notice.