



Performance Routing (PfR) Master Controller Redundancy Configuration

This application note provides an overview on how to configure a Performance Routing (PfR) master controller in a redundant configuration. Additional information and more detailed examples and command output from both deployment scenarios can be found in the *Transport Diversity: Performance Routing (PfR) Design Guide* at www.cisco.com/go/srnd.

There are two deployment models shown in this application note:

- [Dedicated Master Controller Topology, page 3](#)
- [Collocated Master Controller Topology, page 8](#)

Overview

This section provides an overview about the dedicated and collocated master controller topologies.

Dedicated Master Controller

The dedicated master controller topology is typically associated with a campus hub location where deploying routers solely for the purpose of supporting the PfR master controller function is a best practice to support and scale a deployment where a high number (more than several hundred) network prefixes is expected. This means that a primary and a standby routers are dedicated to supporting the master controller function. This means that the master controller function is supported by a dedicated primary router and a second router is configured as a dedicated standby master controller.



Corporate Headquarters:
Cisco Systems, Inc., 170 West Tasman Drive, San Jose, CA 95134-1706 USA

Copyright © 2008 Cisco Systems, Inc. All rights reserved.

Collocated Master Controller

The collocated master controller topology is a common deployment in a large branch office environment where two WAN routers support the branch location. Each WAN router has two or more WAN links defined as PfR external interfaces. The term *collocated* refers to the fact that one border router supports the primary master controller configuration and a second border router is configured as the standby master controller.

Assumptions

These assumptions are applicable to both deployment models:

- The reader is familiar with implementing PfR and the focus of this document is to provide an overview of how to configure a PfR master controller in a standby configuration. This enables WAN optimization to function in the event the primary master controller fails or is taken out of service for maintenance.
- The [Collocated Master Controller Topology](#) section consists simply of a topology diagram and configuration example. The concepts from the [Dedicated Master Controller Topology](#) section should be read and understood before reviewing the [Collocated Master Controller Topology](#) section.

More detailed output and validation of these deployments is available in the *Transport Diversity: Performance Routing (PfR)* at www.cisco.com/go/srnd.

PfR requires authentication of border routers and master controllers is accomplished by the use of a key-chain. A sample configuration is shown below. This required configuration is not shown the configuration examples for the sake of brevity.

```
!  
key chain NGWAN  
  key 10  
    key-string cisco  
!
```

The key-string value of **cisco** is a lab standard and not recommended for a customer deployment.



Tip

In the event the PfR border router loses contact with the master controller, the border router stops managing prefixes or applications. In other words, the fail-safe mode of PfR is to remove any routes injected into the IP routing table or BGP table and to stop policy routing of applications if the master controller is unavailable.

Dedicated Master Controller Topology

To provide availability of the master controller function, a backup (or standby) master controller can be implemented using a Hot Standby Router Protocol (HSRP) virtual IP address as the destination IP address configured in the border routers to identify the master controller. The implementation of PfR is such that the master controller listens on TCP port 3949 (by default, the port number is configurable) and the border routers initiate the TCP connection.

If the border routers lose contact with the master controller, the border routers remove PfR controlled routes and fallback to parent route control. The border routers continuously initiates a connection to the master controller IP address until a session is established.

To implement a standby master controller, do the following:

-
- Step 1** Configure the PfR master controller configuration identically on both primary and standby.
 - Step 2** Implement HSRP on the two master controllers.
 - Step 3** Configure the border routers to use the HSRP virtual IP address to identify the master controller.
-



Note The standby master controller functions in a stateless failover mode. However, standalone (dedicated) master controllers are not in the data plane switching path, so user traffic is not disrupted; it is simply not optimized during failover and recovery.

The general guideline is the standby master controller is effectively on-line and operational beginning with the monitor period following the failure. Therefore, given a 1-minute monitor period, prefixes can be managed within 2 to 3 minutes following a failure.

More specific implementation details are described in the following section.

Operational Overview

In this sample configuration, the master controller is configured to instruct the border routers to learn traffic for a 1-minute interval (**monitor-period 1**) and initiate subsequent monitoring (**periodic-interval 0**) immediately.

The learned prefixes are maintained by the master controller for the period specified by the **expire after time** command value. The number of prefixes learned by the border routers in one interval is specified by the **prefixes** command under the **learn** statement in the master controller configuration. All prefixes learned are sorted by throughput and the highest (top) 'n' throughput prefixes are reported to the master controller.

PfR Keepalive Disabled

If the PfR keepalive is disabled (**no keepalive** configured under the **oer master** section), the state of the TCP session triggers the failover to the standby master controller. At the failure of the primary master controller, the standby master controller becomes the active HSRP router and takes over control of the virtual IP address. The border routers have an open TCP session with the virtual address on TCP port 3949. When the standby master controller receives a TCP packet destined for port 3949 (for which it is listening, but has no active session) from the border router, it sends a TCP RST (reset) in response. The

border routers log the fact the master controller is down, and immediately attempt to re-establish a TCP session to the virtual address on port 3949. In the testing environment, the border routers were able to re-establish communication with the backup master controller within approximately 5 seconds after receiving the TCP RST.

Network prefixes that are currently under control of an PfR border router (static or BGP) are removed from the IP routing or BGP table when the master controller is deemed down. The standby master controller does not have any knowledge of previously learned prefixes, but it begins accumulating the learned prefixes for every monitor period interval subsequent to the failure of the primary master controller.

**Tip**

If testing this failure condition, administratively shutting down the primary master controller interface does not demonstrate a true failure condition. The Cisco IOS sends a TCP RST for the active TCP connection on that interface to the border routers. The border routers respond by tearing down their TCP connection and retry approximately 5 seconds later to establish a connection to the master controller IP address.

PfR Keepalive Enabled

If PfR keepalives are enabled, these PfR specific keepalives are also used to detect a communication failure between the master controller and the border routers. The border router can detect a failure of the master controller after three keepalives are unanswered and deem the master controller down. Once the master controller is down, either because of the TCP RST message received from the standby master controller, or the loss of three keepalives, the next course of action by the border router is the same, attempt to establish a new TCP session with the master controller.

In testing a keepalive value of 1 second was configured and the master controller was declared down 3 seconds after the link was failed between border routers and master controller. The default keepalive value is 5 seconds. The hold-time (dead interval) is 3 times the configured keepalive value.

The use of PfR keepalive is most applicable to verifying communications with a single master controller. As the standby master controller sends a TCP RST following assuming control of the virtual IP address, the use of the PfR keepalive is not required as a recovery mechanism.

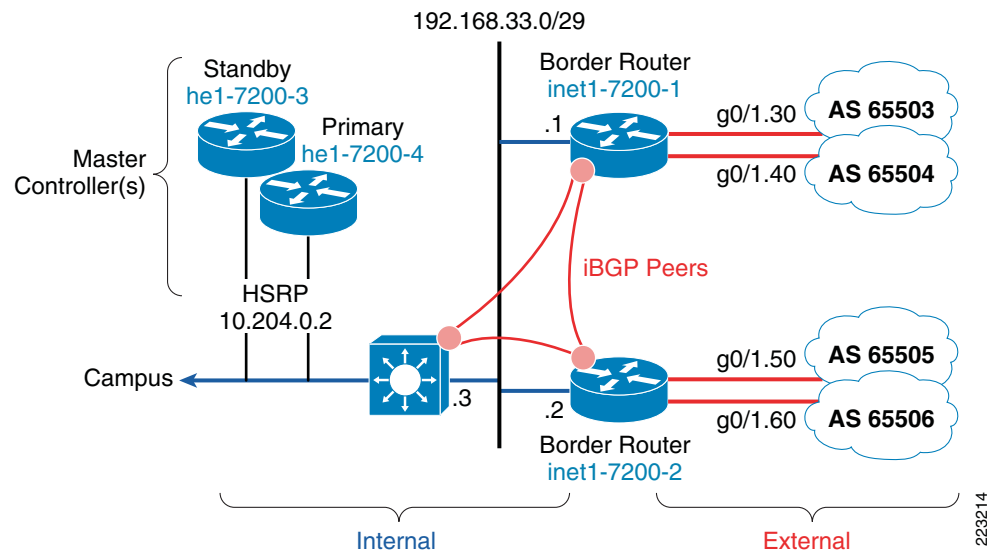
**Tip**

PfR keepalive can be disabled in this configuration. By doing so, it also eliminates the overhead associated with sending and responding to keepalives and can help in scaling.

Topology

Two border routers are implemented, each with two external (exit) links. A primary and standby master controller are implemented as shown in [Figure 1](#).

Figure 1 *Dedicated Master Controller Deployment Example*



A single Layer-3 campus switch is implemented in this test deployment. In most customer implementations, multiple Layer-3 campus switches are deployed for increased availability.

Master Controller Configuration

In the test environment, both primary and standby master controller are configured identically as far as the **oer master** portion of the configuration, which is shown as follows:

```
oer master
no keepalive
!
border 192.168.33.1 key-chain NGWAN
interface GigabitEthernet0/2 internal
interface GigabitEthernet0/1.30 external
interface GigabitEthernet0/1.40 external
!
border 192.168.33.2 key-chain NGWAN
interface GigabitEthernet0/2 internal
interface GigabitEthernet0/1.50 external
interface GigabitEthernet0/1.60 external
!
learn
throughput
delay

periodic-interval 0
monitor-period 1
prefixes 2500
expire after time 4
```

```

    aggregation-type prefix-length 22

mode route control
mode monitor passive
mode select-exit best
periodic 180
!
```

Primary Master Controller HSRP interface

The interface used as the target IP address by the border routers for the primary master controller is shown in the following example. Note that the standby IP address is 10.204.0.2 and this address is used in the border router configuration rather than the actual interface IP address of the primary master controller.

```

interface GigabitEthernet0/2
 ip address 10.204.0.4 255.252.0.0
 no ip redirects
 no ip unreachable
 . . .
 standby 0 ip 10.204.0.2
 standby 0 timers 3 10
 standby 0 priority 110
 standby 0 preempt
 hold-queue 4096 in
 hold-queue 4096 out
!
```

The configured standby priority of 110 is higher than the default value of 100, making this router the preferred active HSRP router.

The default value for HSRP timers is 3-second hello and 10-second dead (*standby timers 3 10*). Given this is a stateless failover configuration and these master controllers are not in the data path plane (they are not switching any packets, only managing the master controller database), there is no appreciable benefit to aggressively tune the HSRP timers to any value lower than the default values. In fact, increasing the HSRP hello and dead interval times is recommended so that a HSRP failover due to CPU busy on the primary master controller (a false positive switchover) is avoided. HSRP failover and recovery within 30 seconds to one minute is sufficient for most deployments.

The hold-queue values of 4096 for both input and output is recommended as all packets on these dedicated master controller routers is process switched. Buffer tuning is also advised to eliminate buffer misses and failures.

Standby Master Controller HSRP Interface

The standby master controller HSRP interface uses the default HSRP priority of 100 and has not been configured to preempt, otherwise, it is identical to the primary master controller configuration.

```

!
interface GigabitEthernet0/2
 ip address 10.204.0.3 255.252.0.0
 . . .

 standby 0 ip 10.204.0.2
 standby 0 timers 3 10
 hold-queue 4096 in
 hold-queue 4096 out
```

```
!
```

Border Router Configuration

Both border routers are identically configured as shown below:

```
!
border
 logging
  local GigabitEthernet0/2
  master 10.204.0.2 key-chain NGWAN
!
```

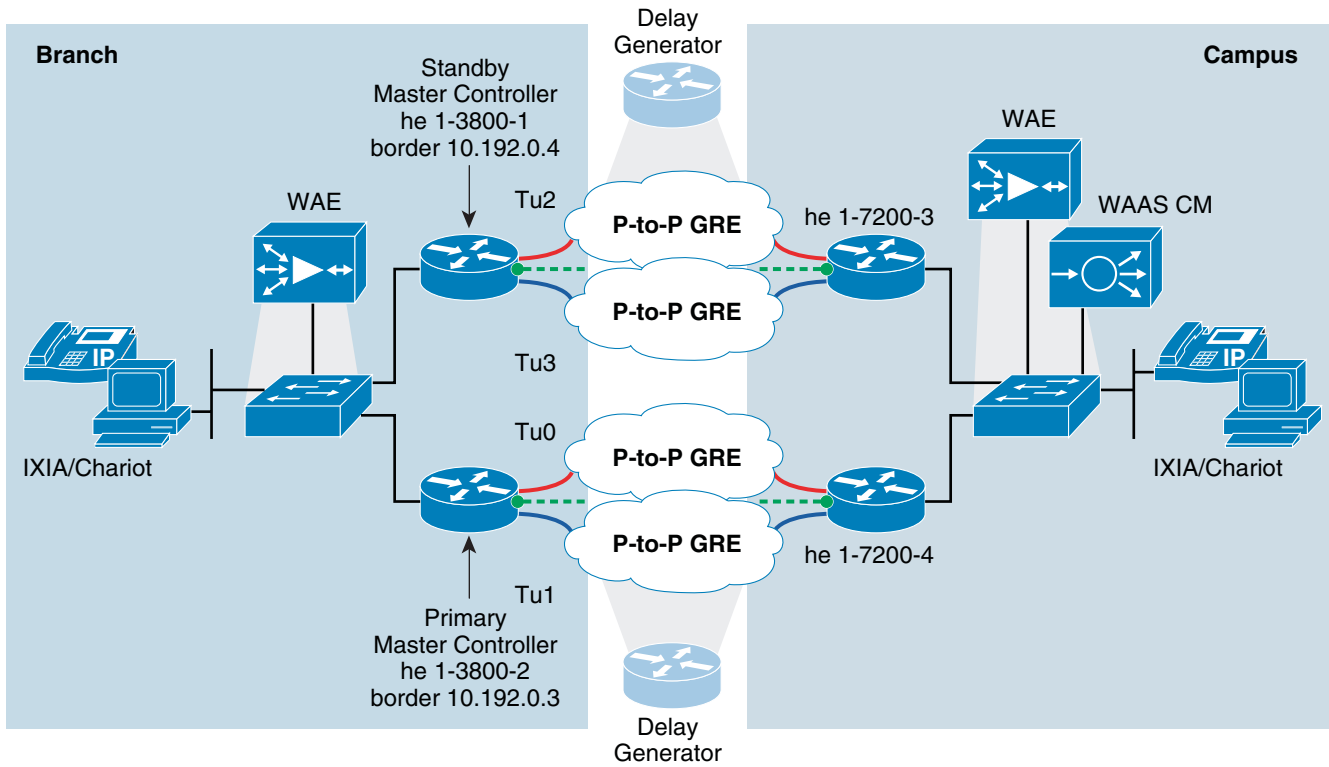
Summary

Implementing a dedicated standby master controller is a simple and effective means of providing stateless failover for the master controller function for a campus hub location. An additional advantage to doing so is that PfR keepalives can be disabled since the standby master controller can notify the border routers indirectly by way of a TCP RST message that the primary master controller has failed and the standby master controller has taken control of the virtual IP address.

Collocated Master Controller Topology

The collocated master controller topology is typically associated with a branch deployment model where the number of network prefixes or applications being managed are substantially less than at a campus hub location. [Figure 2](#) is a reference for the accompanying sample configuration below.

Figure 2 Collocated Master Controller Deployment Example



Configuration files for both branch routers are shown in the following section. The campus routers could be similarly configured as the branch routers, or a dedicated master controller topology could be implemented.

Assumptions

Assume both border routers are configured with HSRP and the virtual IP address is 10.192.0.1. The primary master controller is the active HSRP router. The standby master controller is the standby HSRP router.

Because PfR must have at least two external interfaces (exits) to function, the standby master controller/border router must have at least two exits. This is to meet the minimum of two exit requirement in the event the primary master controller / border router fails or is taken out of service. A typical deployment would have both border routers with two exits as shown in the topology diagram.

In the configuration examples, oer-map BRANCH is omitted for brevity.

Primary Master Controller and Border Router

The system image file is flash:c3845-adventerprisek9-mz.124-11.T3.bin.

```

!
hostname he1-3800-2
!
!
oer master
  policy-rules BRANCH

  logging
  !
  border 10.192.0.3 key-chain NGWAN
    interface GigabitEthernet0/1.1100 internal
    interface Tunnel0 external
    interface Tunnel1 external
  !
  border 10.192.0.4 key-chain NGWAN
    interface GigabitEthernet0/1 internal
    interface Tunnel2 external
    interface Tunnel3 external
  !
  learn
    throughput
    delay
    periodic-interval 0
    monitor-period 1
    prefixes 2500
    aggregation-type prefix-length 29
  mode route control
  mode monitor passive
  mode select-exit best
  periodic 180
!
oer border
  logging
  local GigabitEthernet0/1.1100
  master 10.192.0.1 key-chain NGWAN
!
!end

```

Standby Master Controller and Border Router

The system image file is flash:c3825-adventerprisek9-mz.124-11.T3.

```

!
hostname he1-3800-1
!
!
oer master
  policy-rules BRANCH
  no keepalive
  logging
  !
  border 10.192.0.4 key-chain NGWAN
    interface GigabitEthernet0/1 internal
    interface Tunnel2 external
    interface Tunnel3 external

```

```
!  
border 10.192.0.3 key-chain NGWAN  
  interface Tunnel0 external  
  interface Tunnell1 external  
  interface GigabitEthernet0/1.1100 internal  
!  
learn  
  throughput  
  delay  
  periodic-interval 0  
  monitor-period 1  
  prefixes 2500  
  expire after time 65000  
  aggregation-type prefix-length 29  
  mode route control  
  mode monitor passive  
  mode select-exit best  
  periodic 180  
!  
peer border  
  logging  
  local GigabitEthernet0/1  
  master 10.192.0.1 key-chain NGWAN  
!  
!  
end
```

Summary

By implementing a primary and standby master controller at branch locations, traffic optimization can continue to function in the event one of the two branch routers fails or is taken out of service for maintenance.