ERICSSON

Configuring Inter-Chassis Redundancy

SYSTEM ADMINISTRATOR GUIDE

14/1543-CRA 119 1170/1 Uen D

Contents

1 Contents of This Document

This document describes inter-chassis redundancy (ICR) failover support for PPPoE, Dynamic Host Configuration Protocol (DHCP), and clientless IP service selection (CLIPS) subscriber sessions. In an ICR configuration, if the active SmartEdge® router fails, or a card or port carrying a client session fails, ICR allows the subscriber to re-establish the session on the backup SmartEdge router. During the failover, all traffic packets for the failed client are lost.

2 Restrictions and Requirements

ICR failover is subject to the following SmartEdge hardware restrictions and requirements:

- Supported on all SmartEdge chassis.
- Supported on PPA2 Ethernet cards, but not on PPA2 ATM cards.

ICR failover is subject to the following configuration restrictions and requirements:

- Supported for DHCP clients in DHCP proxy server and internal DHCP server contexts, but not in DHCP relay server contexts.
- Supported for directly connected CLIPS clients in DHCP proxy server and internal DHCP server configurations.
- Supported for relay connected CLIPS clients in DHCP proxy server configurations.
- Not supported on link-group circuits.
- If DHCP client IP addresses are assigned by internal DHCP servers, IP address conflicts can occur between the active and backup SmartEdge chassis and this can only be prevented if the internal pools are set up so they do not overlap. Subscriber IP address pools assigned from external pools (through RADIUS, external DHCP, or both) do not have this problem.
- Some ICR configurations (DHCP proxy server and PPPoE) require multibind interfaces with identical IP addresses on active and backup SmartEdge chassis in the ICR setup:
  - All circuits using multibind interfaces with identical IP addresses must be protected with the `track vrrp` command.
  - Do not enable proxy ARP on multibind interfaces with identical IP addresses; otherwise, you will have multiple SmartEdge chassis responding to ARP requests.
- Not supported on circuit creation on-demand (CCoD) circuits.

A *monitor PVC* is a static 802.1Q PVC that is bound to a VRRP interface using the `bind interface` *`vrrp-if-name context-name`* command. The following configuration restrictions and requirements apply: to monitor PVCs:

- To support ICR failover of PPPoE, DHCP, and CLIPS sessions in the event of the failure of a monitored port, the PVCs carrying these sessions must be configured to track a monitor PVC using the `track vrrp` *`vrrp-id vrrp-if-name context-name`* command.
- You can configure a port to have more than one monitor PVC and partition the configured follower PVCs to track the different monitor PVCs.

- Only statically a configured 802.1Q PVC can be designated as a monitor PVC. ICR does not support designating a port-level monitor PVC or a non-802.1Q PVC as a monitor PVC.
- To protect against port level failures, the monitor and the follower PVCs must be configured over the same port. While ICR supports configuring monitor and the followers on the same card but on different ports, such a configuration does not protect against follower PVC port failures. ICR does not support configuring the monitor on one card and corresponding the followers on a different card.
- A monitor PVC cannot be configured under the following conditions:
  - An explicit range; that is, each monitor PVC must be assigned its PVC identifier individually rather than specifying *ID-n* `through`.
  - An on-demand CCoD range.
  - A transport range.
  - An access link-group.

3 Configuration and Operations Tasks

Use the following CLI commands to configure ICR.

Table 1 ICR Configuration Commands

Task	Root Command	Notes
Create a VRRP router, assign it a VRRP ID and enter VRRP configuration mode	*vrrp*	Configures the current interface as a VRRP interface. Enter the command with the `owner` keyword to configure the VRRP router as the owner of the VRRP interface IP address (VRIP) or with the `backup` keyword to configure it as a backup owner of the VRIP in the event of an ICR failover.
Configures the VRIP of current interface.	*virtual address*	Enter in VRRP configuration mode. The owner of the VRIP responds to ARP requests for this address. In the event of an ICR failover, the backup owner responds. In that way the owner and backup owner are redundant routers for the VRIP.
(Optional) Configures the VRRP election priority for the backup VRRP router.	*priority (VRRP)*	Enter in VRRP configuration mode. A backup VRRP router that has a higher priority value is preferred over a router that has a lower priority value. The priority of an owner VRRP router is 254.
(Optional) Configures the authentication of VRRP exchanges.	*authentication (VRRP)*	Enter in VRRP configuration mode.
See *Configuring VRRP* for additional commands that might be necessary for configuration of VRRP options.		

Table 1 ICR Configuration Commands

Task	Root Command	Notes
Set the IP address used as the giaddr in DHCP packets that the SmartEdge router relays to the external DHCP server.	*giaddr*	Use the `giaddr` command to set the IP address used as the giaddr in DHCP packets that the SmartEdge router relays to the external DHCP server: • The giaddr used by the DHCP proxy must be unique on each backup and active chassis in the ICR setup. • The *`ip-addr`* value must be one of the addresses in the subnet configured for the multibind interface. If this command is not entered, the SmartEdge router uses the primary IP address of the multibind interface as the giaddr. Use the `giaddr` command in DHCP-proxy configuration mode.
Specify a second IP address for the current interface that can be used for ICR operation. Enter this command with the `application` keyword.	*ip address (interface)*	This option allows ICR active and backup SmartEdge routers to use different IP addresses as the giaddr from identically addressed multibind interfaces. The different giaddr IP addresses are a requirement that makes it possible for the identically addressed multibind interfaces of multiple SmartEdge chassis in the ICR operation configuration to communicate with a DHCP server. The application address must be one of the addresses in the subnet configured for the multibind interface to be used for giaddr. Any ARP requests received for application addresses are dropped. Use the `ip address (interface)` command in DHCP-proxy configuration mode.

Table 1 ICR Configuration Commands

Task	Root Command	Notes
Change a static dot1q PVC into a track PVC.	*track (vrrp)*	Use the `track vrrp` command to configure a static 802.1Q PVC as a track PVC; that is, a PVC that tracks the state of a VRRP interface. The track PVC carries traffic when it is active and no traffic when it is in backup state. Only one track option can be defined on a PVC. When a static 802.1Q PVC is configured to track the state of a VRRP interface, all of its child circuits inherit this functionality and also track the state of that same VRRP interface. Use the no form of this command to change a track dot1q PVC into a static PVC. Use the `track vrrp` command in dot1q PVC configuration mode.

Table 1 ICR Configuration Commands

Task	Root Command	Notes
Enable setting subscriber-specific DHCP lease times with the internal DHCP server and DHCP proxy server.	*subscriber dhcp-lease idle-timeout*	Normally, the `idle` *`minutes`* option provides a session idle timeout value to the subscriber currently being configured; that is, the subscriber-specific idle timeout value. However, when the `subscriber dhcp-lease idle-timeout` command is entered for the context, the `idle` *`minutes`* option instead sets the subscriber-specific DHCP lease times. Alternatively, you can configure the subscriber lease time using your RADIUS database and providing a value for RADIUS Attribute #28-Idle Timeout. If you use your RADIUS database, you must also configure the SmartEdge router with the `subscriber dhcp-lease idle-timeout` command. DHCP split lease times apply when the SmartEdge router is configured as a DHCP proxy and the subscriber-specific DHCP lease times are set separately from the external DHCP server's setting for DHCP lease times. The value entered for the subscriber lease time should be less than 50% of the external DHCP server lease time for split lease to be effective. When you configure shorter subscriber-speci fic lease times, DHCP clients have a shorter failover time (to the backup ICR SmartEdge chassis) and allow the DHCP server to reclaim inactive clients' IP address leases more quickly **Caution:** Setting a too low lease time can adversely affect the session bring up rate as well as the recovery rate of DHCP and CLIPS sessions. A too low lease time might cause the SmartEdge router to lose subscribers IP leases. Use the `subscriber dhcp-lease idle-timeout` command in context configuration mode.

Table 1 ICR Configuration Commands

Task	Root Command	Notes
(Required in some configurations) Enable the interface to have multiple circuits bound to it.	*interface (context)*	Enter the `interface` command with the `multibind` keyword in context configuration mode. Some ICR configurations (DHCP proxy and PPPoE) require multibind interfaces with identical IP addresses on all VRRP chassis. All subscriber circuits using these interfaces must be protected with the `track vrrp` command.(1)
Statically bind an 802.1Q PVC to a previously created interface in the specified context.	*bind interface* `vrrp-if-name context-name`	Bind the PVC to a VRRP interface to create a VRRP *monitor PVC*.
(Optional) Specify the default lease time for this DHCP server or one of its subnets.	*default-lease-time*	Use the `default-lease-time` *`seconds`* command in DHCP server configuration mode. DHCP short lease times apply when the SmartEdge router is configured as an internal DHCP server. With short lease times, DHCP clients can detect the failure of the active VRRP node more quickly and thereby have a shorter failover time. **Caution:** Setting a too low lease time can adversely affect the session bring up rate as well as the recovery rate of DHCP and CLIPS sessions. A too low lease time might cause subscribers to lose their IP leases.
(Optional) Specify the maximum allowed time for the lease for this internal DHCP server or one of its subnets.	*max-lease-time*	Use the `max-lease-time` *`seconds`* command in DHCP server configuration mode.
(Optional) Configure the idle session timeout criteria for a subscriber session.	*timeout (subscriber)*	Enter the `timeout (subscriber)` command with the `idle` in subscriber configuration mode. This command sets the DHCP subscriber lease time when `subscriber dhcp-lease idle-timeout` is enabled.

(1) In configurations with identical IP addresses on all VRRP chassis, proxy ARP must not be enabled; otherwise multiple chassis will respond to ARP requests.

4 Configuration Examples

This section contains a example configuration of an owner VRRP router (Section 4.1 on page 11) and a backup VRRP router (Section 4.2 on page 14).

4.1 Example Configuration of an Owner VRRP Router in ICR Setup

This section shows the configuration of an owner VRRP router named **Pacific** that supports redundancy for DHCP, CLIPS, and PPPoE clients. To provide failover resiliency, a backup router is also required. The backup router is configured in Section 4.2 on page 14.

You can also configure a VRRP interface to have both owner and backup VRRP routers; however, each VRRP router must have a unique ID and unique VRIP.

In addition to the monitor PVC in this section, examples of other types of monitor PVCs are found in Section 4.4 on page 16, Section 4.3 on page 16, and Section 4.5 on page 17.

4.1.1 Active VRRP Router

See Section 4.1.2 on page 13 for notes describing the important ICR sections of this example.

```
service multiple-contexts
!
configure
!
context pctx
 domain test advertise
!
interface foo multibind
 ip address 10.0.0.1/16
 ip pool 10.0.0.0/16
!
interface one
 ip address 1.2.3.4/24
 vrrp 7 owner
  virtual-address 1.2.3.4
!
 aaa authentication subscriber none
!
 subscriber default
  ip address pool
!
```

```
context dhcp_ctx
!
 interface dhcp multibind
  ip address 50.0.0.1/16
  dhcp proxy 48000
   giaddr 50.0.0.1
!
 interface if1
  ip address 192.168.42.1/24
  vrrp 22 owner
   virtual-address 192.168.42.1
!
 interface to-server
  ip address 172.168.1.1/24
!
 aaa authentication subscriber none
!
 subscriber default
  dhcp max-addrs 1
!
 dhcp relay option
 dhcp relay server 172.168.1.100
!
port ethernet 1/1
 no shutdown
 encapsulation dot1q
 dot1q pvc 101
  bind interface one pctx
 dot1q pvc 108 encapsulation pppoe
  track vrrp 7 one pctx
  bind authentication chap context pctx maximum 2
!
port ethernet 1/5
 no shutdown
 encapsulation dot1q
 dot1q pvc 2
  bind interface if1 dhcp_ctx
 dot1q pvc 4
  track vrrp 22 if1 dhcp_ctx
  bind subscriber user1@dhcp_ctx
 dot1q pvc 5
  track vrrp 22 if1 dhcp_ctx
  service clips dhcp context dhcp_ctx
!
port ethernet 3/1
 no shutdown
 bind interface to-server dhcp_ctx
!
end
```

4.1.2 Notes

The following command lines are excerpts from the preceding full configuration example:

1. The following lines create two VRRP interfaces (named **one** and **if1**) and two VRRP routers (with IDs **7** and **22**). Each VRRP router is assigned a static IP address and a VRIP address. Both VRRP routers are owners of their VRIP addresses.

```
interface one
 ip address 1.2.3.4/24
 vrrp 7 owner
  virtual-address 1.2.3.4
!
 interface if1
  ip address 192.168.42.1/24
  vrrp 22 owner
   virtual-address 192.168.42.1
```

2. In the following section of the configuration, port **1/1** contains the Dot1q PVC **101**. When this PVC is bound to the VRRP interface **one**, it becomes a *monitor PVC* for the VRRP router **7**. When a failure event occurs in port **1/1**, the circuits tracked by the active VRRP router fail over to the corresponding circuits tracked by the backup VRRP router.

   You can configure a port to have more than one monitor PVC and partition the configured follower PVCs to track the different monitor PVCs.

```
port ethernet 1/1
 no shutdown
 encapsulation dot1q
 dot1q pvc 101
  bind interface one pctx
```

3. The following lines configure PVC **108** to carry PPPoE subscriber traffic. The `track` command specifies that its failover behavior follows the state of VRRP interface **one** of the VRRP router **7**.

```
dot1q pvc 108 encapsulation pppoe
  track vrrp 7 one pctx
  bind authentication chap context pctx maximum 2
```

4. The following lines create the monitor PVC **2** and bind it to VRRP interface **if1**.:

```
port ethernet 1/5
 no shutdown
 encapsulation dot1q
 dot1q pvc 2
  bind interface if1 dhcp_ctx
```

5. The following lines create PVC **4** for the DHCP client and PVC **5** for CLIPS client(s) to carry subscriber traffic. The `track` command specifies the failover behavior of the link is determined by the state of VRRP interface **if1** of the VRRP router **22**:

```
dot1q pvc 4
  track vrrp 22 if1 dhcp_ctx
  bind subscriber user1@dhcp_ctx
 dot1q pvc 5
  track vrrp 22 if1 dhcp_ctx
  service clips dhcp context dhcp_ctx
```

4.2 Example Configuration of Backup Router in ICR Setup

The SmartEdge chassis **Atlantic** is configured with backup VRRP routers in which subscriber circuits are tracked. These subscriber circuits serve as failover backup for the subscriber circuits tracked by the corresponding active router **Pacific**.

The ICR configuration parameters of the backup SmartEdge chassis are identical to the active SmartEdge router except for the giaddr and DHCP server interface addresses.

```
service multiple-contexts
!
context pctx
 domain test advertise
!
 interface foo multibind
  ip address 10.0.0.1/16
  ip pool 10.0.0.0/16
!
 interface one
  ip address 1.2.3.5/24
  vrrp 7 backup
   virtual-address 1.2.3.4
!
 aaa authentication subscriber none
!
 subscriber default
  ip address pool
!
```

```
context dhcp_ctx
!
 interface dhcp multibind
  ip address 50.0.0.1/16
  ip address 50.0.0.2 application
  dhcp proxy 48000
   giaddr 50.0.0.2
!
 interface if1
  ip address 192.168.42.2/24
  vrrp 22 backup
   virtual-address 192.168.42.1
!
 interface to-server
  ip address 172.168.1.2/24
!
 aaa authentication subscriber none
!
 subscriber default
  dhcp max-addrs 1
!
 dhcp relay option
 dhcp relay server 172.168.1.100
!
port ethernet 1/1
 no shutdown
 encapsulation dot1q
 dot1q pvc 101
  bind interface one pctx
 dot1q pvc 108 encapsulation pppoe
  track vrrp 7 one pctx
  bind authentication chap context pctx maximum 2
!
port ethernet 1/5
 no shutdown
 encapsulation dot1q
 dot1q pvc 2
  bind interface if1 dhcp_ctx
 dot1q pvc 4
  track vrrp 22 if1 dhcp_ctx
  bind subscriber user1@dhcp_ctx
 dot1q pvc 5
  track vrrp 22 if1 dhcp_ctx
  service clips dhcp context dhcp_ctx
!
port ethernet 3/1
 no shutdown
 bind interface to-server dhcp_ctx
!
end
```

4.3 Single Monitor PVC Example

The following example shows the configuration of the 802.1Q PVC `200:8` as a monitor PVC:

```
configure
 port ethernet 14/1
  encapsulation dot1q
   dot1q pvc 200 encapsulation 1qtunnel
    dot1q pvc 200:8
    bind interface vrrp-if1 local
! CVLAN 200:8 is now a monitor PVC
```

4.4 Multiple Monitor PVCs and Multiple Track PVCs Example

The following example shows the configuration of two 802.1Q PVCs as monitor PVCs for a single port that is partitioned into two interfaces. The VRRP instance `1` is configured under the `vrrp-1` interface and VRRP instance `100` is configured under the `vrrp-100` interface:

```
configure
 port ethernet 1/1
  encapsulation dot1q
  dot1q pvc 1
   bind interface vrrp-1 local
! PVC 1 is a monitor pvc of port 1/1
   exit
  dot1q pvc 2
   track vrrp 1 vrrp-1 local
   bind authentication chap context local
   exit
  dot1q pvc 3
   track vrrp 1 vrrp-1 local
   bind authentication chap context local
! PVCs 2 and 3 are track PVCs of interface vrrp-1 on port 1/1
! PVCs 2 and 3 are PPPoE circuits
   end
!
!
configure
 port ethernet 1/1
  encapsulation dot1q
  dot1q pvc 101
   bind interface vrrp-100 local
! pvc 101 is also a monitor pvc of port 1/1
   exit
  dot1q pvc 102
   track vrrp 100 vrrp-100 local
   service clips dhcp context local
   exit
  dot1q pvc 103
   track vrrp 100 vrrp-100 local
   service clips dhcp context local
! PVCs 102 and 103 are track PVCs of interface vrrp-100, port 1/1
! PVCs 102 and 103 are CLIPS enabled circuits
```

4.5 SVLAN Monitor PVC with Multiple CVLAN Track PVCs Example

The following example shows the configuration of a 802.1Q PVC as a monitor PVC where . VRRP instance 1 is configured under `vrrp-if1` interface:

```
configure
 port ethernet 12/2
  encapsulation dot1q
  dot1q pvc 30 encapsulation 1qtunnel
   bind interface vrrp-if1 local
! SVLAN 30 is now a monitor PVC
   exit
  dot1q pvc 30:1
   track vrrp 1 vrrp-if1 local
   service clips dhcp context local
   exit
  dot1q pvc 30:2
   track vrrp 1 vrrp-if1 local
   service clips dhcp context local
! PVCs 30:1 and 30:2 are track PVCs of interface
vrrp-if1 on port 12/2
```

Glossary

follower circuit
A subscriber circuit that tracks a VRRP interface to which a monitor circuit is bound; that is, when a port fails in which one or more VRRP monitor circuits and follower circuits are configured, the follower circuits failover to an interface on a backup VRRP router with the same VRIP and same VRRP router ID.

giaddr
A field in a DHCP message that holds the Relay Agent IP address used in booting via a relay agent.

monitor PVC
A static 802.1Q PVC that monitors the status of a port for ICR failover operation. The monitor PVC is bound to the VRRP interface to a VRRP router.

VRRP
Virtual Router Redundancy Protocol

VRRP interface
An interface to a VRRP router on which a VRRP interface IP address (VRIP) is configured.