



Engineering Rules

This section provides engineering rules or guidelines that must be considered prior to configuring the system for your network deployment.

- [Interface and Port Rules, on page 1](#)
- [Subscriber Rules, on page 3](#)
- [Service Rules, on page 3](#)

Interface and Port Rules

The rules discussed in this section pertain to the Ethernet 10/100 line card, the Ethernet 1000 line card and the four-port Quad Gig-E line card and the type of interfaces they facilitate, regardless of the application.

R-P Interface Rules

The following engineering rules apply to the R-P interface:

- An R-P interface is created once the IP address of a logical interface is bound to a PDSN service.
- The logical interface(s) that will be used to facilitate the R-P interface(s) must be configured within an "ingress" context.
- PDSN services must be configured within an "ingress" context.
- At least one PDSN service must be bound to each interface; however, multiple PDSN services can be bound to a single interface if secondary addresses are assigned to the interface.
- Each PDSN service must be configured with the Security Parameter Index (SPI) of the Packet Control Function (PCF) that it will be communicating with over the R-P interface.
- Multiple SPIs can be configured within the PDSN service to allow communications with multiple PCFs over the R-P interface. It is best to define SPIs using a netmask to specify a range of addresses rather than entering separate SPIs. This assumes that the network is physically designed to allow this communication.
- Depending on the services offered to the subscriber, the number of sessions facilitated by the R-P interface can be limited.

Closed R-P Interface Rules

The following engineering rules apply to the Closed R-P interface:

- A Closed R-P interface is created once the IP address of a logical interface is bound to a Closed R-P service.
- The logical interface(s) that will be used to facilitate the Closed R-P interface(s) must be configured within an "ingress" context.
- Closed R-P services must be configured within an "ingress" context.
- Up to 64 Closed R-P services may be configured system wide. All 64 Closed R-P services may be configured in one context.
- At least one Closed R-P service must be bound to each interface; however, multiple Closed R-P services can be bound to a single interface.
- The Closed R-P service must not be bound to the same ip address of an interface as other L2TP LAC or LNS services.
- Each Closed R-P service must be configured with the Peer PCF address of the Packet Control Function (PCF) that it will be communicating with over the Closed R-P interface.
- Multiple Peer PCFs can be configured within the Closed R-P service to allow communications with multiple PCFs over the Closed R-P interface. This assumes that the network is physically designed to allow this communication.
- Depending on the services offered to the subscriber, the number of sessions facilitated by the Closed R-P interface can be limited.

Pi Interface Rules

FA to HA Rules

When supporting Mobile IP, the system can be configured to perform the role of a FA, an HA, or both. This section describes the engineering rules for the Pi interface when using the system as a FA.

The following engineering rules apply to the Pi interface between the FA and HA:

- A Pi interface is created once the IP address of a logical interface is bound to an FA service.
- The logical interface(s) that will be used to facilitate the Pi interface(s) must be configured within the egress context.
- FA services must be configured within the egress context.
- If the system is configured as a FA is communicating with a system configured as an HA, then it is recommended that the name of the context in which the FA service is configured is identical to the name of the context that the HA service is configured in on the other system.
- Each FA service may be configured with the Security Parameter Index (SPI) of the HA that it will be communicating with over the Pi interface.

- Multiple SPIs can be configured within the FA service to allow communications with multiple HAs over the Pi interface. It is best to define SPIs using a netmask to specify a range of addresses rather than entering separate SPIs. This assumes that the network is physically designed to allow this communication.
- Depending on the services offered to the subscriber, the number of sessions facilitated by the Pi interface can be limited.

HA to FA

The following engineering rules apply to the Pi interface between the HA and FA:

- When supporting Mobile IP, the system can be configured to perform the role of a FA, an HA or both. This section describes the engineering rules for the Pi interface when using the system as an HA.
- A Pi interface is created once the IP address of a logical interface is bound to an HA service.
- The logical interface(s) that will be used to facilitate the Pi interface(s) must be configured within an ingress context.
- HA services must be configured within an ingress context.
- If the system configured as an HA is communicating with a system configured as a FA, then it is recommended that the name of the context in which the HA service is configured is identical to the name of the context that the FA service is configured in on the other system.
- Each HA service may be configured with the Security Parameter Index (SPI) of the FA that it will be communicating with over the Pi interface.
- Multiple SPIs can be configured within the HA service to allow communications with multiple FAs over the Pi interface. It is best to define SPIs using a netmask to specify a range of addresses rather than entering separate SPIs. This assumes that the network is physically designed to allow this communication.
- Each HA service must be configured with a Security Parameter Index (SPI) that it will share with mobile nodes.
- Depending on the services offered to the subscriber, the number of sessions facilitated by the Pi interface can be limited in order to allow higher bandwidth per subscriber.

Subscriber Rules

The following engineering rule applies to subscribers configured within the system:

- A maximum of 2,048 local subscribers can be configured per context.
- Default subscriber templates may be configured on a per PDSN or FA service.

Service Rules

The following engineering rules apply to services configured within the system:

- A maximum of 256 services (regardless of type) can be configured per system.



Caution Large numbers of services greatly increase the complexity of management and may impact overall system performance (i.e. resulting from such things as system handoffs). Therefore, it is recommended that a large number of services only be configured if your application absolutely requires it. Please contact your local service representative for more information.

- Up to 2,048 Security Parameter Indices (SPIs) can be configured for a single PDSN service.
- Up to 2,048 MN-HA and 2048 FA-HA SPIs can be supported for a single HA service.
- Up to 2,048 FA-HA SPIs can be supported for a single FA service.
- The system supports unlimited peer FA addresses per HA.
 - The system maintains statistics for a maximum of 8192 peer FAs per HA service.
 - If more than 8192 FAs are attached, older statistics are identified and overwritten.
- The system maintains statistics for a maximum of 4096 peer HAs per FA service.
- There are a maximum of 8 HA assignment tables per context and per chassis.
- The total number of entries per table and per chassis is limited to 256.
- Up to 10,000 LAC addresses can be configured per LNS service.
- Even though service names can be identical to those configured in different contexts on the same system, this is not a good practice. Having services with the same name can lead to confusion, difficulty troubleshooting problems, and make it difficult understanding outputs of show commands.
- Up to 10,000 PCF addresses can be configured per PDSN ClosedRP service.
- Up to 64 Closed R-P services may be configured system wide. All 64 Closed R-P services may be configured in one context.