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

The rules discussed in this section pertain to both the Ethernet 10/100, the Ethernet 1000 Line Card and the four-port Quad Gigabit Ethernet Line Card, known as the Quad Gig-E or QGLC and the type of interfaces they facilitate, regardless of the application.

Pi Interface Rules

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:

Default subscriber templates may be configured on a per HA service.

Service Rules

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

☞

Important Given capacities do not apply to the XT2 platform.

⚠

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.

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

- 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.

- Single HA service shall support more than one enterprise.

- Total number of service addresses per VPN context limited to 512.