

Engineering Rules

This appendix provides engineering guidelines for configuring the system to meet network deployment requirements.

Note The Engineering Rules listed in this appendix reflect the maximum capacity of StarOS. The actual limits of VPC running in a VM are governed by the amount of vCPU and vMemory capacity allocated to the instance.

- CLI Session Rules, on page 1
- VPC-DI Interface and Port Rules, on page 1
- Context Rules, on page 3
- Subscriber Rules, on page 6
- Service Rules, on page 6
- Access Control List (ACL) Engineering Rules, on page 7
- ECMP Groups, on page 7
- VPC-DI Network Dependencies, on page 8
- VPN Scaling Requirements, on page 9

CLI Session Rules

Multiple CLI session support is based on the amount of available memory. The internal Resource Manager reserves enough resources to support a minimum of six CLI sessions at all times. One of the six sessions is further reserved for use exclusively by a CLI session on the serial interface.

Additional CLI sessions beyond the pre-reserved limit are permitted if sufficient resources are available. If the Resource Manager is unable to reserve resources for a CLI session beyond those that are pre-reserved, users with administrator-privileges are prompted to create the new CLI session, even without reserved resources.

VPC-DI Interface and Port Rules

The rules discussed in this section pertain to the vNIC Ethernet ports used for context traffic on Service Function (SF) VMs.

vNIC Ethernet Ports

- Give all logical interfaces a unique name to identify the interface from others in the same context. Logical interfaces in different contexts may have the same name.
- A single virtual port can support multiple logical interfaces when you configure VLAN tags for that port. You can use VLAN tagging to bind a single port to multiple logical interfaces that reside in different contexts.
- Assign all logical interfaces a valid IP address and subnet.
  - Give each logical interface within a context a unique IP address. Logical interfaces in different contexts can have the same IP address.
  - If multi-homing is supported on the network, you can assign all logical interfaces a single primary IP address and up to 16 secondary IP addresses.
  - Shared subnets are configurable for both primary and secondary IP addresses. For example, this configuration is valid:

```
interface intf1
   ip address 190.20.20.2 255.255.255.0
   ipv6 address 2001::34/64 secondary
   exit
interface intf2
   ip address 190.20.20.3 255.255.255.0
   ipv6 address 2001::35/64 secondary
   exit
interface intf3
   ipv6 address 2001::36/64
   ip address 190.20.20.4 255.255.255.0 secondary
   exit
```

- You can configure a logical interface in only one context, but you can configure multiple interfaces (up to 512) in a single context.
- You can apply a maximum of 256 access control list (ACL) rules to a single logical interface.
- All ports are identified by their *slot*/*port*. For VPC-DI, the *slot* corresponds to a CF or SF VM within the virtual chassis (VPC-DI instance). The hypervisor assigns a unique slot number to each VM during initial configuration of the VPC-DI instance. Slots 1 and 2 are assigned to the CF VMs slot numbers 3 to 32 are assigned to SF VMs. The CF only supports port 1.

  Each SF supports 12 vNICs numbered 1 to 12 with corresponding virtual Ethernet ports numbered 10 to 21. SF port number 10 must be configured.
- Each vNIC port for subscriber traffic may contain up to a maximum of 1,024 VLAN tags.
- A logical interface is limited to using a single VLAN on a single physical port, identified by its <*card*/*slot*/*port*>.

Packet Data Network (PDN) Interface Rules

The following engineering rules apply to the interface to the packet data network (PDN):

- Configure the logical interfaces used to facilitate the PDN interface within the egress context.

- The default is to use a single interface within the egress context to facilitate the PDN interface.
- You can configure multiple interfaces in the egress context by using static routes or dynamic routing protocols.
- You may also configure next-hop default gateways.

Context Rules

- A maximum of 63 contexts may be configured per chassis.Enabling demux functions on an MIO card reduces the maximum number of contexts to 10.
- Interfaces per Context
  - With the Demux MIO/UMIO/MIO2 feature enabled, up to 64 interfaces can be configured within a single context.
  - 512 Ethernet+PPP+tunnel interfaces
  - 32 ipv6ip tunnel interfaces
  - 511 GRE tunnels (2,048 GRE tunnels per chassis)
  - 256 loopback interfaces
- IP Addresses and IP Address Pools
  - Up to 2,000 IPv4 address pools can be configured within a single context.
  - Up to 256 IPv6 pools can be configured within a single context.
  - Up to a combined total of 5,000 IPv4 and IPv6 addresses can be configured per chassis.
  - Each context supports up to 32,000,000 static IP pool addresses. You can configure a maximum total of 96,000,000 static IP pool addresses per chassis. Each static IP pool can contain up to 500,000 addresses.
  - Each context supports up to 16,000,000 dynamic IP pool addresses. You can configure a maximum total of 32,000,000 dynamic IP pool addresses per chassis. Each dynamic IP pool can contain up to 500,000 addresses.



Important The actual number of IP Pools supported per context and chassis depends on how many addresses are being used and how they are subnetted.



Important Each address in the pool requires approximately 60 bytes of memory. The amount of memory required, however, depends on a number of factors such as the pool type, and hold-timer usage. Therefore, in order to conserve available memory, you may need to limit the number of pools depending on the number of addresses to be configured and the number of installed application cards.

- The maximum number of simultaneous subscriber sessions is controlled by the installed capacity license for the service(s) supported.
- The maximum number of static address resolution protocol (ARP) entries per context is 128.
- The maximum number of domains per context is 2,048.
- ASN-GW services configured within the same context cannot communicate with each other.
- Routes
  - Up to 1,200 static routes per context (48,000 per chassis).
  - 6,000 pool routes per context (6,000 per chassis)
  - 24,000 pool explicit host routes per context (24,000 per chassis)

    40,000 pool explicit host routes per context in 21.28.x and above. For more information on extending the explicit route limit to 40,000, see Command Line Interface Reference, Modes C - D, StarOS Guide.
  - 64 route maps per context
- BGP
  - 64,000 BGP prefixes can be learned/advertised per context (64,000 per chassis)
  - 64 EBGP peers can be configured per context (512 per chassis)
  - 16 IBGP peers per context
  - 512 BGP/AAA monitors per context in support of Interchassis Session Recovery (ICSR)
- OSPF
  - 200 OSPF neighbors per chassis
  - 10,000 OSPF routes per context (64,000 per chassis)
- MPLS

  Until Release 21.6
  - 16 label distribution protocol (LDP) sessions per context
  - Up to 8,000 incoming label map (ILM) entries per context (48, 000 per chassis)
  - Combine Table size of 128,000 next hop label forwarding entries (NHLFE) and 64,000 prefixes that could potentially yield:
    - 1,000 forwarding equivalence class (FEC) entries per context (4,000 per chassis) - with 32 paths
    - 2,000 forwarding equivalence class (FEC) entries per context (8,000 per chassis) - with 16 paths
    - 16,000 forwarding equivalence class (FEC) entries per context (64,000 per chassis) - with 2 paths
    - 64,000 forwarding equivalence class (FEC) entries per context (64k per chassis) - with 1 path

- *Release 21.7 and higher*
  - 16 label distribution protocol (LDP) sessions per context
  - Up to 8,000 incoming label map (ILM) entries per context (48,000 per chassis)
  - Combine Table size of 256,000 next hop label forwarding entries (NHLFE) and 64,000 prefixes that could potentially yield:
    - 1,000 forwarding equivalence class (FEC) entries per context (4,000 per chassis) - with 64 paths
    - 2,000 forwarding equivalence class (FEC) entries per context (8,000 per chassis) - with 32 paths
    - 32,000 forwarding equivalence class (FEC) entries per context (64,000 per chassis) - with 2 paths
    - 64,000 forwarding equivalence class (FEC) entries per context (64,000 per chassis) - with 1 path

- VRF
  - 300 virtual routing and forwarding (VRF) tables per context (2,048 VRFs per chassis) [256 VRFs per context with demux functions enabled on the MIO card]
  - APN limit is 2,048 per chassis; VRF limits and APN limits should be identical.
  - 64,000 IP routes

- NEMO (Network Mobility)
  - 512K prefixes/framed routes per chassis and up to 16 dynamically learned prefixes per MR (Mobile Router)

- 128 AAA servers per context for a default AAA server group. The servers can be configured as accounting, authentication, charging servers, or any combination thereof.
- You can configure up to 800 AAA server groups per context with following limitations:
  - 128 servers per AAA server group (accounting, authentication, charging server, or any combination thereof)
  - 1,600 servers per context in AAA Server group mode (accounting, authentication, charging server, or any combination thereof)
  - 800 NAS-IP address/NAS identifier (one primary and one secondary per server group) per context
- Up to 12 charging gateway functions (CGFs) for GTPP accounting can be configured per context.
- Up to 16 bidirectional forwarding detection (BFD) sessions per context (64 per chassis)



Important Refer to *Engineering Rules* in your product administration guide for additional information on product-specific operating limits.

Subscriber Rules

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

- Configure a maximum of 2,048 local subscribers per context.
- You may configure attributes for each local subscriber.
- The system creates a default subscriber for each context when the context is made. Configure attributes for each default subscriber. If a AAA-based subscriber is missing attributes in the authentication reply message, the default subscriber attributes in the context where the subscriber was authenticated are used.



Important Default is not used when local authentication (for local subscribers) is performed.

- Configure default subscriber templates on a per AAA realm (domain aliases configured within a context) basis.
- Configure default subscriber templates on a per PDSN, FA, ASN-GW, or HA service.
- For AAA authenticated subscribers, the selection of local subscriber template to use for setting attributes is in the following order:
  - If the username (NAI) matches any local domain name and the domain name has a local subscriber name configured, that local subscriber template is used.
  - If the first case fails, and if the serving service has a default username configured, that subscriber template is used.
  - If the first two cases fail, the default subscriber template in the AAA context is used.

Service Rules

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

- Configure a maximum of 256 services (regardless of type) per system.



Caution Large numbers of services greatly increase the complexity of management and may affect overall system performance. Therefore, you should not configure a large number of services unless your application absolutely requires it. Please contact your Cisco service representative for more information.

- The total number of entries per table and per chassis is limited to 256.
- Although you can use service names that are identical to those configured in different contexts on the same system, this is not a good practice. Services with the same name can lead to confusion and difficulty in troubleshooting problems, and make it difficult to understand the output of **show** commands.

Access Control List (ACL) Engineering Rules

The following rules apply to Access Control Lists:

- The maximum number of rules per ACL is 256.
- The maximum number of ACL rules applied per port is 128.
- The maximum number of ACL rules applied per context is 2048.
- The maximum number of ACL rules per IPSec policy is 1.
- The maximum number of IPSec ACL rules per context is 1,024.
- The maximum number of IPSec ACL rules per crypto map is 8.
- The maximum number of ACLs you can configure per context is limited by the number of rules allowed within each ACL. If each ACL contained the maximum number of rules (256), the maximum number of ACLs per context is 8 (256 X 8 ACLs = 2048 ACL rules per context).
- The maximum number of ACLs applied to an IP access group is 1, whether it is configured for a port or context. Since the maximum number of IP access groups you can apply to an interface or context is 16, the following calculations apply:
  - For each interface/port: 8 rules per ACL multiplied by 16 IP access groups = 128 (the ACL rules limit per port)
  - For each context: 16 rules per ACL multiplied by 16 IP access groups = 256 (the ACL rules limit per context)
- Setting up a maximum of 1200 ACLs on an ASR5500 system is advised. Depending on the amount of memory available in the sct proclet, more ACLs can be configured in the system.

ECMP Groups

The maximum number of Equal Cost Multiple Path (ECMP) groups are as follows:

- For releases prior to 21.4, StarOS supports a maximum of 64000 groups.
- For release 21.4 and higher, StarOS supports a maximum of 32000 groups.

Note

- *max_num* is an integer from 1 through 10

  Release 21.4x

  - QVPC-DI: 64
  - QVPC-SI: 64
  - ASR 5500: 24

- Save your configuration as described in the Verifying and Saving Your Configuration chapter.

VPC-DI Network Dependencies

This section outlines the services and network devices required to use functionality of StarOS gateways in a VPC-DI instance.

Routers

Routers supporting the VPC-DI network must support:

- IPv4 and IPv6 interfaces (NOTE: Within each routing context, StarOS requires each IP interface to be in a unique subnet.)
- Static routes
- *Optional:* OSPFv2 and/or OSPFv3 (Requires multicast support for IPv4/IPv6)
- eBGP (External Border Gateway Protocol)
  - Multihop
  - Multipath
  - IPv4 address Family
  - IPv6 address Family
  - *Optional:* VPNv4 and VPNv6 for Enterprise style configurations
  - Via eGBP these routers would learn IP pools, routes, service interface addresses, etc.
- Routing Policies to filter the routes.
- MPLS encapsulation and forwarding (NOTE: required only for enterprise style configurations.)
- ICMP and ICMPv6 support
- BFD (Bidirectional Forwarding Detection) [for quicker detection of failures]
- ECMP (Equal Cost Multiple Path)

External Network Dependencies

- RADIUS or Diameter Services
- DNS Services
- NTPv4 Services
- SNMP services
- syslog services
- SSH/SFTP services (for bulkstats and administration)
- *Optional:* TACACS+ services
- Billing Services (CDR/EDR)
- VLAN encapsulation
- ICSR requirements:
  - Dedicated connectivity between primary and backup virtual chassis
  - eBGP
  - Layer 2 L2 Switch functionality between two virtual chassis (for processing GARP/Mac Address Changes, when chassis state is changed)

VPN Scaling Requirements

The following VPN scaling numbers are supported for specific releases.

Parameter	Scaling Number (Release 12.x, 14.x)	Scaling Number (Release 15.x, 16.x)	Scaling Number (Release 17.x, 18.x, 19.x, 20.x +)
BFD Sessions	16 per context 64 per chassis	16 per context 64 per chassis	16 per context 64 per chassis
Context level ACLs	256 per context	256 per context	256 per context
Dynamic pool addresses	16 million per context 32 million per chassis	16 million per context 32 million per chassis	16 million per context 32 million per chassis
IPv4 pools per context	2000 per context 5000 per chassis (combined IPv4 and IPv6)	2000 per context 5000 per chassis (combined IPv4 and IPv6)	2000 per context 5000 per chassis (combined IPv4 and IPv6)
IPv6 pools per context	32 per context 5000 per chassis (combined IPv4 and IPv6)	256 per context 5000 per chassis (combined IPv4 and IPv6)	256 per context 5000 per chassis (combined IPv4 and IPv6)
Number of BGP prefixes	16,000 per context 64,000 per chassis.	32,000 per context 64,000 per chassis.	64,000 per context 64,000 per chassis.

Parameter	Scaling Number (Release 12.x, 14.x)	Scaling Number (Release 15.x, 16.x)	Scaling Number (Release 17.x, 18.x, 19.x, 20.x +)
Number of contexts	63 (though PSC migrations do not work well beyond 32 contexts)	63 (though PSC migrations do not work well beyond 32 contexts) Note the information about "Demux on MIO Cards" at the end of this section.	63 (though PSC migrations do not work well beyond 32 contexts) Note the information about "Demux on MIO Cards" at the end of this section.
Number of dynamically learned prefixes per MR	8	16	16
Number of EBGP peers	64 per context 512 per chassis	64 per context 512 per chassis	64 per context 512 per chassis
Number of FEC entries	8000 labels per context 48,000 per chassis	8000 labels per context 48,000 per chassis	8000 labels per context 48,000 per chassis
Number of IBGP peer	16 per context	16 per context	16 per context
Number of ILM entries	8000 labels per context 48,000 per chassis	8000 labels per context 48,000 per chassis	8000 labels per context 48,000 per chassis
Number of interfaces	512 ethernet+ppp+tunnel interfaces per context 32 IPv6 IP tunnel interfaces per context Upto 511 GRE tunnels/context and 2048 GRE tunnels/chassis 256 loopback interfaces per context	512 ethernet+ppp+tunnel interfaces per context 32 IPv6 IP tunnel interfaces per context Upto 511 GRE tunnels/context and 2048 GRE tunnels/chassis 256 loopback interfaces per context Note the information about "Demux on MIO Cards" at the end of this section.	512 ethernet+ppp+tunnel interfaces per context 32 IPv6 IP tunnel interfaces per context Upto 511 GRE tunnels/context and 2048 GRE tunnels/chassis 256 loopback interfaces per context Note the information about "Demux on MIO Cards" at the end of this section.
Number of LDP sessions	16 per context	16 per context	16 per context
Number of NEMO prefixes/Framed routes	256,000 per chassis	512,000 per chassis	512,000 per chassis
Number of OSPF neighbors	Up to 200 per chassis	Up to 200 per chassis	Up to 200 per chassis
Number of OSPF routes	Up to 10,000 per context 64,000 per chassis	Up to 10,000 per context 64,000 per chassis	Up to 10,000 per context 64,000 per chassis

Parameter	Scaling Number (Release 12.x, 14.x)	Scaling Number (Release 15.x, 16.x)	Scaling Number (Release 17.x, 18.x, 19.x, 20.x +)
Number of pool explicit host routes	5000 per context (6000 per chassis)	5000 per context (6000 per chassis)	5000 per context (6000 per chassis) in 17.x and 18.[1234] 24000 per context (24000 per chassis) in 18.5 and above 40,000 per context in 21.28.x and above. For more information on how to extend the explicit route limit configuration to 40,000, see Command Line Interface Reference, Modes C - D, StarOS Guide.
Number of pool routes	6000 per context (6000 per chassis)	6000 per context (6000 per chassis)	6000 per context (6000 per chassis)
Number of routes (excluding framed-routes)	64,000 per context	64,000 per context	64,000 per context
Number of secondary addresses per interface	16	16	16
Number of Static routes	1200 per context	1200 per context	1200 per context
Number of VLANs	4000 per chassis	4000 per chassis	4000 per chassis
Number of VRFs	250 per context 2048 per chassis APN limit is 1024/chassis and does not match the VRF limit.	300 per context 2048 per chassis **Note** • VRF Limits and APN limits are assumed to be identical. • Note the "Demux on MIO Cards" section at the end of this section.	300 per context 2048 per chassis Note: VRF Limits and APN limits are assumed to be identical. Note the "Demux on MIO Cards" section at the end of this section.
Number of routes (all kinds of routes including framed routes)	64,000 per context	64,000 per context	64,000 per context
Route maps	64 per context	64 per context	64 per context

Parameter	Scaling Number (Release 12.x, 14.x)	Scaling Number (Release 15.x, 16.x)	Scaling Number (Release 17.x, 18.x, 19.x, 20.x +)
Static pool addresses	32 million per context 96 million per chassis	32 million per context 96 million per chassis	32 million per context 96 million per chassis

Demux on MIO Cards

When enabling Demux on MIO cards, VPN resources are combined on the MIO cards with the controller processes, thus reducing the resources available for all VPN tasks. This results in reducing some of the limits (mentioned in the previous section) when MIO cards are demux-enabled.