



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, page 1](#)
- [VPC Interface and Port Rules, page 1](#)
- [Context Rules, page 2](#)
- [Subscriber Rules, page 5](#)
- [Service Rules, page 5](#)
- [Access Control List \(ACL\) Engineering Rules, page 6](#)
- [ECMP Groups, page 6](#)

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 Interface and Port Rules

The rules discussed in this section pertain to the vNIC Ethernet ports designated via the hypervisor for subscriber traffic.

vNIC Ethernet Ports

- Give all hypervisorassigned logical interfaces a unique name to differentiate the interface from others in the same context.
- A single virtual port can support multiple hypervisorassigned 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.
- Each vNIC port for subscriber traffic may contain up to a maximum of 1,024 VLAN tags (maximum of 4,000 VLANs per VPC chassis).
- All hypervisorassigned logical interfaces must have a valid IP address and subnet.
 - If multihoming is supported on the network, you can assign all logical interfaces a single primary IP address and up to 16 secondary IP addresses.
- You configure a single StarOS logical (named interface per context. That named interface can have up to 512 ethernet+ppp+tunnel interfaces.
- Different StarOS contexts can share the same logical (named interface.
- You can apply a maximum of 256 access control list (ACL rules to a StarOS logical interface.
- In StarOS all ports are identified by their <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
 - *Prior to Release 15.0:* Up to 16 interfaces can be configured within a single context.
 - *For Release 15.0 and higher:* 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.
 - *Prior to Release 15.0:* Up to 32 IPv6 pools can be configured within a single context.
 - *For Release 15.0 and higher:* 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)
 - *Releases prior to 18.5:* 5,000 pool explicit host routes per context (6,000 per chassis)
 - *Release 18.5 and higher:* 24,000 pool explicit host routes per context (24,000 per chassis)
 - 64 route maps per context
- BGP

- *Releases 12 and 14:* 16,000 BGP prefixes can be learned/advertised per context (64,000 per chassis)
- *Releases 15 and 16:* 32,000 BGP prefixes can be learned/advertised per context (64,000 per chassis)
- *Releases 17, 18 and higher:* 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
 - 16 label distribution protocol (LDP) sessions per context
 - 8,000 forwarding equivalence class (FEC) entries per context (48, 000 per chassis)
 - Up to 8,000 incoming label map (ILM) entries per context (48, 000 per chassis)
- VRF
 - *Prior to Release 15.0:* 250 virtual routing and forwarding (VRF) tables per context (1,024 or 2,048 [release 14.0+] VRFs per chassis)
 - *Release 15.0 and higher:* 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)
 - *Prior to Release 15.0:* 256K prefixes/framed routes per chassis and up to 8 dynamically learned prefixes per MR (Mobile Router)
 - *Release 15.0 and higher:* 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 128.
- The maximum number of ACL rules applied per port is 128.
- The maximum number of ACL rules applied per context is 1,024.
- 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 (128), the maximum number of ACLs per context is 8 (128 X 8 ACLs = 1,024 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: 64 rules per ACL multiplied by 16 IP access groups = 1,024 (the ACL rules limit per context)

ECMP Groups

The maximum number of ECMP groups are as follows:

- *For releases prior to 17.0*, StarOS supports a maximum of 512 groups.
- *For release 17.0 and higher*, StarOS supports a maximum of 2048 groups.