ERICSSON

SmartEdge OS Hardening Guide

SYSTEM ADM. GUIDE

89/1543-CRA 119 1170/1 Uen E1

Contents

1 General Information

1.1 Scope

This document describes how to harden the SmartEdge OS in general and various services running on it by following certain precautions and configuring it to protect against security threats. SmartEdge OS security is verified by targeted tests with industry-leading security test tools.

1.2 Audience

This document is intended for operators and people responsible for configuration and maintenance of the SmartEdge OS and services running on it.

1.3 Terminology

ACL	Access Control List
BGF	Border Gateway Function
SBC	Session Border Controller
OAM	Operation Administration and Maintenance
DoS	Denial of Service
DDoS	Distributed Denial of Service
RPF	Reverse Path Forwarding

2 Overview

The SmartEdge® platform is a multiservice edge router providing forwarding and advanced Layer 2 (L2) to Layer 7 (L7) services to carriers around the world. With the rapid expansion of the Internet-connected devices, bandwidth, and multimedia (data, voice, and video)-security has become an important aspect of handling internet traffic. A SmartEdge router deployed at the edge of the network is directly exposed to various types of security attacks. SmartEdge OS comprehensive security features protect the SmartEdge router and other nodes in the core network from various attacks.

3 Security from the Ground Up

Ericsson applies general security design rules and requirements that must be met by all Ericsson nodes. The SmartEdge OS security implementation is a multilayered strategy that provides protection at various components and modules in the system. The strategy includes the following main components:

- Secure OS
- Secure operation, administration, and maintenance
- Secure forwarding plane
- Secure Service layer
- Packet filtering
- Security protocols
- Security alarms
- Security logging
- Media plane attack prevention

In addition, the SmartEdge OS is continuously tested and validated against industry-standard security testing tools to make sure that the system is fully compliant with security requirements and industry standards.

3.1 Secure OS

The SmartEdge OS provides a set of security functions in multiple areas such as packet filters, logging, access control, routing protocols, and the kernel.

3.1.1 Kernel and OS Hardening

Administrative ACLs are used to filter traffic sent to the kernel. BOOTP, Finger, Identd, PAD, HTTP-Server, and HTTPS-Server services are not supported.

3.2 Secure OAM

3.2.1 Access Control

You can access the SmartEdge OS by a directly connected console port or by telnet or ssh sessions to the management port. Access is permitted only to

successfully authentication users, based on username and password. The user database used for authentication can be locally managed on the SmartEdge platform or centrally managed on a RADIUS or TACACS server.

3.2.2 Local User Management

You can manage local user accounts through the command-line interface (CLI). All user accounts must have a password, which is stored encrypted in the configuration file. After you log in the first time, you can change your password using CLI commands.

3.2.3 Privilege Levels

In SmartEdge OS, user privilege levels determine the set of commands accessible to that user. Users with the default privilege level (6) cannot configure the system but can modify some ACL rule conditions. Access to higher privilege levels is password protected.

3.2.4 Login Sessions

Logins are permitted only by successful authentication. Idle sessions are disconnected by default after 10 minutes. The idle timeout interval is configurable.

To discourage brute-force login attempts, the session is terminated after three authentication failures. Successful and failed logins as well as logouts are logged to the system log to provide security trails.

3.2.5 Secure Access

The OAM traffic can be both physically and logically separated from other traffic by using separate network interfaces.

The SmartEdge platform supports secure protocols such as SSH, SCP, and SFTP. Traffic over line cards can be secured with IPsec by using the ASE card.

SSHv2 provides better security than SSHv1.

You can disable SSHv1 by using the following command at the global level:

```
(config)#no ssh server v1
```

You can disable unsecured access to the SmartEdge OS by using the following commands at the local context:

```
(config-ctx)#no service telnet
(config-ctx)#no service ftp client
(config-ctx)#no service tftp
```

```
(config-ctx)#no service rcp client
```

You can enable the secure alternatives by using the following commands at the local context:

```
(config-ctx)# service sftp
(config-ctx)# service scp
```

3.2.6 Logging

Event logs and alarms raised by different modules on the SmartEdge platform are handled by the logging infrastructure. Logs can be filtered based on the log level and directed to multiple destinations such as the system console, local storage, and remote syslog servers. You can use tunnel mode IPsec to transmit logs to syslog servers securely.

Malicious traffic logs are visible to the logging infrastructure, but ACL drop logs are not.

The security audit trail logs successful logins, logouts, and failed login attempts.

3.2.7 Secure SNMP

The SmartEdge OS supports SNMPv1, SNMPv2, and SNMPv3. You can configure community strings only by using the CLI. DES encryption is supported for SNMPv3. For security reasons, use SNMPv3 whenever possible.

3.3 L2 Security

The SmartEdge platform includes various features that provide L2 security and protect against various L2 attacks.

3.3.1 Ports & VLANs

All ports on the SmartEdge router are disabled by default, and have to be explicitly enabled and bound to an interface or associated with a bridge instance to be functional. By default, routing protocols are not enabled on any interfaces. You must configure a VLAN explicitly by setting the port encapsulation to 802.1q and creating PVCs. Merely setting the encapsulation to 802.1q does not result in a default VLAN being created.

3.3.2 Bridging

3.3.2.1 MAC Flooding

The SmartEdge platform has several mechanisms that protect against MAC flooding attacks:

- Disabling of MAC learning
- Configuration of static MAC entries
- Support for disallowed MAC lists
- Rate limits for broadcast traffic, multicast traffic, and traffic to unknown destinations
- Configuration of circuits as restricted (no MAC learning combined with static MAC entries)

3.3.2.2 ARP

By default, the SmartEdge router accepts gratuitous ARP messages and is vulnerable to ARP poisoning. You can mitigate this problem by enabling secure ARP. You can rate-limit ingress ARP messages by configuring a QoS rate-limit policy and applying the policy to relevant circuits.

3.4 L3 Security

The SmartEdge platform supports various features that provide protection from L3 attacks. Besides detecting malicious traffic, the SmartEdge router counts and logs malicious traffic and raises alarms when malicious traffic exceeds a configured threshold.

3.4.1 Routing

We recommend that you use loopback interfaces for all routing protocols. However, no restriction exists for configuring routing protocols to use nonloopback interfaces. By default, routing protocols are not enabled on any interface.

Manually distributed keys, stored encrypted in the configuration file, are used for authentication. Authentication is implemented for all unicast IPv4 protocols. BGP peers are authenticated using MD5.

You can configure prefix lists. The maximum number of prefixes accepted is configurable by BGP peer or by peer group and address family.

The BGP process receives packets only on sockets bound to a specific local address and connected to a particular destination address. You can configure administrative ACLs in the BGP context to provide additional security.

3.4.1.1 Keychain Definition

SmartEdge OS IGPs support MD5 authentication using keychains, as in the following example:

```
context igp-auth
!
 no ip domain-lookup
!
 interface one
  ip address 20.1.1.1/24
!
 interface two
  ip address 20.1.2.1/24
 no logging console
!
 router ospf 1
  fast-convergence
  area 0.0.0.0
   interface one
    authentication md5 igp-auth
   interface two
    authentication md5 igp-auth-2
!
 router rip 1
  interface one
   authentication md5 igp-auth
  interface two
   authentication md5 igp-auth-2
!
 router isis 1
  address-family ipv4 unicast
 !
  interface one
   ! not bound to any circuit
   authentication key-chain igp-auth
   address-family ipv4 unicast
   address-family ipv6 unicast
!
 key-chain igp-auth key-id 1
  key-string encrypted 1332D19ABE0212296D0FC2AA0BE83874
!
 key-chain igp-auth-2 key-id 1
  key-string encrypted 307F6191B6E81DF06D0FC2AA0BE83874
!
```

Example 1

3.4.1.2 BGP Authentication

MD5 authentication can be configured for BGP at either the neighbor or peer-group level:

```
context bgp-auth
!
 no ip domain-lookup
 no logging console
!
 router bgp 1
!
  peer-group group-1 external
    password encrypted 5B6BB7367E0B4818FBC0458C112BC2F1
   address-family ipv4 unicast
!
  neighbor 30.1.1.2 external
    password encrypted 2C4431DD321FCA7E9AE114DBF5E29241
   address-family ipv4 unicast
```

Example 2

3.4.1.3 BGP Prefix Lists

You can use prefix lists to limit the routes BGP learns or advertises at the neighbor or peer-group level. The following example limits routes learned for both IPv4 and IPv6:

```
context bgp-prefix-list
!
 no ip domain-lookup
 no logging console
!
 ip prefix-list bgp-v4-prefix-list
  seq 10 permit 50.0.0.0/8 ge 8
!

router bgp 1
!
  peer-group our-group internal
   address-family ipv4 unicast
      prefix-list bgp-v4-prefix-list in
!
  neighbor 50.1.1.2 internal
   address-family ipv4 unicast
      prefix-list bgp-v4-prefix-list in
!
```

Example 3

3.4.1.4 Keychain Graceful Key Rollover

Using keychains, you can gracefully switch from one key to the next. In the following example, the key with *key-id 1* is accepted for an entire day after the key with *key-id 2* is sent. This allows a new key to be deployed without synchronizing the configurations.

```
!
 key-chain graceful-rollover key-id 1
  key-string encrypted 181A990F5A9E4DD09308445556F0F589
  accept-lifetime 2008:12:31:00:00 2010:01:01:23:59
  send-lifetime 2009:01:01:00:00 2009:12:31:23:59
!
 key-chain gracefull-rollover key-id 2
  key-string encrypted 48BDAA0754E0C9A121832D9A5BBCB2839308445556F0
  accept-lifetime 2009:12:31:00:00 2011:01:01:23:59
  send-lifetime 2010:01:01:00:00 2010:12:31:23:59
```

Example 4

3.4.2 Packet Filtering Using ACLs

You can filter packets by configuring IP ACLs. When you apply an ACL to an interface, packets received and sent over the interface are subject to the rules specified in the ACL. ACLs applied at the context level are called administrative ACLs; only packets sent to the kernel are subject to those ACLs. The format and function of ACLs are the same regardless of whether they are applied to kernel-bound packets or traffic sent or received over interfaces.

IPv4 and IPv6 ACL rules can either permit or deny a packet based on the following match criteria:

- IP header fields—Source address, destination address, protocol, DSCP, ToS, total length
- TCP—Port, flags (RST or ACK)
- UDP—Port
- ICMP—Type, code
- IGMP—Type

3.4.3 Administrative ACLs

You can configure administrative ACLs used in any context to protect the control plane from unwanted traffic.

The following example allows only BGP traffic in the bgp-only context:

```
context bgp-only
!
 no ip domain-lookup
!
 interface bgp-local-address loopback
  ip address 30.1.1.1/32
 no logging console
!
 ip access-list bgp-only
  seq 10 permit tcp any eq bgp host 30.1.1.1
  seq 20 permit tcp any host 30.1.1.1 eq bgp
!
 admin-access-group bgp-only in
```

Example 5

3.4.4 Malicious Traffic Detection

Some L3 security checks are performed implicitly by the forwarding plane, and others are performed when enabled through configuration. Malicious traffic detection is performed using a combination of implicit and configured checks.

3.4.5 Malicious Traffic Counters

The forwarding plane maintains counters for packets dropped due to implicit checks, ACLs, reassembly failures, and so on. Some counters are maintained at the circuit level, and others are maintained at the context level. The malicious traffic counters correspond to the malicious traffic alarms. Related counters are organized into a counter group (category) for display purposes.

Counter groups are maintained on a per-context basis. The supported counter groups are:

- Malformed-IP
- Malformed-Layer4
- Filtered
- Spoofed
- Failed-Reassembly
- Other

The Other group represents malicious-traffic drops not counted against the rest of the categories.

The following table shows how the drop reasons are grouped and the counter collection point.

Table 1

Drop Reason	Drop Counter Category	Collection Point
IP version other than v4 or v6	Malformed-IP	Circuit
Invalid IP header length	Malformed-IP	Circuit
Invalid IP total length	Malformed-IP	Circuit
Invalid IP checksum	Malformed-IP	Circuit
Invalid ICMP checksum	Malformed-Layer4	Context
Invalid UDP length	Malformed-Layer4	Context
Invalid TCP flag combinations (IPv4)	Filtered	Circuit
Fragments	Filtered	Circuit
IP options (IPv4)	Malformed-IP, Filtered	Circuit
IP filter ACL drops	Filtered	Circuit
Src Addr == Dst Addr	Filtered, Spoofed	Circuit
All ICMP or specific ICMP packets	Filtered	Circuit
Source/Destination address in a bogus network: 0.0.0.0/8, 255.255.255.255/32, 127.0.0.0/8, 224.0.0.0/4	Filtered Spoofed Null Route	Circuit Circuit Context
Reverse route	Spoofed	Circuit
Reassembly failures	Failed-Reassembly	Context
BGF realm drops	Other	Realm

4 Secure Service Layer

4.1 BGF

In addition to the core security functionality supported by the SmartEdge platform, BGF supports various security features in the media plane that protect the core network from security attacks. The following sections describe BGF security features.

4.1.1 Control Plane

BGF features are controlled primarily through the H.248 interface over SCTP towards the Ericsson SGC. This link is secure because the IP address is not published, and BGF acts as a client connecting to the Ericsson SGC published IP address and port number. Some security features are also enabled through the CLI.

4.1.2 Media Plane

The media plane supports both IPv4 and IPV6 pinholes. Any traffic to and from ports in the configured port range that do not have a pinhole setup are dropped.

The following media planespecific checks are performed implicitly:

- Dropping RTP packets with a version other than 2
- Dropping RTP packets less than the minimum length (12 bytes)
- Dropping TCP packets with invalid flag combinations (SYN + FIN is considered illegal)
- Dropping packets destined for invalid pinholes
- Dropping TCP/UDP packets that are longer than the configured L4 packet size
- Dropping UDP packets for a TCP pinhole and TCP packets for a UDP pinhole

4.1.2.1 Topology and IP Address Hiding

For all media traffic relayed by BGF, endpoints are exposed to only the media addresses of BGF. The media address and port number in SIP/SDP are modified by the SGC with the SDP provided by BGF in H.248 signaling. This prevents endpoints from knowing any core or other endpoint IP addresses and prevents direct attacks on them.

If the media plane receives a packet for an invalid pinhole (not adhering to the filtering criteria on the forwarding plane), it is dropped. If the port on which the media was received falls within the configured port range on the forwarding plane, and no port table entry is added on that port, the packet is dropped. If the received port is not within the range of the configured port range on forwarding plane, then the packet is sent to BSD and not forwarded. The packets that are sent to BSD are treated as non-VoIP traffic because they are not within the configured port range. These packets are counted as part of the general punted packet count on the SmartEdge OS.

ACL infrastructure allows you to configure policies to drop UDP or TCP traffic based on destination IP address and port number. You can manually configure the SmartEdge OS using ACLs to drop packets outside the configured VoIP port ranges for a given IP to prevent the packets being sent to BSD. TTL values for all forwarded packets are reset to 255 to hide network topology.

4.1.2.2 DoS and DDoS Prevention Mechanisms

The following mechanisms help prevent DoS and DDoS attacks:

- Rate limiting of traffic (control plane) destined to SmartEdge interface addresses is performed on the forwarding plane on a per-protocol basis (DHCP, ICMP, PPP, and so on). This is not configured through the CLI.
- Administrative ACLs (BSD kernel ACLs) prevent traffic from line cards from reaching the BSD kernel. This is configured through the CLI, per context.
- For a configured loopback IP address, received packets that are outside the configured port range are dropped by the forwarding plane.
- For the configured port range, packets received on any of the ports not allocated for a media pinhole are dropped by the forwarding plane.
- For a configured pinhole, SGC can specify the IP address and mask along with a port range to limit access before latching to a remote media address.
- DoS and DDoS prevention mechanisms exist in each application and protocol daemon. These mechanisms are tested using industry-standard vulnerability analysis tools such as Nmap, Nessus, and MuSecurity.

4.1.2.3 Protocol Validation

BGF supports pinholes for UDP- and TCP-based protocols. If a pinhole is created for a UDP-based protocol, and if TCP traffic was received on it, the TCP traffic is dropped and counted. Similarly, for a pinhole created for a TCP-based protocol, received UDP traffic is dropped and counted. If the pinhole is for RTP or RTCP traffic, any packets that do not have RTP version 2 are dropped and increment the malicious-traffic counters.

4.1.2.4 Bandwidth Policing

When a call is set up, the codecs used for the session are negotiated. BGF performs bandwidth policing on ingress traffic per established media stream. For each stream, all bandwidth policing parameters are provided by the SGC using (Tman package) H.248.53.

5 Security Protocols

IPsec tunnel mode is supported on ASE cards. Only traffic exchanged over line cards can be secured by IPsec; the NetBSD kernel does not support IPsec. However, SmartEdge OS control traffic can be secured by IPsec if the traffic travels over line-card ports.

IPv4 VPN traffic encryption is supported, as are the AES and 3DES encryption algorithms. Detection of anomalies and attacks is not supported.

The SmartEdge IPsec implementation is fully compliant with RFCs 2403 and 2405 and partially compliant with RFCs 4301-4309.

6 Security Alarms

6.1 Malicious Traffic Alarms

A single context-wide alarm against the aggregate of all the drop category counters is supported. You can configure a global alarm threshold that applies to all contexts.

An alarm is raised when the context-wide aggregate drop counter reaches or exceeds the configured high watermark within the configured interval. A raised alarm is cleared when the counter reaches or falls below the configured low watermark.

7 Security Logging

7.1 Malicious Traffic Logging

Malicious traffic logging is disabled by default. You can configure the forwarding plane to log packets dropped due to the various implicit and configured checks You can enable logging for each of the counter groups.

You can enable forwarding plane malicious-traffic logging independently of ACL packet logging, service troubleshooting, and forwarding plane packet logging.

Loggd supports the storage of traffic logs on local files or on syslog servers. The default file format is binary; text file format is supported.

Loggd enforces the restriction that malicious traffic log files be different from regular event log files and that the malicious traffic log file is unique for each context. Malicious traffic logs cannot be printed on the console. As with regular event log files, loggd can use up to seven files to store traffic logs. The maximum log file size is 1 MB; older log files are archived with gzip. When 80% of the total malicious-traffic log file capacity (7 MB) has been used, loggd generates an alert message indicating that the file capacity is about to be reached. When 100% of the capacity has been used, loggd generates an alert message indicating that the file capacity has reached the maximum.

8 Media Plane Attack Prevention

The following media plane attacks can be prevented using built-in (implicit) security checks and configurable security settings.

- **IP short header, UDP short header,** and**TCP short header attacks** - Packets with IP header length less than 20 bytes are dropped. This check is implicit.
- **Malformed IP header** - The following implicit checks help prevent malformed IP header attacks:
  - Drop packets with IP version other than IPv4 or IPv6
  - Drop packets with total length larger than the packet length
  - Drop packets with invalid IP header checksums
- **Unknown protocols** - For BGF media traffic, an implicit check ensures that only the negotiated protocols through H.248 signaling are allowed (either UDP or TCP). All other packets are dropped. In addition, you can configure ACLs to permit or deny protocols based on protocol ID.
- **IP fragmentation attacks** - You can use ACLs to disable fragmentation. For more information, refer to Configuring Malicious Traffic Detection and Monitoring, Reference [1]

  Non-locally-destined fragments are forwarded if they pass the implicit and configured security checks. If fragments are permitted and the DF bit is set, the fragment is accepted and retained unless **clear-df** has been configured on the egress interface.

  Locally destined fragments are reassembled. The forwarding plane has built-in protections to prevent fragments from using excessive resources. The forwarding plane discards all fragments associated with a fragment chain if the missing fragments have not been received within 5 seconds of the first fragment being received.
- **IP option attacks** - Only RR and RA options are allowed for forwarded packets. You can use an ACL to drop packets with any IPv4 options. For more details, refer to Configuring Malicious Traffic Detection and Monitoring, Reference [1].
- **TCP echo protocol attack** and **fraggle attack** - You can configure an ACL to drop these packets, and an implicit check exists to drop packets outside the configured port range for BGF traffic. (The echo port is outside the configurable port range.) For more details, see Configuring Malicious Traffic Detection and Monitoring, Reference [1].

- **Invalid TCP flag combinations** - For IPv4 traffic, you can configure an ACL to drop packets with the following TCP flag combinations: SYN+FIN, FIN+URG+PSH, and all zero. For BGF traffic, the SYN+FIN combination is disallowed. For more details, refer to Configuring Malicious Traffic Detection and Monitoring, Reference [1].

- **ICMP attacks** - Use ACLs to drop all ICMP packets or permit only ICMP packets of certain types. An implicit ICMP rate limit by the forwarding plane exists for ICMP packets targeted to the control plane.

- **TCP packet oversized** - The maximum L4 packet size is configurable for BGF traffic. For more details, refer to Configuring Malicious Traffic Detection and Monitoring, Reference [1].

- **Answering TCP packets from a multicast address** - You can configure ACLs to prevent this.

- **Listening sockets** - Well-known ports for protocols are open only if configured so.

- **UDP DoS with same source as destination IP address** - To prevent this, enable RPF checks (both strict and loose) on the interface. You can also use ACLs. For more details, refer to Configuring Malicious Traffic Detection and Monitoring, Reference [1].

- **TCP RST attacks**, **SYN attacks with IP spoofing**, and **SYN/ACK attacks with IP spoofing** - To prevent these attacks, enable RPF checks and use IPsec (with an ASE card). For more details, refer to Configuring Malicious Traffic Detection and Monitoring, Reference [1].

- **UDP flooding** - Prevent UDP flooding by enabling bandwidth policing for BGF traffic.

- **UDP attack on diagnostic ports (Pepsi attack)** - Prevent these attacks by configuring appropriate ACLs. In addition, diagnostic ports are outside the configurable port range for BGF traffic and therefore are dropped. For more details, refer to Configuring Malicious Traffic Detection and Monitoring, Reference [1].

- **Invalid source and destination** - Configure appropriate ACL rules to drop packets with an invalid source or destination address, such as 0.0.0.0 or 255.255.255.255.

- **LAND attack (TCP src addr/port == dst addr/port)** - For forwarded packets, you can use ACL or strict RPF to disallow src addr == dst addr. For more details, refer to Configuring Malicious Traffic Detection and Monitoring, Reference [1].

- **IP spoofing attack** - Enable RPF checks for IPv4 interfaces. For IPv6 subscriber circuits, enable IP source address validation. For more details, refer to Configuring Malicious Traffic Detection and Monitoring, Reference [1].

9 Security Testing

Ericsson uses industry-leading security test tools, such as Mu Dynamics, to constantly validate SmartEdge OS L2 to L7 services.

10 Known Vulnerabilities

None

Reference List

Ericsson

[1] *Configuring Malicious Traffic Detection and Monitoring*, 87/1543-CRA 119 1170/1 Uen