Security

In This Chapter

This chapter provides information to configure security parameters. Topics in this chapter include:

Authentication, Authorization, and Accounting

This chapter describes authentication, authorization, and accounting (AAA) used to monitor and control network access on routers. Network security is based on a multi-step process. The first step, authentication, validates a user's name and password. The second step is authorization, which allows the user to access and execute commands at various command levels based on profiles assigned to the user.

Another step, accounting, keeps track of the activity of a user who has accessed the network. The type of accounting information recorded can include a history of the commands executed, the amount of time spent in the session, the services accessed, and the data transfer size during the session. The accounting data can then be used to analyze trends, and also for billing and auditing purposes.

You can configure routers to use local, Remote Authentication Dial In User Service (RADIUS), or Terminal Access Controller Access Control System Plus (TACACS+) security to validate users who attempt to access the router by console, Telnet, or FTP. You can select the authentication order which determines the authentication method to try first, second, and third.

The router supports the following security features:

- RADIUS can be used for authentication, authorization, and accounting.
- TACACS+ can be used for authentication, authorization, and accounting.
- Local security can be implemented for authentication and authorization.

Figure 1 depicts end user access-requests sent to a RADIUS server. After validating the user names and passwords, the RADIUS server returns an access-accept message to the users on ALA-1 and ALA-2. The user name and password from ALA-3 could not be authenticated, thus access was denied.

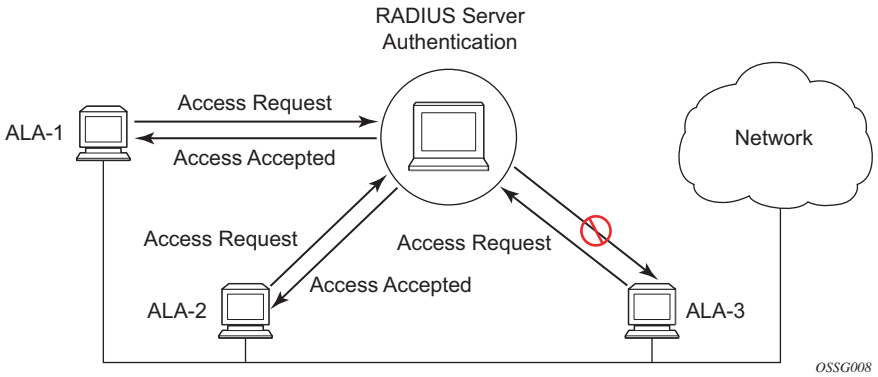

Figure 1: RADIUS Requests and Responses

Authentication

Authentication validates a user name and password combination when a user attempts to log in.

When a user attempts to log in through the console, Telnet, SSH, SCP, or FTP, the client sends an access request to a RADIUS, TACACS+, or local database.

Transactions between the client and a RADIUS server are authenticated through the use of a shared secret. The secret is never transmitted over the network. User passwords are sent encrypted between the client and RADIUS server which prevents someone snooping on an insecure network to learn password information.

If the RADIUS server does not respond within a specified time, the router issues the access request to the next configured servers. Each RADIUS server must be configured identically to guarantee consistent results.

If any RADIUS server rejects the authentication request, it sends an access reject message to the router. In this case, no access request is issued to any other RADIUS servers. However, if other authentication methods such as TACACS+ and/or local are configured, then these methods are attempted. If no other authentication methods are configured, or all methods reject the authentication request, then access is denied.

For the RADIUS server selection, round-robin is used if multiple RADIUS servers are configured. Although, if the first alive server in the list cannot find a user-name, the router does not re-query the next server in the RADIUS server list and denies the access request. It may get authenticated on the next login attempt if the next selected RADIUS server has the appropriate user-name. It is recommended that the same user databases are maintained for RADIUS servers in order to avoid inconsistent behavior.

The user login is successful when the RADIUS server accepts the authentication request and responds to the router with an access accept message.

Implementing authentication without authorization for the routers does not require the configuration of VSAs (Vendor Specific Attributes) on the RADIUS server. However, users, user access permissions, and command authorization profiles must be configured on each router.

Any combination of these authentication methods can be configured to control network access from a router:

- Local Authentication on page 24
- RADIUS Authentication on page 24
- TACACS+ Authentication on page 29

Local Authentication

Local authentication uses user names and passwords to authenticate login attempts. The user names and passwords are local to each router not to user profiles.

By default, local authentication is enabled. When one or more of the other security methods are enabled, local authentication is disabled. Local authentication is restored when the other authentication methods are disabled. Local authentication is attempted if the other authentication methods fail and local is included in the authentication order password parameters.

Locally, user names and password management information can be configured. This is referred to as local authentication. Remote security servers such as RADIUS or TACACS+, are not enabled.

RADIUS Authentication

Remote Authentication Dial-In User Service (RADIUS) is a client/server security protocol and software that enables remote access servers to communicate with a central server to authenticate dial-in users and authorize access to the requested system or service.

RADIUS allows you to maintain user profiles in a shared central database and provides better security, allowing a company to set up a policy that can be applied at a single administered network point.

RADIUS Server Selection

The RADIUS server selection algorithm is used by different applications:

- RADIUS operator management
- RADIUS authentication for Enhanced Subscriber Management
- RADIUS accounting for Enhanced Subscriber Management
- RADIUS PE-discovery

In all these applications, up to 5 RADIUS servers pools (per RADIUS policy, if used) can be configured.

The RADIUS server selection algorithm can work in 2 modes, either Direct mode or Round-robin mode.

Direct Mode

The first server is used as the primary server. If this server is unreachable, the next server, based on the server index, of the server pool is used. This continues until either all servers in the pool have been tried or an answer is received.

If a server is unreachable, it will not be used again by the RADIUS application for the next 30 seconds to allow the server to recover from its unreachable state. After 30 seconds the unreachable server is available again for the RADIUS application. If in these 30 seconds the RADIUS application receives a valid response for a previously sent RADIUS packet on that unreachable server, the server will be available for the RADIUS application again, immediately after reception of that response.

Round-Robin Mode

The RADIUS application sends the next RADIUS packet to the next server in the server pool. The same server non-reachability behavior is valid as in the Direct mode.

Server Reachability Detection

A server is reachable, when the operational state UP, when a valid response is received within a timeout period which is configurable by the retry parameter on the RADIUS policy level.

A server is treated as not-reachable, when the operational state down, when the following occurs:

- A timeout — If a number of consecutive timeouts are encountered for a specific server. This number is configurable by the retry parameter on RADIUS policy level.
- A send failed — If a packet cannot be sent to the RADIUS server because the forwarding path towards the RADIUS server is broken (for example, the route is not available, the is interface shutdown, etc.), then, no retry mechanism is invoked and immediately, the next server in line is used.

A server that is down can only be used again by the RADIUS algorithm after 30 seconds, unless, during these 30 seconds a valid RADIUS reply is received for that server. Then, the server is immediately marked UP again.

The operational state of a server can also be “unknown” if the RADIUS application is not aware of the state of the RADIUS server (for example, if the server was previously down but no requests had been sent to the server, thus, it is not certain yet whether the server is actually reachable).

Application Specific Behavior

Operator Management

The server access mode is fixed to Round-Robin (Direct cannot be configured for operator management). A health-check function is available for operator management, which can optionally be disabled. The health-check polls the server once every 10 seconds with an improbable user name. If the server does not respond to this health-check, it will be marked down.

If the first server in the list cannot find a user, the next server in the RADIUS server list is not queried and access is denied. If multiple RADIUS servers are used, it is assumed they all have the same user database.

RADIUS Authentication

If the first server in the list cannot find a user, the next server in the RADIUS server list is not queried and access is denied. If multiple RADIUS servers are used, it is assumed they all have the same user database.

RADIUS Challenge/Response Interactive Authentication

Challenge-response interactive authentication is used for key authentication where the Radius server is asking for the valid response to a displayed challenge. The challenge packet includes a challenge to be displayed to the user, such as a unique generated numeric value unlikely ever to be repeated. Typically this is obtained from an external server that knows what type of authenticator is in the possession of the authorized user and can therefore choose a random or non-repeating pseudorandom number of apropriate length.

The user then enters the challenge into his device (or software) and it calculates a response, which the user enters into the client which forwards it to the RADIUS server via an access request. If the response matches the expected response, the RADIUS server allows the user access, otherwise it rejects the response.

RADIUS challenge/response mode is enabled using the CLI interactive-authentication command in the config>system>security>radius context. RADIUS interactive authentication is disabled by default. The option needs to be enabled via CLI.

Enabling interactive authentication under CLI does not mean that the system uses RADIUS challenge/response mode by default. The configured password authentication-order parameter is used. If the authentication-order parameter is local RADIUS, the system will first attempt to

login the user via local authentication. If this fails, the system will revert to RADIUS and challenge/response mode. The authentication-order will precede the RADIUS interactive-authentication mode.

Even if the authentication-order is RADIUS local, the standard password prompt is always displayed. The user enters a username and password at this prompt. If RADIUS interactive-authentication is enabled the password does not have to be the correct password since authentication is accomplished using the RADIUS challenge/response method. The user can enter any password. The username and password are sent to the RADIUS server, which responds with a challenge request that is transmitted back to the node by the RADIUS server. Once the user enters the challenge response, the response is authenticated by the RADIUS server to allow node access to the user.

For example, if the system is configured with system security authentication-order set to local RADIUS, at the login prompt the user can enter the username “admin” and the corresponding password. If the password for local authentication does not match, the system falls into RADIUS authentication mode. The system checks the interactive-authentication configuration and if it is enabled it enters into challenge/response mode. It sends the username and password to the RADIUS server, and the server sends the challenge request back to the node and to the user where it appears as a challenge prompt onscreen. A challenge received from the RADIUS server typically contains a string and a hardware token that can be used to generate a password on the users’ local personal token generator. For example, the RADIUS server might send the challenge prompt "Enter response for challenge 12345:" to SR OS. The string “12345” can be entered in the local token generator which generates the appropriate challenge response for the entered string. This challenge response can then be entered on the SR-OS prompt for authorization.

Once the user enters the correct challenge response it is authenticated via the RADIUS server. The server authenticates the user and the user gains access to the node.

If session timeout and Idle timeout values are configured on the RADIUS server, these are used to govern the length of time before SR-OS cancels the challenge prompt. If the user is idle longer than the received idle-timeout (seconds) from the RADIUS server, and/or if the user does not press ENTER before the received session-timeout (seconds).

Note: For SSH only the session-timeout value is used. The SSH stack cannot track character input into the login prompt until the enter key is pressed.

Note: If the idle/session attribute is not available or if the value is set to a very large number, SR OS uses the smallest value set in "configure system login-control idle-timeout" and the idle/session timeout attribute value to terminate the prompt. If the “login-control idle-timeout” is set to 0 (equivalent to infinite), the maximum idle-timeout (24-hours) is used for the calculation.

SR-OS displays the log-in attempts/failure per user in the "show system security user user-name" screen. If the RADIUS rejects a challenge response, it counts as a failed login attempt and a new prompt is displayed. The number of failed attempts is limited by the value set for "configure

system security password attempt." An incorrect challenge response results in a failure count against the password attempts.

RADIUS Accounting

The RADIUS accounting application will try to send all the concerned packets of a subscriber host to the same server. If that server is down, then the packet is sent to the next server and, from that moment on, the RADIUS application uses that server to send its packets for that subscriber host.

RADIUS PE-Discovery

If the first server in the list cannot find a user, the next server in the RADIUS server list is not queried and access is denied. If multiple RADIUS servers are used, it is assumed they all have the same user database.

The RADIUS PE-discovery application makes use of a 10 second time period instead of the generic 30 seconds and uses a fixed consecutive timeout value of 2 (see Server Reachability Detection on page 25).

As long as the Session-Timeout (attribute in the RADIUS user file) is specified, it is used for the polling interval. Otherwise, the configured polling interval will be used (60 seconds by default).

TACACS+ Authentication

Terminal Access Controller Access Control System, commonly referred to as TACACS is an authentication protocol that allows a remote access server to forward a user's logon password to an authentication server to determine whether access can be allowed to a given system. TACACS is an encryption protocol and therefore less secure than the later Terminal Access Controller Access Control System Plus (TACACS+) and RADIUS protocols.

TACACS+ and RADIUS have largely replaced earlier protocols in the newer or recently updated networks. TACACS+ uses Transmission Control Protocol (TCP) and RADIUS uses the User Datagram Protocol (UDP). TACACS+ is popular as TCP is thought to be a more reliable protocol. RADIUS combines authentication and authorization. TACACS+ separates these operations.

Authorization

SR OS routers support local, RADIUS, and TACACS+ authorization to control the actions of specific users. Any combination of these authorization methods can be configured to control actions of specific users:

- Local Authorization on page 30
- RADIUS Authorization on page 30
- TACACS+ Authorization on page 31

Local authorization and RADIUS authorization operate by applying a profile based on user name and password configurations once network access is granted. The profiles are configured locally as well as VSAs on the RADIUS server. See Vendor-Specific Attributes (VSAs) on page 53.

Local Authorization

Local authorization uses user profiles and user access information after a user is authenticated. The profiles and user access information specifies the actions the user can and cannot perform.

By default, local authorization is enabled. Local authorization is disabled only when a different remote authorization method is configured, such as TACACS+ or RADIUS authorization.

You must configure profile and user access information locally.

RADIUS Authorization

RADIUS authorization grants or denies access permissions for a router. Permissions include the use of FTP, Telnet, SSH (SCP), and console access. When granting Telnet, SSH (SCP) and console access to the router, authorization can be used to limit what CLI commands the user is allowed to issue and which file systems the user is allowed or denied access.

Once a user has been authenticated using RADIUS (or another method), the router can be configured to perform authorization. The RADIUS server can be used to:

- Download the user profile to the router
- Send the profile name that the node should apply to the router.

Profiles consist of a suite of commands that the user is allowed or not allowed to execute. When a user issues a command, the authorization server looks at the command and the user information and compares it with the commands in the profile. If the user is authorized to issue the command, the command is executed. If the user is not authorized to issue the command, then the command is not executed.

Profiles must be created on each router and should be identical for consistent results. If the profile is not present, then access is denied.

Table 3 displays the following scenarios:

- Remote (RADIUS) authorization cannot be performed if authentication is done locally (on the router).
- The reverse scenario is supported if RADIUS authentication is successful and no authorization is configured for the user on the RADIUS server, then local (router) authorization is attempted, if configured in the authorization order.

When authorization is configured and profiles are downloaded to the router from the RADIUS server, the profiles are considered temporary configurations and are not saved when the user session terminates.

Table 3: Supported Authorization Configurations

	Router	RADIUS Supplied Profile
Routerconfigured user	Supported	Not Supported
RADIUS server configured user	Supported	Supported
TACACS+ server configured user	Supported	Not Supported

When using authorization, maintaining a user database on the router is not required. User names can be configured on the RADIUS server. User names are temporary and are not saved in the configuration when the user session terminates. Temporary user login names and their associated passwords are not saved as part of the configuration.

TACACS+ Authorization

TACACS+ authorization operates in one of three ways:

- All users who authenticate via TACACS+ can use a single common default profile that is configured on the SR OS Router, or
- Each command attempted by a user is sent to the TACACS+ server for authorization

- The operator can configure local profiles and map **tacplus priv-lvl** based authorization to those profiles (the **use-priv-lvl** option)

To use a single common default profile to control command authorization for TACACS+ users, the operator must configure the **tacplus use-default-template** option and configure the parameters in the **user-template tacplus_default** to point to a valid local profile.

If the default template is not being used for TACACAS+ authorization and the **use-priv-lvl** option is not configured, then each CLI command issued by an operator is sent to the TACACS+ server for authorization. The authorization request sent by SR OS contains the first word of the CLI command as the value for the TACACS+ cmd and all following words become a cmd-arg. Quoted values are expanded so that the quotation marks are stripped off and the enclosed value are seen as one cmd or cmd-arg.

Examples

Here is a set of examples, where the following commands are typed in the CLI:

```
- "show"
- "show router"
- "show port 1/1/1"
- "configure port 1/1/1 description "my port"
```

This results in the following AVPairs:

```
cmd=show

cmd=show
cmd-arg=router

cmd=show
cmd-arg=port
cmd-arg=1/1/1

cmd=configure
cmd-arg=port
cmd-arg=1/1/1
cmd-arg=description
cmd-arg=my port
```

For TACACS+ authorization, SR OS sends the entire CLI context in the **cmd** and **cmd-arg** values. Here is a set of examples where the CLI context is different:

```
- *A:dut-c# configure service
- *A:dut-c>config>service# vprn 555 customer 1 create
- *A:dut-c>config>service>vprn$ shutdown
```

This results in the following AVPairs:

```
cmd =configure
cmd-arg=service

cmd=configure
cmd-arg=service
cmd-arg=vprn
cmd-arg="555"
cmd-arg=customer
cmd-arg=1
cmd-arg=create

cmd=configure
cmd-arg=service
cmd-arg=vprn
cmd-arg="555"
cmd-arg=customer
cmd-arg=1
cmd-arg=create
cmd-arg=shutdown
```

Accounting

When enabled, RADIUS accounting sends command line accounting from the router to the RADIUS server. The router sends spars using UDP packets at port 1813 (decimal).

The router issues an accounting request packet for each event requiring the activity to be recorded by the RADIUS server. The RADIUS server acknowledges each accounting request by sending an accounting response after it has processed the accounting request. If no response is received in the time defined in the timeout parameter, the accounting request must be retransmitted until the configured retry count is exhausted. A trap is issued to alert the NMS (or trap receiver) that the server is unresponsive. The router issues the accounting request to the next configured RADIUS server (up to 5).

User passwords and authentication keys of any type are never transmitted as part of the accounting request.

RADIUS Accounting

Accounting tracks user activity to a specified host. When RADIUS accounting is enabled, the server is responsible for receiving accounting requests and returning a response to the client indicating that it has successfully received the request. Each command issued on the router generates a record sent to the RADIUS server. The record identifies the user who issued the command and the timestamp.

Accounting can be configured independently from RADIUS authorization and RADIUS authentication.

TACACS+ Accounting

The OS allows you to configure the type of accounting record packet that is to be sent to the TACACS+ server when specified events occur on the device. The accounting **record-type** parameter indicates whether TACACS+ accounting start and stop packets be sent or just stop packets be sent. Start/stop messages are only sent for individual commands, not for the session.

When a user logs in to request access to the network using Telnet or SSH, or a user enters a command for which accounting parameters are configured, or a system event occurs, such as a reboot or a configuration file reload, the router checks the configuration to see if TACACS+ accounting is required for the particular event.

If TACACS+ accounting is required, then, depending on the accounting record type specified, sends a start packet to the TACACS+ accounting server which contains information about the event.

The TACACS+ accounting server acknowledges the start packet and records information about the event. When the event ends, the device sends a stop packet. The stop packet is acknowledged by the TACACS+ accounting server.

Security Controls

You can configure routers to use RADIUS, TACACS+, and local authentication to validate users requesting access to the network. The order in which password authentication is processed among RADIUS, TACACS+ and local passwords can be specifically configured. In other words, the authentication order can be configured to process authorization through TACACS+ first, then RADIUS for authentication and accounting. Local access can be specified next in the authentication order in the event that the RADIUS and TACACS+ servers are not operational.

Table 4: Security Methods Capabilities

Method	Authentication	Authorization	Accounting*
Local	Y	Y	N
TACACS+	Y	Y	Y
RADIUS	Y	Y	Y

* Local commands always perform account logging using the **config log** command.

When a Server Does Not Respond

A trap is issued if a RADIUS + server is unresponsive. An alarm is raised if RADIUS is enabled with at least one RADIUS server and no response is received to either accounting or user access requests from any server.

Periodic checks to determine if the primary server is responsive again are not performed. If a server is down, it will not be contacted for 5 minutes. If a login is attempted after 5 minutes, then the server is contacted again. When a server does not respond with the health check feature enabled, the server's status is checked every 30 seconds. Health check is enabled by default. When a service response is restored from at least one server, the alarm condition is cleared. Alarms are raised and cleared on Alcatel-Lucent's Fault Manager or other third party fault management servers.

The servers are accessed in order from lowest to highest specified index (from 1 to 5) for authentication requests until a response from a server is received. A higher indexed server is only queried if no response is received, implying a lower indexed server is not available. If a response from the server is received, no other server is queried.

Access Request Flow

In Figure 2, the authentication process is defined in the `config>system>security>password` context. The authentication order is determined by specifying the sequence in which password authentication is attempted among RADIUS, TACACS+, and local passwords. This example uses the authentication order of RADIUS, then TACACS+, and finally, local. An access request is sent to RADIUS server 1. One of two scenarios can occur. If there is no response from the server, the request is passed to the next RADIUS server with the next lowest index (RADIUS server 2) and so on, until the last RADIUS server is attempted (RADIUS server 5). If server 5 does not respond, the request is passed to the TACACS+ server 1. If there is no response from that server, the request is passed to the next TACACS+ server with the next lowest index (TACACS+ server 2) and so on.

If a request is sent to an active RADIUS server and the user name and password is not recognized, access is denied and passed on to the next authentication option, in this case, the TACACS+ server. The process continues until the request is either accepted, denied, or each server is queried. Finally, if the request is denied by the active TACACS+ server, the local parameters are checked for user name and password verification. This is the last chance for the access request to be accepted.

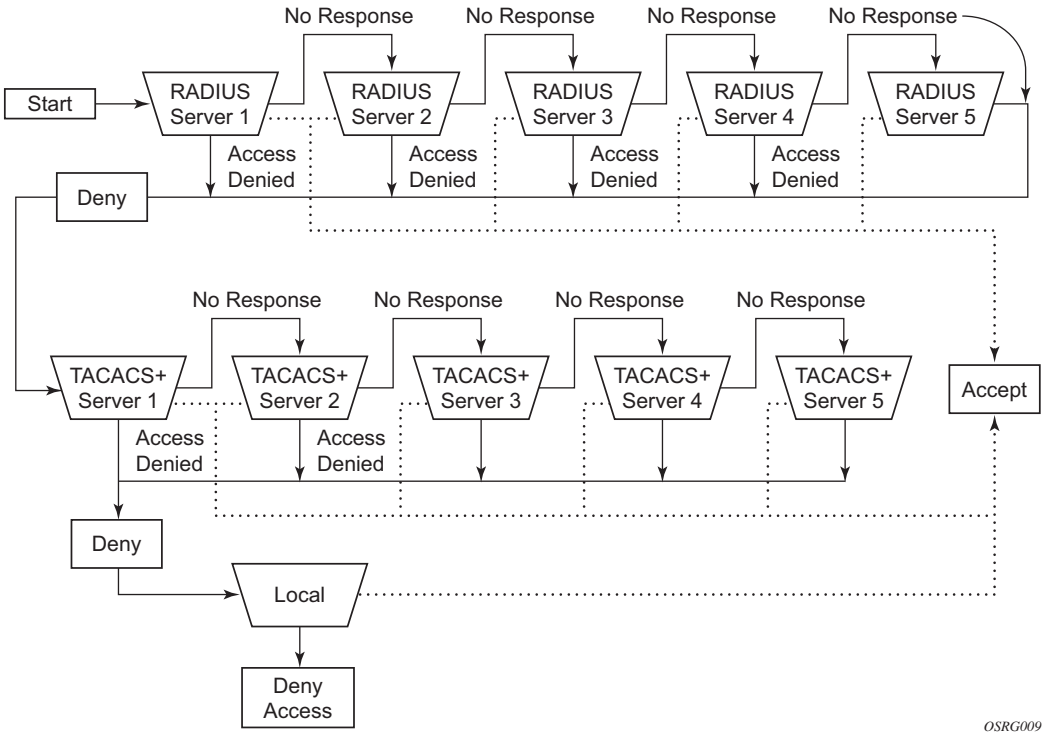

Figure 2: Security Flow

CPU Protection

SR OS provides several rate limiting mechanisms to protect the CPM/CFM processing resources of the router:

- CPU Protection: A centralized rate limiting function that operates on the CPM to limit traffic destined to the CPUs.
- Distributed CPU Protection: A control traffic rate limiting protection mechanism for the CPM/CFM that operates on the line cards (hence 'distributed').

CPU protection protects the CPU of the node that it is configured on from a DOS attack by limiting the amount of traffic coming in from one of its ports and destined to the CPM (to be processed by its CPU) using a combination of the configurable limits.

Some of the limits are configured globally for the node, and some of the limits are configured in CPU Protection profiles which are assigned to interfaces.

The following limits are configured globally for the node:

- link-specific rate — Applies to the link-specific protocols LACP (ethernet LAG control) and LMI (ATM, Ethernet and Frame Relay). The rate is a per-link limit (each link in the system will have LACP/LMI packets limited to this rate).
- port-overall-rate – Applies to all control traffic each port. The rate is a per-port limit (each port in the system will have control traffic destined to the CPM limited to this rate).
- protocol-protection — Blocks network control traffic for unconfigured protocols. If IS-IS is not configured on an IP interface all IS-IS-related traffic will be dropped and not reach the CPU.

The following limits are configured within CPU Protection policies (1-255). CPU Protection policies are created, configured, and then assigned to interfaces.

- overall-rate — Applies to all control traffic destined to the CPM (all sources) received on the interface (only where the policy is applied). This is a per-interface limit. Control traffic received above this rate will be discarded.
- per-source-rate — Used to limit the control traffic destined to the CPM from each individual source. This per-source-rate is only applied when an object (SAP) is configured with a cpu-protection policy and also with the optional mac-monitoring or ip-src-monitoring keywords. A source is defined as a *SAP, Source MAC Address* tuple for mac-monitoring and as a *SAP, Source IP Address* tuples for **ip-src**-monitoring. Only certain protocols (as configured under *included-protocols* in the cpu protection policy) are limited (per source) when the **ip-src**-monitoring keyword is used.
- out-profile-rate – Applies to all control traffic destined to the CPM (all sources) received on the interface (only where the policy is applied). This is a per-interface

limit. Control traffic received above this rate will be marked as discard eligible and is more likely to be discarded if there is contention for CPU resources.

A three-color marking mechanism uses a green, yellow and red marking function. This allows greater flexibility in how traffic limits are implemented. A CLI command within the DoS protection policy called **out-profile-rate** maps to the boundary between the green (accept) and yellow (mark as discard eligible) regions. The **overall-rate** command marks the boundary between the yellow and red (drop) regions point for the associated policy (Figure 3).

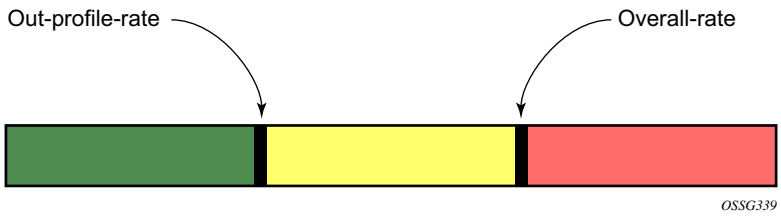

Figure 3: Profile Marking

There are two default CPU protection policies. They are modifiable, but cannot be deleted.

Policy 254:

- This is the default policy that is automatically applied to access interfaces
- Traffic above 6000 pps is discarded
- overall-rate = 6000
- per-source-rate = max
- out-profile-rate = 6000

Policy 255:

- This is the default policy that is automatically applied to Network interfaces
- Traffic above 3000 pps is marked as discard eligible, but is not discarded unless there is congestion in the queueing towards the CPU
- overall-rate = max
- per-source-rate = max
- out-profile-rate = 3000

All traffic destined to the CPM and that will be processed by its CPU will be subject to the limit specified. Therefore, if there is a protocol running on the violating interface, then protocol traffic on that interface will be affected. The objective of CPU protection is to limit the amount of traffic that the CPU will process at an early stage, therefore, the good and bad

traffic coming in cannot be distinguished when it arrives at a rate higher than the user-configured limit.

If the overall rate is set to 1000 pps and as long as the total traffic that is destined to the CPM and intended to be processed by the CPU is less than or equal to 1000 pps, all traffic will be processed. If the rate exceeds 1000 pps, then protocol traffic is discarded (or marked as discard eligible in the case of the out-profile-rate) and traffic on the interface is affected.

This protects all the other interfaces on the system and make sure that a violation from one interface does not affect the rest of the box.

The protocol-protection configuration is not a rate (just an enable/disable configuration). When enabled, this feature causes the network processor on the CPM to discard all packets received for protocols that are not configured on the particular interface. This helps mitigate DoS attacks by filtering invalid control traffic before it hits the CPU. The system automatically populates and maintains a per-interface list of configured (such as valid) protocols (based on interface config, etc). For example, if an interface does not have IS-IS configured, then protocol-protection will discard any IS-IS packets received on that interface.

Some protocols are not bound to a specific interface, for example, BGP. SR-OS will discard packets for these protocols if the protocol is not configured anywhere in the system. Note that protection for the following protocols is achieved using the per-peer-queueing feature of SR-OS: BGP, T-LDP, LDP, MSDP.

Protocols controlled by the protocol-protection mechanism include:

- OSPFv2
- OSPFv3
- IS-IS
- RSVP-TE
- RIP
- PIM
- MLD
- IGMP
- L2TP
- PPP

Note: If PIM or PIM snooping is not configured on any interfaces/SAPs then all PIM packets will be discarded. If PIM or PIM snooping is configured on an interface/SAP, then multicast PIM messages are filter based on PIM being enabled on that particular interface. All unicast PIM messages are sent to the CPU to be processed.

The CPU protection features are supported on the following platforms:

- 7750 SR-7/SR-12
- 7450 ESS-6/ESS-7/ESS-12
- 7950 XRS

The CPU protection features are **not** supported on the following platforms:

- 7450 ESS-1
- 7710 SR-c4/c12
- 7750 SR-c4/c12

CPU Protection Extensions ETH-CFM

CPU protection has been extended to provide the ability to explicitly limit the amount of ETH-CFM traffic that arrives at the CPU for processing. ETH-CFM packets that are redirected to the CPU by either a Management Endpoint (MEP) or a Management Intermediate Point (MIP) will be subject to the configured limit of the associated policy. Up to four CPU protection policies may include up to ten individual eth-cfm specific entries. The eth-cfm entries allow the operator to apply a packet per second rate limit to the matching combination of level and opcode, for eth-cfm packet that are redirected to the CPU. Any eth-cfm traffic that is redirected to the CPU by a Management Point (MP) that does not match any entries of the applied policy is still subject to the overall rate limit of the policy itself. Any eth-cfm packets that are not redirected to the CPU are not subject to this function and are treated as transit data, subject to the applicable QoS policy.

The operator first creates a CPU Policy and includes the required eth-cfm entries. Overlap is allowed for the entries within a policy, first match logic is applied. This means ordering the entries in the proper sequence is important to ensure the proper behavior is achieved. Even thought the number of eth-cfm entries is limited to ten, the entry numbers have a valid range from 1-100 to allow for ample space to insert policies between one and other.

Ranges are allowed when configuring the Level and the OpCode. Ranges provide the operator a simplified method for configuring multiple combinations. When more than one Level or OpCode is configured in this manner the configured rate limit is applied separately to each combination of level and OpCode match criteria. For example, if the Levels are configured with using a range of 5-7 and the OpCode is configured for 3,5 with a rate of 1. That restricts all possible combinations on that single entry to a rate of 1 packet per second. In this example six different match conditions are programmed behind the scene.

Table 5: Ranges versus Levels and OpCodes

Level	OpCode	Rate
5	3	1
5	5	1
6	3	1
6	5	1
7	3	1
7	5	1

Once the policy is created it must be applied to a SAP/Binding within a service for these rates to take affect. This means the rate is on a per SAP/Binding basis. Only a single policy may be applied to a SAP/Binding. The “eth-cfm-monitoring” option must be configured in order for the eth-cfm entries to be applied when the policy is applied to the SAP/Binding. If this option

is not configured, eth-cfm entries in the policy will be ignored. It is also possible to apply a policy to a SAP/Binding configuring “eth-cfm-monitoring” which does not have an MP. In this case, although these entries are enforced, no packets are being redirect to the CPU due to the lack of an MP.

By default, rates are applied on a per peer basis. This means each individual peer is subject to the rate. However, it is suggested that the “aggregate” option be configured to apply the rate to the sum total of all peers. MIPs for example only respond to Loopback Messages and Linktrace Messages. These are typically on demand functions and per peer rate limiting is likely not required thus making the aggregate function a more appealing model.

“eth-cfm-monitoring” and “mac-monitoring” are mutually exclusive and cannot be configured on the same SAP/Binding “mac-monitoring” is used in combination with the traditional CPU protection and is not specific to the eth-cfm rate limiting feature describe here.

When an MP is configured on a SAP/Binding within a service which allows an external source to communicate with that MP, for example a User to Network Interface (UNI), it is suggested that “eth-cfm-monitoring” with the “aggregate” option be configured on all SAP/Bindings to provide the highest level of rate control.

The example below shows a sample configuration for a policy and the application of that policy to a SAP in a VPLS service configured with a MP.

Policy 1 entry 10 limits all eth-cfm traffic redirected to the CPU for all possible combinations to 1 packet per second. Policy 1 entry 20 limits all possible combinations to a rate of zero, dropping all request which match any combination. If entry 20 did not exist then only rate limiting of the entry 10 matches would occur and any other eth-cfm packets redirected to the CPU would not be bound by a CPU protection rate.

```
config>sys>security>cpu-protection#
  policy 1
    eth-cfm
      entry 10 level 5-7 opcode 3,5 rate 1
      entry 20 level 0-7 opcode 0-255 rate 0

config>service>vpls#
  sap 1/1/4:100
    cpu-protection 1 eth-cfm-monitoring aggregate
    eth-cfm
      mip
    no shutdown
```

IOM1s are restricted to Down MEPs and ingress MIP for this feature. This feature is not supported on UP MEPs and egress MIPs for this IOM type.

ETH-CFM Ingress Squelching

CPU protection provides a granular method to control which ETH-CFM packets are processed. As indicated in the previous section, a unique rate can be applied to ETH-CFM packets classifying on specific MD-Level and specific OpCode and applied to both ingress (Down MEP and ingress MIP) and egress (Up MEP and egress MIP) extraction. That function is to protect the CPU upon extraction when a Management Point (MP) is configured.

It is also important to protect the ETH-CFM architecture deployed in the service provider network. The protection scheme here varies form CPU protection. This model is used to prevent ETH-CFM frames at the service provider MD-levels from gaining access to the network even when extraction is not in place. ETH-CFM squelching allows the operator to achieve this goal using a simple method to drop all ETH-CFM packets at or below the configured MD-level. The ETH-CFM squelch feature is ingress only.

Figure 4 shows a typical ETH-CFM hierarchical model with a Subscriber ME (6), Test ME (5), EVC ME (4) and an Operator ME (2). This model provides the necessary transparency at the different levels of the architecture. For security reasons, it may be necessary to prevent errant levels from entering the service provider network at the UNI, ENNI, or other untrusted interconnection points. Configuring squelching at level four on both UNI-N interconnection ensures that ETH-CFM packets matching the SAP or binding delimited configuration will silently discard ETH-CFM packets at ingress.

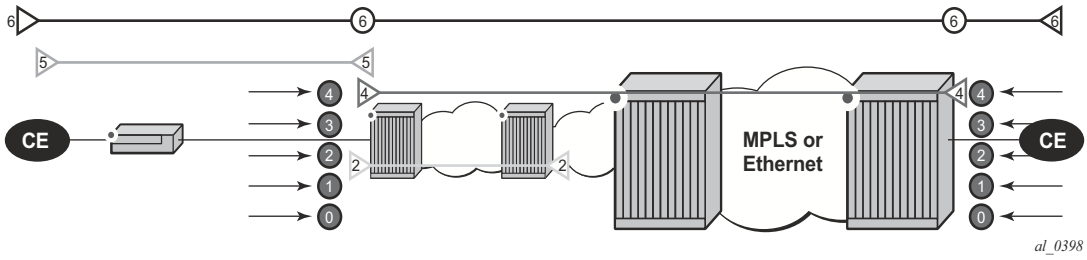

Figure 4: ETH-CFM Hierarchical Model

Squelching configuration uses a single MD-level [0..7] to silently drop all ETH-CFM packets matching the SAP or binding delimited configuration at and below the specified MD-level. In Figure 4, a squelch level is configured at MD-level 4. This means the configuration will silently discard MD-levels 0,1,2,3 and 4, assuming there is a SAP or binding match.

Note: Extreme caution must be used when deploying this feature.

The operator is able to configure Down MEPs and ingress MIPs that conflict with the squelched levels. This also means that any existing MEP or MIP processing ingress CFM packets on a SAP on Binding where a squelching policy is configured will be interrupted as

soon as this command is entered into the configuration. These MPs will not be able to receive any ingress ETH-CFM frames because squelching is processed before ETH-CFM extraction.

CPU Protection Extensions for ETH-CFM are still required in the model above because the Subscriber ME (6) and the Test ME (5) are entering the network across an untrusted connection, the UNI. ETH-CFM squelching and CPU Protection for ETH-CFM can be configured on the same SAP or binding. Squelching is first in the process order followed by CPU Protection for ETH-CFM.

MPs configured to support primary VLAN are not subjected to the squelch function. Primary VLAN based MPs, supported only on Ethernet SAPs, are extractions that take into consideration an additional VLAN beyond the SAP configuration.

The difference in the two protection mechanisms is shown in the Table 6. CPU Protection is used to control access to the CPU resources when processing is required. Squelching is required when the operator is protecting the ETH-CFM architecture from external sources.

Table 6: CPU PRotection and Squelching

Description	CPU Protection Extension for ETH-CFM	ETH-CFM Squelching
Ingress Filtering	Yes	Yes
Egress Filtering	Yes	No
Granularity	Specified Level AND OpCode	Level (At and below)
Rate	Configurable Rate (includes 0=drop all)	Silent Drop
Primary VLAN Support	Rate shared with SAP delineation	Not exposed to squelch
Extraction	Requires MEP or MIP to extract	No MEP or MIP required

As well as including the squelching information under the **show service** *service-id* **all**, display output the **squelch-ingress-level** key has been added to the **sap-using** and **sdp-using show** commands.

```
show service sap-using squelch-ingress-levels
===============================================================================
ETH-CFM Squelching
===============================================================================
PortId             SvcId      Squelch Level
-------------------------------------------------------------------------------
6/1/1:100.*        1          0 1 2 3 4 5 6 7
```

```
lag-1:100.*        1          0 1 2 3 4
6/1/1:200.*        2          0 1 2
lag-1:200.*        2          0 1 2 3 4 5
-------------------------------------------------------------------------------
Number of SAPs: 4
-------------------------------------------------------------------------------
===============================================================================
show service sdp-using squelch-ingress-levels
===============================================================================
ETH-CFM Squelching
===============================================================================
SdpId              SvcId       Type Far End                Squelch Level
-------------------------------------------------------------------------------
12345:4000000000   2147483650   Spok 1.1.1.1               0 1 2 3 4
===============================================================================
```

Extreme caution must be used when deploying this feature.

Distributed CPU Protection (DCP)

SR OS provides several rate limiting mechanisms to protect the CPM/CFM processing resources of the router:

- CPU Protection: A centralized rate limiting function that operates on the CPM to limit traffic destined to the CPUs. This feature is described elsewhere in this guide.
- Distributed CPU Protection: A control traffic rate limiting protection mechanism for the CPM/CFM that operates on the line cards (hence 'distributed').

Distributed CPU Protection (DCP) offers a powerful per-protocol-per-object (examples of objects are SAPs and network interfaces) rate limiting function for control protocol traffic that is extracted from the data path and sent to the CPM. The DCP function is implemented on the router line cards that allows for high levels of scaling and granularity of control.

The DCP rate limiting is configured via policies that are applied to objects (for example, SAPs).

The basic types of policers in DCP are:

- Enforcement Policers — An instance of a policer that is policing a flow of packets comprised of a single (or small set of) protocols(s) arriving on a single object (for example, SAP). Enforcement policers perform a configurable action (for example, discard) on packets that exceed configured rate parameters. There are two basic sub-types of enforcement policers:
  - → Static policers — always instantiate.
  - → Dynamic policers — only instantiated (allocated from a free pool of dynamic policers) when a local monitor detects non-conformance for a set of protocols on a specific object.
- Local Monitors — A policer that is primarily used to measure the conformance of a flow comprised of multiple protocols arriving on a single object. Local monitors are used as a trigger to instantiate dynamic policers.

The use of dynamic policers reduces the number of policers required to effectively monitor and control a set of protocols across a large set of objects since the per-protocol-per-object dynamic policers are only instantiated when an attack or misconfiguration occurs, and they are only instantiated for the affected objects.

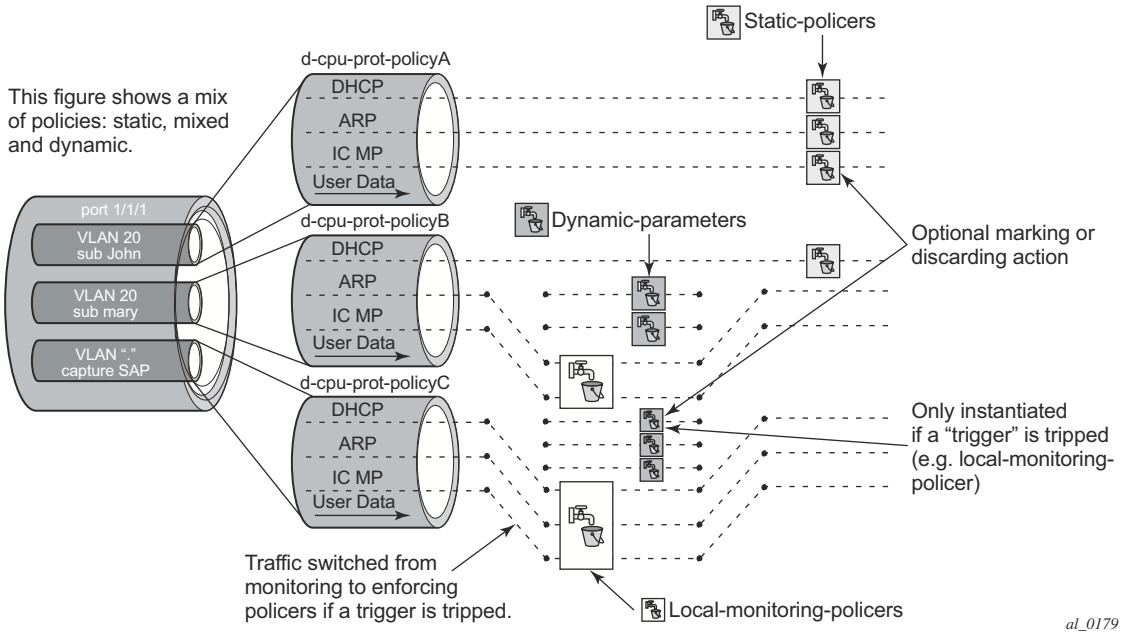

Figure 5: Per SAP per Protocol Static Rate Limiting with DCP

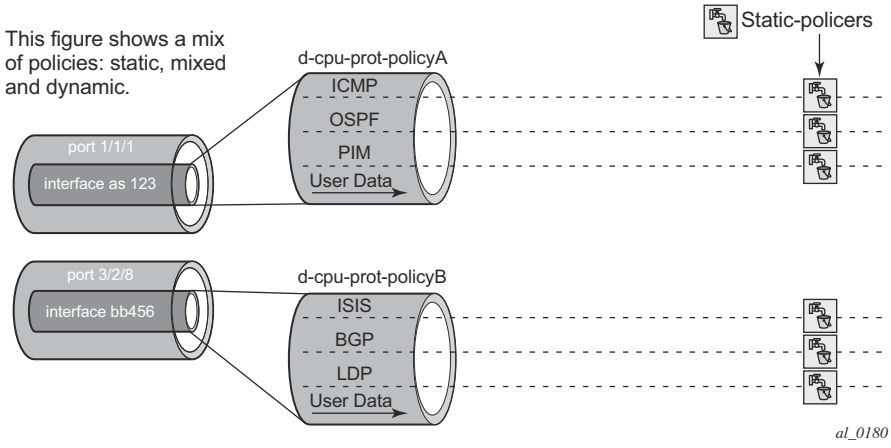

Figure 6: Per Network Interface per Protocol Static Rate Limiting with DCP

Applicability of Distributed CPU Protection

dist-cpu-protection (DCP) policies can be applicable to the following types of objects:

- most types of SAPs, including capture SAPs and SAPs on pseudo wires, but it is not applicable to b-vpls saps (b-saps).
- Network Interfaces, but not to any other type of interface. A DCP policy can be configured at the interface sap instead.

Control packets that are both forwarded (which means they could be subject to normal QoS policy policing) and also copied for extraction are not subject to Distributed CPU Protection (including in the all-unspecified bucket). This includes traffic snooping (for example, PIM in VPLS) as well as control traffic that is flooded in an R-VPLS instance and also extracted to the CPM such as ARP, ISIS and VRRP. Centralized per SAP/interface cpu-protection can be employed to rate limit or mark this traffic if desired.

Control traffic that arrives on a network interface, but inside a tunnel (for example, SDP, LSP, PW) and logically terminates on a service (that is, traffic that is logically extracted by the service rather than the network interface layer itself) will bypass the DCP function. The control packets in this case will not be subject to the DCP policy that is assigned to the network interface on which the packets arrived. This helps to avoid customer traffic in a service from impacting other services or the operator's infrastructure.

Control packets that are extracted in a vprn service, where the packets arrived into the node via a vpls SAP (that is, r-vpls scenario), will use the DCP policy and policer instances associated with the vpls SAP. In this case the DCP policy that an operator creates for use on VPLS SAPs, for VPLSs that have a l3-interface bound to them (r-vpls), may have protocols like OSPF, ARP, configured in the policy.

Log Events, Statistics, Status and SNMP support

A comprehensive set of log events are supported for DCP in order to alert the operator to potential attacks or misconfigurations and to allow tuning of the DCP settings. Refer to the NOTIFICATION-TYPE objects with “Dcp” in the names in the following MIBs for details:

- TIMETRA-CHASSIS-MIB
- TIMETRA-SAP-MIB
- TIMETRA-VRTR-MIB

The log events can also be seen in the CLI using the following **show log event-control | match Dcp** command

DCP throttles the rate of DCP events to avoid event floods when multiple parallel attacks or problems are occurring.

Many of the DCP log events can be individually enabled or disabled at the DCP policy level (in the DCP policy config) as well as globally in the system (in log event-control).

If needed when a DCP log event indicates a SAP, and that SAP is an MSAP, the operator can determine which subscriber(s) is/are on a specific MSAP by using the **show service active-subs** command and then filtering (“| match”) on the msap string.

Statistics and status related to DCP are available both via:

- CLI
- SNMP — See various tables and objects with “Dcp” or “DCpuProt” in their name in the TIMETRA-CHASSIS-MIB¸ TIMETRA-SECURITY-MIB, TIMETRA-SAP-MIB and TIMETRA-VRTR-MIB

DCP Policer Resource Management

The policer instances are a limited h/w resource on a given forwarding plane. DCP policers (static, dynamic, local-monitor) are consumed from the overall forwarding plane policer resources (from the ingress resources if ingress and egress are partitioned). Each per-protocol policer instantiated reduces the number of FP child policers available for other purposes.

When DCP is configured with dynamic enforcement, then the operator must set aside a pool of policers that can be instantiated as dynamic enforcement policers. The number of policers reserved for this function are configurable per card/fp. The policers in this pool are not available for other purposes (normal SLA enforcement).

Static enforcement policers and local monitoring policers use policers from the normal/global policer pool on the card/fp. Once a static policer is configured in a DCP policy and it is referenced by a protocol in the policy, then this policer will be instantiated for each object (SAP or network interface) that is created and references the policy. If there is no policer free on the associated card/fp, then the object will be blocked from being created. Similarly for local monitors: once a local monitoring policer is configured and referenced by a protocol, then this policer will be instantiated for each object that is created and references the policy. If there is no policer free, then the object will be blocked from being created.

Dynamic enforcement policers are allocated as needed (when the local monitor detects non-conformance) from the reserved dynamic-enforcement-policer-pool.

When a DCP policy is applied to an object on a LAG, then a set of policers is allocated on each forwarding plane (on each line card that contains a member of the LAG). The LAG mode is ignored and the policers are always shared by all ports in the LAG on that forwarding plane on the SAP/interface. In other words, with link-mode lag a set of DCP policers are not allocated per port in the LAG on the SAP.

In order to support large scale operation of DCP, and also to avoid overload conditions, a polling process is used to monitor state changes in the policers. This means there can be a delay between when an event occurs in the data plane and when the relevant state change or event notification occurs towards an operator, but in the meantime the policers are still operating and protecting the control plane.

Operational Guidelines and Tips

The following points offer various optional guidelines that may help an operator decide how to leverage Distributed CPU Protection.

- The rates in a policy assigned to a capture SAP should be higher than those assigned to MSAPs that will contain a single subscriber. The rates for the capture sap policy should allow for a burst of MSAP setups.
- To completely block a set of specific protocols on a given SAP, create a single static policer with a rate of 0 and map the protocols to that policer. Dynamic policers and local monitors can't be used to simultaneously allow some protocols but block others (the non-zero rates in the monitor would let all protocols slip through at a low rate).
- During normal operation it is recommended to configure "log-events" (no verbose keyword) for all static-policers, in the dynamic-parameters of all protocols and for all local-monitoring-policers. he verbose keyword can be used selectively during debug, testing, tuning and investigations.
- Packet based rate limiting is generally recommended for low rate subscriber based protocols whereas kbps rate limiting is recommended for higher rate infrastructure protocols (such as BGP).
- It is recommended to configure an exceed-action of low-priority for routing and infrastructure protocols. Marked packets are more likely to be discarded if there is congestion in the control plane of the router, but will get processed if there is no contention for CPU resources allowing for a work-conserving behavior in the CPM.
- In order to assign a different dist-cpu-protection policy to a specific MSAP (instance) or to all MSAPs for a specific msap policy, the operator can assign a new dist-cpu-protection policy to the MSAP policy and then use the **eval-msap** tool:

  A:nodeA>tools>perform# subscriber-mgmt eval-msap
  - eval-msap { policy <msap-policy-name> | msap <sap-id> }

  Note that any new MSAPs will also be assigned the new dist-cpu-protection policy.
- If needed, an operator can determine which subscriber is on a specific MSAP by using the **show service active-subs** command and then filtering ("| match") on the msap string.
- If protocol X is trusted, and using the "all-unspecified" protocol is not required, then simply avoid creating protocol X in the policy configuration.
- If protocol X is trusted, but the all-unspecified bucket is required, then there are two options:
  - → avoid creating protocol X so that it is treated as part of the all-unspecified bucket (but account for the packets from X in the all-unspecified rate and local-mon rate), or
  - → create protocol X and configure it to bypass.

Vendor-Specific Attributes (VSAs)

The software supports the configuration of Alcatel-Lucent-specific RADIUS attributes. These attributes are known as vendor-specific attributes (VSAs) and are discussed in RFC 2138. VSAs must be configured when RADIUS authorization is enabled. It is up to the vendor to specify the format of their VSA. The attribute-specific field is dependent on the vendor's definition of that attribute. The Alcatel-Lucent-defined attributes are encapsulated in a RADIUS vendor-specific attribute with the vendor ID field set to 6527, the vendor ID number.

Note that the PE-record entry is required in order to support the RADIUS Discovery for Layer 2 VPN feature. Note that a PE-record is only relevant if the RADIUS Discovery feature is used, not for the standard RADIUS setup.

The following RADIUS vendor-specific attributes (VSAs) are supported by Alcatel-Lucent.

- `timetra-access <ftp> <console> <both>` — This is a mandatory command that must be configured. This command specifies if the user has FTP and /or console (serial port, Telnet, and SSH) access.
- `timetra-profile <profile-name>` — When configuring this VSA for a user, it is assumed that the user profiles are configured on the local router and the following applies for local and remote authentication:

  1. The `authentication-order` parameters configured on the router must include the `local` keyword.
  2. The user name may or may not be configured on the router.
  3. The user must be authenticated by the RADIUS server
  4. Up to 8 valid profiles can exist on the router for a user. The sequence in which the profiles are specified is relevant. The most explicit matching criteria must be ordered first. The process stops when the first complete match is found.

  If all the above mentioned conditions are not met, then access to the router is denied and a failed login event/trap is written to the security log.
- `timetra-default-action <permit-all|deny-all|none>` — This is a mandatory command that must be configured even if the `timetra-cmd` VSA is not used. This command specifies the default action when the user has entered a command and no entry configured in the `timetra-cmd` VSA for the user resulted in a match condition.
- `timetra-cmd <match-string>` — Configures a command or command subtree as the scope for the match condition.

  The command and all subordinate commands in subordinate command levels are specified.

Other Security Features

Secure Shell (SSH)

Secure Shell Version 1 (SSH) is a protocol that provides a secure, encrypted Telnet-like connection to a router. A connection is always initiated by the client (the user). Authentication takes places by one of the configured authentication methods (local, RADIUS, or TACACS+). With authentication and encryption, SSH allows for a secure connection over an insecure network.

The OS allows you to configure Secure Shell (SSH) Version 2 (SSH2). SSH1 and SSH2 are different protocols and encrypt at different parts of the packets. SSH1 uses server as well as host keys to authenticate systems whereas SSH2 only uses host keys. SSH2 does not use the same networking implementation that SSH1 does and is considered a more secure, efficient, and portable version of SSH.

SSH runs on top of a transport layer (like TCP or IP), and provides authentication and encryption capabilities.

The OS has a global SSH server process to support inbound SSH and SCP sessions initiated by external SSH or SCP client applications. The SSH server supports SSHv1. Note that this server process is separate from the SSH and SCP client commands on the routers which initiate outbound SSH and SCP sessions.

Inbound SSH sessions are counted as inbound telnet sessions for the purposes of the maximum number of inbound sessions specified by Login Control. Inbound SCP sessions are counted as inbound ftp sessions by Login Control.

When SSH server is enabled, an SSH security key is generated. The key is only valid until either the node is restarted or the SSH server is stopped and restarted (unless the preserve-key option is configured for SSH). The key size is non-configurable and set at 1024 bits. When the server is enabled, both inbound SSH and SCP sessions will be accepted provided the session is properly authenticated.

When the global SSH server process is disabled, no inbound SSH or SCP sessions will be accepted.

When using SCP to copy files from an external device to the file system, the SCP server will accept either forward slash (“/”) or backslash (“\”) characters to delimit directory and/or filenames. Similarly, the SCP client application can use either slash or backslash characters, but not all SCP clients treat backslash characters as equivalent to slash characters. In particular, UNIX systems will often times interpret the backslash character as an “escape” character which does not get transmitted to the SCP server. For example, a destination

directory specified as “cf1:\dir1\file1” will be transmitted to the SCP server as “cf1:dir1file1” where the backslash escape characters are stripped by the SCP client system before transmission. On systems where the client treats the backslash like an “escape” character, a double backslash “\\” or the forward slash “/” can typically be used to properly delimit directories and the filename.

Two cipher lists, the client-cipher-list and the server-cipher-list, can be configured for negotiation of the best compatible ciphers between the the client and server. The two cipher lists can be created and managed under the security ssh sub menu. The client-cipher-list is used when SR OS is acting as ssh client and the server-cipher-list is used when the SR OS is acting as a server. The first cipher matched on the lists between the client and server is the preferred cipher for the session.

SSH PKI Authentication

The SR OS supports Secure Shell Version 2, but user authentication appears to be limited to using a username and password.

Note: SSHv1 is not supported when the node is running in FIPS-140-2 mode.

SSH also supports public key authentication whereby the client can provide a signed message that has been encrypted by his private key. As long as the server has been previously configured to know the client's public key, the server can authenticate the client.

Using Public Key authentication (also known as Public Key Infrastructure - PKI) can be more secure than the existing username/password method for a few reasons:

- A user will typical re-use the same password with multiple servers. If the password is compromised, the user must reconfigure the password on all affected servers.
- A password is not transmitted between the client and server using PKI. Instead the sensitive information (the private key) is kept on the client. Therefore it is less likely to be compromised.

This feature includes server side support for SSHv2 public key authentication. It does not include a key generation utility.

Support for PKI should be configured in the system level configuration where one or more public keys may be bound to a username. It should not affect any other system security or login functions.

Key Generation

Before SSH can be used with PKI, someone must generate a public/private key pair. This is typically supported by the SSH client software. For example, PuTTY supports a utility called PuTTYgen that will generate key pairs.

SSHv2 supports both RSA and DSA keys. The Digital Signature Algorithm is a U.S Federal Government standard for digital signatures. PuTTYGen can be used to generate either type of key. The SR OS currently supports only RSA keys.

Assume the client is using PuTTY. First the user generates a key pair using PuTTYgen. The user sets the key type (SSH-1 RSA, SS-2 RSA, or SSH-2 DSA) and sets the number of bits to be used for the key (default = 1024). The user can also configure a passphrase that will be used to store the key locally in encrypted form. If the passphrase is configured the user must enter the passphrase in order to use the private key. Thus, it is a password for the private key. If the passphrase is not used the key is stored in plaintext locally.

Next the user must configure the server to use his public key. This typically requires the user to add the public key to a file on the server. For example, if the server is using OpenSSH, the key must be added to the ssh/authorized_keys file. On the SR OS, the user can program the public Key via Telnet/SSH or SNMP.

Per Peer CPM Queuing

System-level security is crucial in service provider networks to address the increased threat of Denial-of-Service (DoS) attacks.

Control Processor Module Queuing (CPMQ) implements separate hardware-based queues which are allocated on a per-peer basis. CPMQ allocates a separate queue for each LDP and BGP peer and ensures that each queue is served in a round-robin fashion. This mechanism guarantees fair and “non-blocking” access to shared CPU resources across all peers. This would ensure, for example, that an LDP-based DoS attack from a given peer would be mitigated and compartmentalized so that not all CPU resources would be dedicated to the otherwise overwhelming control traffic sent by that specific peer.

CPMQ, using the “per-peer-queuing” command, ensures that service levels would not (or only partially be) impacted in case of an attack from a spoofed LDP or BGP peer IP address.

Per Peer CPM Queueing is supported on the 7750 SR-7/12 and 7750 SR-c12 platforms. -1.

CPM Filters and Traffic Management

Alcatel-Lucent routers have traffic management and queuing hardware dedicated to protecting the control plane.

CPM/CFM filters are supported on the following platforms: 7950 XRS, 7750 SR-7/SR-12/ SR-c12, and 7710 SR-c4/SR-c12. The filters can be used to drop or accept packets, as well as allocate dedicated hardware shaping (CPM) queues for traffic directed to the control processors.

Users can allocate dedicated CPM hardware queues for certain traffic designated to the CPUs and can set the corresponding rate-limit for the queues. CPM queueing is supported on the following platforms: 7950 XRS, 7750 SR-7/SR-12, and 7750 SR-c12.

CPM filters and queues control all traffic going in to the CPM from IOMs/XMAs, including all routing protocols. CPM filters apply to packets from all network and access ports, but not to packets from a management Ethernet port. CPM packet filtering and queuing is performed by network processor hardware using no resources on the main CPUs.

There are three filters that can be configured as part of the CPM filter policy: IP (v4) fitter, IPv6 filter and MAC filter.

The SROS filter implementation exits the filter when the first match is found and execute the actions according to the specified action. For this reason, entries must be sequenced correctly from most to least explicit. When both mac-filter and ip-filter/ipv6-filter are to be applied to a given traffic, mac-filter is applied first.

An entry of an IP(v4), IPv6, MAC CPM filters must have at least one match criteria defined to be active. A default action can be specified for CPM filter policy that applies to each of IP, IPv6, MAC filters that are in a **no shutdown** state as long as the CPM filter policy has at least one active filter entry in any of the IP(v4), IPv6, and MAC filters.

TTL Security for BGP and LDP

The BGP TTL Security Hack (BTSH) was originally designed to protect the BGP infrastructure from CPU utilization-based attacks. It is derived on the fact that the vast majority of ISP eBGP peerings are established between adjacent routers. Since TTL spoofing cannot be performed, a mechanism based on an expected TTL value can provide a simple and reasonably robust defense from infrastructure attacks based on forged BGP packets.

While TSH is most effective in protecting directly connected peers, it can also provide a lower level of protection to multi-hop sessions. When a multi-hop BGP session is required, the expected TTL value can be set to 255 minus the configured range-of-hops. This approach can provide a qualitatively lower degree of security for BGP (for example, a DoS attack could, theoretically, be launched by compromising a box in the path). However, BTSH will catch a vast majority of observed distributed DoS (DDoS) attacks against eBGP. For further information, refer to draft-gill-btsh-xx.txt, *The BGP TTL Security Hack (BTSH)*.

TSH can be used to protect LDP peering sessions as well. For details, see draft-chen-ldp-ttl-xx.txt, *TTL-Based Security Option for LDP Hello Message*.

The TSH implementation supports the ability to configure TTL security per BGP/LDP peer and evaluate (in hardware) the incoming TTL value against the configured TTL value. If the incoming TTL value is less than the configured TTL value, the packets are discarded and a log is generated.

Exponential Login Backoff

A malicious user may attempt to gain CLI access by means of a dictionary attack using a script to automatically attempt to login as an “admin” user and using a dictionary list to test all possible passwords. Using the exponential-backoff feature in the **config>system>login-control** context the OS increases the delay between login attempts exponentially to mitigate attacks.

A malicious user may attempt to gain CLI access by means of a dictionary attack using a script to automatically attempt to login as an “admin” user and using a dictionary list to test all possible passwords. Using the exponential-backoff feature in the config>system>login-control context the OS increases the delay between login attempts exponentially to mitigate attacks.

When a user tries to login to a router using a Telnet or an SSH session, there are a limited number of attempts allowed to enter the correct password. The interval between the unsuccessful attempts change after each try (1, 2 and 4 seconds). If the system is configured for user lockout, then the user will be locked out when the number of attempts is exceeded.

However, if lockout is not configured, there are three password entry attempts allowed after the first failure, at fixed 1, 2 and 4 second intervals, in the first session, and then the session terminates. Users do not have an unlimited number of login attempts per session. After each failed password attempt, the wait period becomes longer until the maximum number of attempts is reached.

The OS terminates after four unsuccessful tries. A wait period will never be longer than 4 seconds. The periods are fixed and will restart in subsequent sessions.

Note that the **config>system>login-control>**[**no**] **exponential-backoff** command works in conjunction with **the config>system>security>password>attempts** command which is also a system wide configuration.

For example:

```
*A:ALA-48>config>system# security password attempts
  - attempts <count> [time <minutes1>] [lockout <minutes2>]
  - no attempts

 <count>              : [1..64]
 <minutes1>           : [0..60]
 <minutes2>           : [0..1440]
```

Exponential backoff applies to any user and by any login method such as console, SSH and Telnet.

Refer to Configuring Login Controls on page 105. The commands are described in Login, Telnet, SSH and FTP Commands on page 125.

User Lockout

When a user exceeds the maximum number of attempts allowed (the default is 3 attempts) during a certain period of time (the default is 5 minutes) the account used during those attempts will be locked out for a pre-configured lock-out period (the default is 10 minutes).

An security event log will be generated as soon as a user account has exceeded the number of allowed attempts and the **show>system>security>user** command can be used to display the total number of failed attempts per user.

The account will be automatically re-enabled as soon as the lock-out period has expired. The list of users who are currently locked-out can be displayed with *show system security user lockout*.

A lock-out for a specific user can be administratively cleared using the *admin user x clear-lockout*.

Encryption

Data Encryption Standard (DES) and Triple DES (3DES) are supported for encryption.

- DES is a widely-used method of data encryption using a private (secret) key. Both the sender and the receiver must know and use the same private key.
- 3DES is a more secure version of the DES protocol.

802.1x Network Access Control

The Alcatel-Lucent OS supports network access control of client devices (PCs, STBs, etc.) on an Ethernet network using the IEEE. 802.1x standard. 802.1x is known as Extensible Authentication Protocol (EAP) over a LAN network or EAPOL.

TCP Enhanced Authentication Option

The TCP Enhanced Authentication Option, currently covered in draft-bonica-tcp-auth-05.txt, *Authentication for TCP-based Routing and Management Protocols*, extends the previous MD5 authentication option to include the ability to change keys without tearing down the session, and allows for stronger authentication algorithms to be used.

The TCP Enhanced Authentication Option is a TCP extension that enhances security for BGP, LDP and other TCP-based protocols. This includes the ability to change keys in a BGP or LDP session seamlessly without tearing down the session. It is intended for applications where secure administrative access to both the end-points of the TCP connection is normally available.

TCP peers can use this extension to authenticate messages passed between one another. This strategy improves upon current practice, which is described in RFC 2385, *Protection of BGP Sessions via the TCP MD5 Signature Option*. Using this new strategy, TCP peers can update authentication keys during the lifetime of a TCP connection. TCP peers can also use stronger authentication algorithms to authenticate routing messages.

Packet Formats

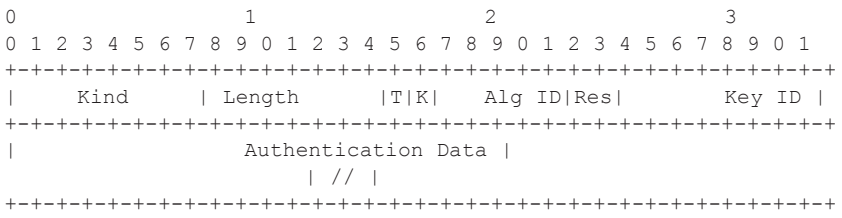

Option Syntax

- Kind: 8 bits

  The Kind field identifies the TCP Enhanced Authentication Option. This value will be assigned by IANA.
- Length: 8 bits

  The Length field specifies the length of the TCP Enhanced Authentication Option, in octets. This count includes two octets representing the Kind and Length fields.

  The valid range for this field is from 4 to 40 octets, inclusive.

  For all algorithms specified in this memo the value will be 16 octets.
- T-Bit: 1 bit

  The T-bit specifies whether TCP Options were omitted from the TCP header for the purpose of MAC calculation. A value of 1 indicates that all TCP options other than the Extended Authentication Option were omitted. A value of 0 indicates that TCP options were included.

  The default value is 0.
- K-Bit: 1 bit

  This bit is reserved for future enhancement. Its value MUST be equal to zero.
- Alg ID: 6 bits

  The Alg ID field identifies the MAC algorithm.

- Res: 2 bits

  These bits are reserved. They MUST be set to zero.

  Key ID: 6 bits

  The Key ID field identifies the key that was used to generate the message digest.
- Authentication Data: Variable length
- The Authentication Data field contains data that is used to authenticate the TCP segment. This data includes, but need not be restricted to, a MAC. The length and format of the Authentication Data Field can be derived from the Alg ID.
- The Authentication for TCP-based Routing and Management Protocols draft provides and overview of the TCP Enhanced Authentication Option. The details of this feature are described in draft-bonica-tcp-auth-04.txt.

Keychain

The keychain mechanism allows for the creation of keys used to authenticate protocol communications. Each keychain entry defines the authentication attributes to be used in authenticating protocol messages from remote peers or neighbors, and it must include at least one key entry to be valid. Through the use of the keychain mechanism, authentication keys can be changed without affecting the state of the associated protocol adjacencies for OSPF, IS-IS, BGP, LDP, and RSVP-TE.

Each key within a keychain must include the following attributes for the authentication of protocol messages:

- key identifier
- authentication algorithm
- authentication key
- direction
- start time

In addition, additional attributes can be optionally specified, including:

- end time
- tolerance

Table 7 shows the mapping between these attributes and the CLI command to set them.

Table 7: Keychain Mapping

Definition	CLI
The key identifier expressed as an integer (0...63)	config>system>security>keychain>direction>bi>entry config>system>security>keychain>direction>uni>receive>entry config>system>security>keychain>direction>uni>send>entry
Authentication algorithm to use with key[i]	config>system>security>keychain>direction>bi>entry with algorithm *algorithm* parameter. config>system>security>keychain>direction>uni>receive>entry with algorithm *algorithm* parameter. config>system>security>keychain>direction>uni>send>entry with algorithm *algorithm* parameter.
Shared secret to use with key[i].	config>system>security>keychain>direction>uni>receive>entry with shared secret parameter config>system>security>keychain>direction>uni>send>entry with shared secret parameter config>system>security>keychain>direction>bi>entry with shared secret parameter

Table 7: Keychain Mapping (Continued)

Definition	CLI
A vector that determines whether the key[i] is to be used to generate MACs for inbound segments, outbound segments, or both.	config>system>security>keychain>direction
Start time from which key[i] can be used.	config>system>security>keychain>direction>bi>entry>begin-time config>system>security>keychain>direction>uni>send>entry >begin-time
End time after which key[i] cannot be used by sending TCPs.	Inferred by the begin-time of the next key (youngest key rule).
Start time from which key[i] can be used.	config>system>security>keychain>direction>bi>entry>begin-time config>system>security>keychain>direction>bi>entry>tolerance config>system>security>keychain>direction>uni>receive>entry >begin-time config>system>security>keychain>direction>uni>receive>entry >tolerance
End time after which key[i] cannot be used	config>system>security>keychain>direction>uni>receive>entry>end-time

The following table details which authentication algorithm can be used in association with specific routing protocols.

Table 8 shows the mapping between these attributes and the CLI command to set them.

Table 8: Security Algorithm Support Per Protocol

Protocol	Clear Text	MD5	HMAC-MD5	HMAC-SHA-1-96	HMAC-SHA-1	HMAC-SHA-256	AES-128-CMAC-96
OSPF	Yes	Yes	No	Yes	Yes	Yes	No
IS-IS	Yes	No	Yes	No	Yes	Yes	No
RSVP	Yes	No	Yes	No	Yes	No	No
BGP	No	Yes	No	Yes	No	No	Yes
LDP	No	Yes	No	Yes	No	No	Yes

Configuration Notes

This section describes security configuration caveats.

General

- If a RADIUS or a TACACS+ server is not configured, then password, profiles, and user access information must be configured on each router in the domain.
- If a RADIUS authorization is enabled, then VSAs must be configured on the RADIUS server.