

Extreme Fabric Automation Security Guide

2.7.1

9037430-01 Rev AA
June 2022



Copyright © 2022 Extreme Networks, Inc. All rights reserved.

Legal Notice

Extreme Networks, Inc. reserves the right to make changes in specifications and other information contained in this document and its website without prior notice. The reader should in all cases consult representatives of Extreme Networks to determine whether any such changes have been made.

The hardware, firmware, software or any specifications described or referred to in this document are subject to change without notice.

Trademarks

Extreme Networks and the Extreme Networks logo are trademarks or registered trademarks of Extreme Networks, Inc. in the United States and/or other countries.

All other names (including any product names) mentioned in this document are the property of their respective owners and may be trademarks or registered trademarks of their respective companies/owners.

For additional information on Extreme Networks trademarks, see: www.extremenetworks.com/company/legal/trademarks

Open Source Declarations

Some software files have been licensed under certain open source or third-party licenses. End-user license agreements and open source declarations can be found at: <https://www.extremenetworks.com/support/policies/open-source-declaration/>



Table of Contents

Preface.....	4
Text Conventions.....	4
Documentation and Training.....	5
Help and Support.....	6
Subscribe to Product Announcements.....	6
Send Feedback.....	6
About this Document.....	8
What's New in this Document.....	8
Security Features in EFA.....	9
EFA User Authentication and Authorization.....	9
Authentication.....	9
Authorization.....	10
Security troubleshooting.....	11
EFA RBAC Policy Enforcement.....	11
Assign and View EFA Roles.....	15
Configure an External LDAP Server.....	15
EFA Certificates.....	16
BGP MD5 Authentication.....	18
BGP MD5 Authentication on Fabric Links.....	19
BGP MD5 Authentication on Edge Links.....	28
Security Hardening.....	39
Overview.....	39
System Hardening for CIS-CAT Assessments.....	40
The iptables Policy.....	43
Secure the Grub Boot Loader.....	45
Install the Linux Audit System.....	46
Install and Use OSSEC.....	47
Secure NTP and EFA.....	49
Secure DNS and EFA.....	50
Detect Rootkits with rkhunter.....	51



Preface

Read the following topics to learn about:

- The meanings of text formats used in this document.
- Where you can find additional information and help.
- How to reach us with questions and comments.

Text Conventions

Unless otherwise noted, information in this document applies to all supported environments for the products in question. Exceptions, like command keywords associated with a specific software version, are identified in the text.

When a feature, function, or operation pertains to a specific hardware product, the product name is used. When features, functions, and operations are the same across an entire product family, such as ExtremeSwitching switches or SLX routers, the product is referred to as *the switch* or *the router*.

Table 1: Notes and warnings






Icon	Notice type	Alerts you to...
	Tip	Helpful tips and notices for using the product
	Note	Useful information or instructions
	Important	Important features or instructions
	Caution	Risk of personal injury, system damage, or loss of data
	Warning	Risk of severe personal injury

Table 2: Text

Convention	Description
screen displays	This typeface indicates command syntax, or represents information as it is displayed on the screen.
The words <i>enter</i> and <i>type</i>	When you see the word <i>enter</i> in this guide, you must type something, and then press the Return or Enter key. Do not press the Return or Enter key when an instruction simply says <i>type</i> .
Key names	Key names are written in boldface, for example Ctrl or Esc . If you must press two or more keys simultaneously, the key names are linked with a plus sign (+). Example: Press Ctrl+Alt+Del
Words in italicized type	Italics emphasize a point or denote new terms at the place where they are defined in the text. Italics are also used when referring to publication titles.
NEW!	New information. In a PDF, this is searchable text.

Table 3: Command syntax

Convention	Description
bold text	Bold text indicates command names, keywords, and command options.
<i>italic</i> text	Italic text indicates variable content.
[]	Syntax components displayed within square brackets are optional. Default responses to system prompts are enclosed in square brackets.
{ x y z }	A choice of required parameters is enclosed in curly brackets separated by vertical bars. You must select one of the options.
x y	A vertical bar separates mutually exclusive elements.
< >	Nonprinting characters, such as passwords, are enclosed in angle brackets.
...	Repeat the previous element, for example, <i>member</i> [<i>member</i> ...].
\	In command examples, the backslash indicates a “soft” line break. When a backslash separates two lines of a command input, enter the entire command at the prompt without the backslash.

Documentation and Training

Find Extreme Networks product information at the following locations:

[Current Product Documentation](#)

[Release Notes](#)

[Hardware and software compatibility](#) for Extreme Networks products

[Extreme Optics Compatibility](#)

[Other resources](#) such as white papers, data sheets, and case studies

Extreme Networks offers product training courses, both online and in person, as well as specialized certifications. For details, visit www.extremenetworks.com/education/.

Help and Support

If you require assistance, contact Extreme Networks using one of the following methods:

Extreme Portal

Search the GTAC (Global Technical Assistance Center) knowledge base; manage support cases and service contracts; download software; and obtain product licensing, training, and certifications.

The Hub

A forum for Extreme Networks customers to connect with one another, answer questions, and share ideas and feedback. This community is monitored by Extreme Networks employees, but is not intended to replace specific guidance from GTAC.

Call GTAC

For immediate support: (800) 998 2408 (toll-free in U.S. and Canada) or 1 (408) 579 2826. For the support phone number in your country, visit: www.extremenetworks.com/support/contact

Before contacting Extreme Networks for technical support, have the following information ready:

- Your Extreme Networks service contract number, or serial numbers for all involved Extreme Networks products
- A description of the failure
- A description of any actions already taken to resolve the problem
- A description of your network environment (such as layout, cable type, other relevant environmental information)
- Network load at the time of trouble (if known)
- The device history (for example, if you have returned the device before, or if this is a recurring problem)
- Any related RMA (Return Material Authorization) numbers

Subscribe to Product Announcements

You can subscribe to email notifications for product and software release announcements, Field Notices, and Vulnerability Notices.

1. Go to [The Hub](#).
2. In the list of categories, expand the **Product Announcements** list.
3. Select a product for which you would like to receive notifications.
4. Select **Subscribe**.
5. To select additional products, return to the **Product Announcements** list and repeat steps 3 and 4.

You can modify your product selections or unsubscribe at any time.

Send Feedback

The Information Development team at Extreme Networks has made every effort to ensure that this document is accurate, complete, and easy to use. We strive to improve our documentation to help you in your work, so we want to hear from you. We welcome all feedback, but we especially want to know about:

- Content errors, or confusing or conflicting information.

- Improvements that would help you find relevant information.
- Broken links or usability issues.

To send feedback, do either of the following:

- Access the feedback form at <https://www.extremenetworks.com/documentation-feedback/>.
- Email us at documentation@extremenetworks.com.

Provide the publication title, part number, and as much detail as possible, including the topic heading and page number if applicable, as well as your suggestions for improvement.



About this Document

[What's New in this Document on page 8](#)

This document describes security features in EFA.

What's New in this Document

The following table describes information added to this guide for the Extreme Fabric Automation 2.7.1 software release.

Table 4: Summary of changes

Description	Link
Added new topic that describes about the EFA certificates and external certificates.	EFA Certificates on page 16



Security Features in EFA

[EFA User Authentication and Authorization](#) on page 9

[BGP MD5 Authentication](#) on page 18

The following provides an overview of the security features in EFA.

This document provides security hardening guidance for Extreme Fabric Automation (EFA), with an emphasis on the installation and usage of open source security tools to achieve a hardened operational security stance. Some basic knowledge of security principles and operations of the Linux operating system and associated technologies is assumed.

Note: This document assumes that EFA has been installed in secure mode and is operational. For details on how to achieve this, refer to the [Extreme Fabric Automation Administration Guide, 2.7.0](#), which includes details on EFA security options and commands.

The following security features are describes in the sections below:

- **Authentication and authorization:** Explains how EFA users are validated with Unix authentication or LDAP and managed with Role-based Access Control (RBAC).
- **BGP MD5 authentication on edge links:** How to authenticate all the BGP peer and peer-group used for edge connectivity.
- **BGP MD5 authentication on fabric links:** How to use MD5 for BGP connections across all fabric links.

EFA User Authentication and Authorization

In addition, at installation time, starting with EFA 2.5.0, all EFA users of services such as MySQL and RabbitMQ are assigned random passwords that are stored in EFA configuration files. This satisfies the requirement to enforce the change of default passwords, and no two EFA installations share identical passwords.

For more information, see [EFA RBAC Policy Enforcement](#) on page 11 and [Assign and View EFA Roles](#) on page 15.

Authentication

EFA validates users and their credentials with the following mechanisms:

- Unix authentication (local and remote) on the host where EFA is installed. Host credentials are the default validation method if LDAP validation fails.
- External LDAP server. Users configured in LDAP use their LDAP credentials to log in to EFA.

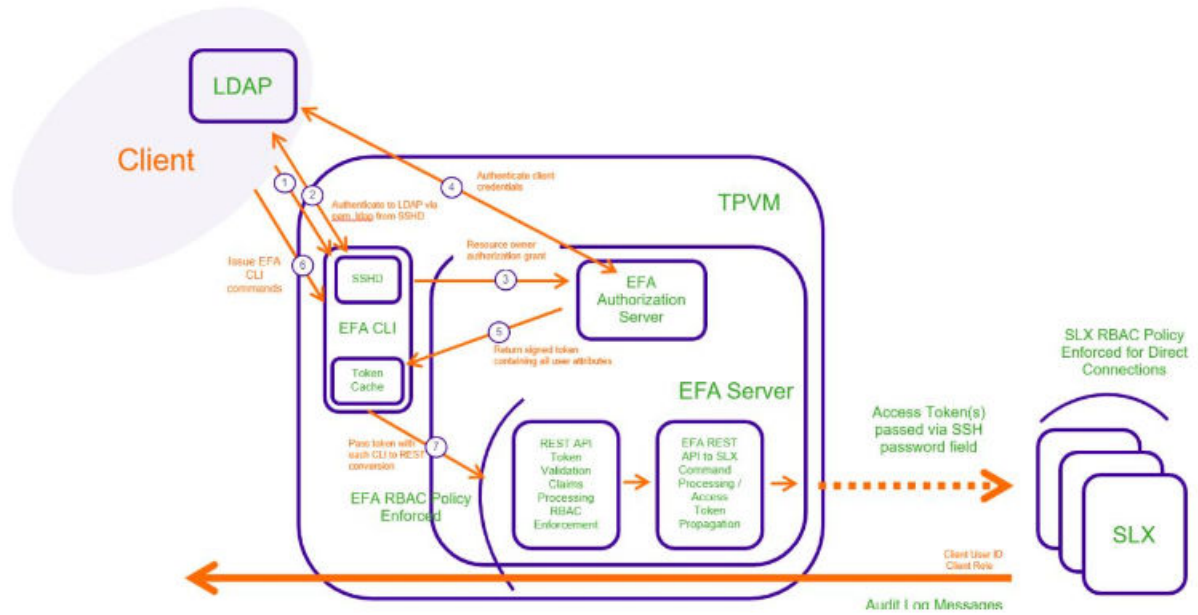


Figure 1: LDAP authentication example

Operational or maintenance tasks are propagated to SLX devices through OAuth2 and JWT access tokens. TLS is used for connections with SLX devices. The OpenStack ML2 plugin also uses TLS and OAuth2 tokens. When EFA is installed in secure mode, traffic to northbound interfaces uses TLS. For more information about secure mode, see the "EFA Installation Modes" topic in the [Extreme Fabric Automation Deployment Guide, 2.7.0](#).

Authorization

After EFA is deployed, the installing user has the role of SystemAdmin and has complete access to EFA functionality. For installation on TPVM, this user has the user name of 'extreme.' By default, no other host OS users can access EFA unless the SystemAdmin assigns the appropriate roles. RBAC occurs on EFA and API.

LDAP supports three modes for fetching the roles assigned to a user.

- The role is available as an attribute in the user Distinguished Name (DN) entry. Group attribute definition is not needed.
- The user has a "memberOf" attribute or any appropriate group DN attribute to identify the groups assigned to the user. Assign the corresponding LDAP group to a role in EFA.
- LDAP groups have user entries in their group definitions. Assign the LDAP groups to roles in EFA.

Security troubleshooting

Use the following logs to troubleshoot authentication, authorization, or RBAC issues.

Table 5: Security log locations

Log source	Filepath
EFA server	/var/log/efa/auth/auth-server.log /var/log/efa/rbac/rbac-server.log
EFA TPVM	/apps/efa_logs/auth/auth-server.log /apps/efa_logs/rbac/rbac-server.log
SLX device	/var/log/pam-oauth2.log

Use the following commands to see lists of commands that were run during a specified time, such as when an RBAC error occurred. This sort of information can help you identify potential causes.

- **efa auth execution show**
- **efa rbac execution show**
- **efa inventory execution show**

EFA RBAC Policy Enforcement

EFA implements an RBAC (Role-based Access Control) policy governing access to northbound REST APIs.

The RBAC policy is enforced at the northbound interface, immediately after validation of the access token. An error message is returned if an RBAC permissions check fails.

RBAC and REST URI Matrix

The RBAC policy is expressed in a permissions matrix indexed by RBAC role and REST URI, in which each matrix element enumerates the permitted HTTP methods.

Table 6: RBAC and REST Matrix

	Role A	Role B	Role C
REST URI 1	GET	GET	GET, POST, PUT, PATCH, DELETE
REST URI 2	GET, POST	GET, POST, PUT	GET, POST, PUT, PATCH, DELETE
REST URI 3	GET, POST	GET, POST	GET, POST, PUT, PATCH, DELETE

RBAC Roles

Roles can be populated into the upstream LDAP instance.

Table 7: Role definitions

Role	Description
FabricAdmin	<ul style="list-style-type: none"> Registers devices to the fabric Configures fabric parameters Validates all devices in the fabric Configures switches for IP fabric with overlay and without overlay Creates tenants Creates networks inside tenants, such as VRF, EPG, and PO Performs fabric debug activities Has privileges for OpenStack, Hyper-V, and vCenter operations
SecurityAdmin	Performs user management, PKI, and key management operations
NetworkOperator	<ul style="list-style-type: none"> Has view-only privileges for fabric configurations, information for tenants and inventory, and all ecosystem information Cannot make changes in the system
SystemDebugger	<ul style="list-style-type: none"> Has privileges to perform supportsave and system backup, and to view the running system configurations Has privileges to perform fabric debug operations Sets debug levels for services Has privileges to collect execution logs from services
SystemAdmin	Has complete privileges to all operations in the system
<Tenant>Admin * Created dynamically per tenant	Performs tenant administration within the assigned tenant, such as the following: <ul style="list-style-type: none"> Adding networks to the tenant Configuring network parameters Configuring switches with tenant-specific information Cannot perform actions for any other tenant

* Tenant Administrator roles are added dynamically to the system when a tenant is created. The name of the role is of the format <Tenant-name>Admin. For example, if a tenant with the name “RegionOne” is created, the role created for the Tenant Administrator is “RegionOneAdmin”.

**Note**

You cannot create custom roles.

*Role Permissions***Table 8: Role permissions**

Allowed Privileges	System Admin	Fabric Admin	Tenant Admin	Network Operator	Security Admin	System Debugger
Create, clone, delete fabric in the system	✓	✓				
Register, unregister devices in fabric, configure IP fabric on the device	✓	✓				
Show IP fabric physical, underlay, overlay topology, IP fabric configs and devices in IP fabric	✓	✓		✓		
Debug fabric operations	✓	✓				✓
Inventory, asset service operations	✓	✓				
Run CLI access on the device	✓	✓				
Create, delete, update tenants	✓	✓				
Create, delete EPG, PO, VRFs inside tenant	✓	✓	✓			
Add, remove port, port channels to and from EPG	✓	✓	✓			
Add, remove network policies to EPG	✓	✓	✓			
Detach network from EPG	✓	✓	✓			
Identify drift in device configuration	✓	✓				
Set tenant debug level	✓	✓	✓			✓
Show OpenStack networks, PO, subnets, tenant, ports, router, router-interface	✓	✓	✓	✓		
Create, delete, clean up OpenStack networks	✓	✓	✓			
Create, delete OpenStack subnets	✓	✓	✓			
Create, delete OpenStack ports	✓	✓	✓			

Table 8: Role permissions (continued)

Allowed Privileges	System Admin	Fabric Admin	Tenant Admin	Network Operator	Security Admin	System Debugger
Create, delete OpenStack router	✓	✓	✓			
Create, delete router interfaces	✓	✓	✓			
Delete OpenStack asset (DebugDeleteOSSAsset)	✓	✓	✓			✓
View vCenter details, events, ESXI details, physical links, virtual links, disconnected links, get server settings	✓	✓	✓	✓		
Register, delete, update vCenter	✓	✓	✓			
Set vCenter debug level	✓	✓	✓			✓
Update vCenter polling frequency, dead link clearing time	✓	✓	✓			
View SCVMM server details, service settings, physical links, virtual links	✓	✓	✓	✓		
Register, delete, update SCVMM server	✓	✓	✓			
Update SCVMM server polling frequency	✓	✓	✓			
User management, assign roles to users, configure LDAP, view available roles in the system	✓				✓	
Notification service (add, delete subscribers)	✓	✓				
Execution log view	✓	✓	✓ (only Tenant)	✓	✓ (only Auth and RBAC)	✓
Support save collection	✓	✓	✓	✓	✓	✓
Backup and restore operation	✓	✓				✓ (only backup)
Install certificates	✓	✓			✓	

Assign and View EFA Roles

You can assign a role to a user and to an LDAP group.

About This Task

For more information about EFA roles, see [EFA RBAC Policy Enforcement](#) on page 11.

Procedure

1. To assign a role to a user, run the following command.

```
# efa auth rolemapping add --name fabricuser --role FabricAdmin --type user
Successfully added the role mapping
```

In this example, a user named fabricuser was assigned the role of FabricAdmin.

2. To assign a role to an LDAP group, run the following command.

```
# efa auth rolemapping add --name "cn=viewer,dc=extr,dc=com" --role NetworkOperator
--type group
Successfully added the role mapping.
```

In this example, a group named "cn=viewer,dc=extr,dc=com" was assigned the role of NetworkOperator.

3. To view all role assignments, run the following command.

```
# efa auth rolemapping show
ID  Name      Role           Type
1   efauser   SystemAdmin    USER
2   fabricuser FabricAdmin     USER
3   viewer    NetworkOperator GROUP
```

4. To delete a role assignment, run the following command.

```
# efa auth rolemapping remove --id 3
Deleted role mapping successfully
```

In this example, the role for the user with ID 3 was removed.

Configure an External LDAP Server

You can configure an LDAP server for user validation and to fetch user groups.

About This Task

For more information about commands and supported parameters, see [Extreme Fabric Automation Command Reference, 2.7.0](#).

Procedure

1. To configure an external LDAP server, run the following command.

```
# efa auth ldapconfig add --name ldapconfig -- host 10.x.x.x --bind-user-
name cn=admin,dc=extrnet,dc=com --bind-user-password password --user-search-
base ou=people,dc=extrnet,dc=com
```

This example configures the bind user name and password and the DN of the node from which searches start.

2. To configure an LDAP server in a TPVM (for the TPVM Ubuntu OS), run the **tpvm config ldap** command from the SLX-OS command line.

EFA Certificates

The following tables provide information about EFA certificates and external certificates.

Certificates of EFA

Certificate	Location in TPVM deployment	Location in server deployment	Description	Default Validity Period	Impact on the system	Renewal Procedure	Alarm/ Notification
SSL/TLS Certificate of EFA	/apps/efadata/certs/own/tls.crt	/opt/efadata/certs/own/tls.crt	The certificate of EFA server for secure communication with the clients. The same certificate is used on port 443 (default EFA), 8078 (monitor service of EFA), 6514 (syslog listener on EFA)	Until EFA 2.7.0 - Expires in 1 year from installation. From EFA 2.7.1 - Expires in 3 years from installation. Reset after every subinterface creation/upgrade	If the certificate expires, then the server communication with SSL verification enabled will fail. Disables syslog messages from the devices	https://extremeportal.force.com/ExtrArticleDetail?an=000102994 https://extremeportal.force.com/ExtrArticleDetail?an=000104838	2.7.x: Notification is sent to EFA subscribers from 30 days to expiry and warning message on every login from 7 days to expiry.
Intermediate CA Certificate of EFA	/apps/efadata/certs/ca/extreme-ca-cert.pem	/opt/efadata/certs/ca/extreme-ca-cert.pem	The certificate of Certificate Authority, which is the issuer of client and server certificates of EFA and HTTPS certificate of SLX. Same certificate is seen as SyslogCA on SLX	Valid till Feb 17 2030 GMT		Not available	Not available

Certificate	Location in TPVM deployment	Location in server deployment	Description	Default Validity Period	Impact on the system	Renewal Procedure	Alarm/ Notification
Root CA Certificate of EFA	/apps/efadata/certs/ca/extreme-ca-root.pem	/opt/efadata/certs/ca/extreme-ca-root.pem	The certificate of Certificate Authority, which is the issuer of Intermediate CA certificate	Valid till Feb 15 2040 GMT		Not available	Not available
HTTPS Certificate of SLX	/apps/efadata/certs/slx-<IP>.extremenetworks.com-cert.pem	/opt/efadata/certs/slx-<IP>.extremenetworks.com-cert.pem	The certificate of SLX Web Server (Apache) for secure communication with the device from EFA	Expires in 2 years from installation	System will not use encryption for HTTPS requests	Not available	From 2.4.x: Notification is sent to EFA subscribers from 30 days of expiry.
K3s Certificate - EFA internal	/apps/rancher/k3s/server/tls/	/var/lib/rancher/k3s/server/tls/	EFA uses k3s for management of services. This certificate is for secure communication of k3s with clients	Expires in 1 year from installation. Reset after every upgrade of EFA		https://extremeportal.force.com/ExtrArticleDetail?an=000102994	Not available

Certificate	Location in TPVM deployment	Location in server deployment	Description	Default Validity Period	Impact on the system	Renewal Procedure	Alarm/ Notification
Host Authentication Service Certificate - EFA internal	/apps/bin/hostauth-certs/cert.pem	/usr/local/bin/hostauth-certs/cert.pem	The server certificate of host authentication service on EFA	Until EFA 2.7.0 - Expires in 3 years from the date of release. From EFA 2.7.1 - Expires in 10 years from the date of release	System will not use encryption for HTTPS requests	Refreshed on upgrade of EFA to next release	Not available
JWT Signing/ Verification - EFA internal	/apps/efadata/certs/cert.crt.pem	/opt/efadata/certs/cert.crt.pem	The RSA public key for JWT verification. This is also used to send user context from EFA to SLX. Same certificate is seen as Oauth certificate on SLX	Expires in 10 years from the date of installation	Disables login to EFA	Not available	Not available

External Certificates

Certificate	Location in TPVM deployment	Location in server deployment	Description
Syslog CA	Notification service database in mariadb	Notification service database in mariadb	Connect to external Syslog server for sending notifications over RELP.
Webhook CA	Notification service database in mariadb	Notification service database in mariadb	Connect to external Syslog server for sending notifications over webhooks.
LDAP CA	Auth service database in mariadb	Auth service database in mariadb	Connect to external LDAP server

BGP MD5 Authentication

The Border Gateway Protocol (BGP) is an exterior gateway protocol designed to exchange routing and reachability information among autonomous systems on the internet. BGP works well in practice, and its

simplicity and resilience enable it to play a fundamental role within the global internet. However, BGP inherently provides few performance or security protections. The following table provides a list of some of the threats against BGP.



Note

BGP depends on TCP as its transport protocol. Therefore, it is vulnerable to the same security attacks as any TCP-based protocol.

Threats against BGP	Description
Denial of Service (DoS)	A malicious host sends unexpected or unwanted BGP traffic to a neighbor in an attempt to saturate control plane resources, which results in not having enough resources to process legitimate BGP traffic on the neighbor.
Route Manipulation	A malicious host modifies the contents of a BGP routing table, diverting traffic, and preventing it, without the sender's knowledge, from reaching its intended destination.
Route Hijacking	A rogue BGP neighbor maliciously advertises a victim's networks to redirect some or all of victim's traffic to itself.
Misconfiguration (non-malicious)	An unintentionally misconfigured BGP router could affect the Internet's BGP routing table, possibly leading to network outages and, worse, unauthorized access to the network traffic.

BGP authentication enables the routers to share information only if they can verify that they are communicating to a trusted source, based on a password. Successful authentication between BGP neighbors proves that the neighbors are legitimate and trusted, verifies communications between those neighbors, and ensures that only routes learned from legitimate neighbors are added to the routing table. Authentication must be enabled on both sides of the peering session and the same password must be used on both peers.



Note

For BGP MD5 passwords, the ASCII characters 0-32 are not supported. In addition, special handling is required for MD5 passwords that contain certain special characters.

Examples

MD5 password provided through CLI	Actual MD5 password
'~`!@#\$%^&*()_+={[] \\"<>"/''''	~`!@#\$%^&*()_+={[] \\"<>/'
'a'''a"	a"a
'a''''a"	a''''a"
'a''''''a'''	a''''a'

BGP MD5 Authentication on Fabric Links

EFA provides secure TCP using MD5 for BGP connections across all fabric links. You are able to configure or modify the MD5 password at any time.

This feature enables you to provide an `md5-password` as a fabric setting that further becomes configured on all the fabric links, that is, BGP peer-groups and individual neighbors without peer-groups on the SLX devices, so that the peer sessions are established using MD5 authentication. This will be applied to both Clos and non-Clos fabrics. By default, MD5 authentication on the fabric is disabled. Any new fabric or the fabrics upgraded from previous releases will have the MD5 authentication disabled.

Configure BGP MD5 Password Create, Update, and Clear

Use the `efa fabric setting update` command to set or clear the MD5 password on a new fabric.

About This Task

Here is the `efa fabric setting update` command:

```
efa fabric setting update --name <fabric-name> --md5-password-enable <yes/no> --md5-
password <password>
```

If the command is entered with `md5-password-enable` as “yes” but without the `md5-password` option, then a prompt is displayed to input string and the password string entered using the prompt is not displayed on the screen.



Note

When providing a password string in the command line, that is using `efa fabric setting update --name <fabric-name> --md5-password <password>`, if the string contains special characters, then you must enclose the string in single quotes. For example `efa fabric setting update --name fabric1 --md5-password 'pass%!'`. Enclosing the password string in single quotes is not required when the string is entered using the prompt.

After setting the `md5-password`, you must configure the fabric, using the command `efa fabric configure --name <fabric-name>`, to apply this MD5 password on fabric devices so that the BGP neighbor sessions are authenticated.

When you configure the `md5-password` on a fabric that has just been created, or a fabric that has not yet been configured, there is no change in the device `app-state`. However, if the `md5-password` is set after the fabric is configured, there is a new `app-state`, the fabric setting is refreshed (and devices will be set to), indicating the fabric properties have been modified and the fabric has to be reconfigured to apply the new settings. As part of fabric configuration, when the MD5 password was successfully configured on all the fabric links on a device, the `app state` on that device will go back to `cfg-in-sync` state.

Procedure

1. Run the `efa fabric setting update --name <fabric-name> --md5-password-enable <yes/no> --md5-password <password>` command to set the MD5 password.

```
efa fabric setting update --name fabric1 --md5-password-enable yes
Please supply a password for BGP MD5 authentication on fabric links:
```

- Run the **efa fabric configure --name <fabric-name>** command to apply this MD5 password on fabric devices so that the BGP neighbor sessions are authenticated.

To create or update MD5 authentication:

```
efa fabric configure --name fabric1
```

To clear MD5 authentication:

```
efa fabric setting update --name fabric1 --md5-password-enable no
efa fabric configure --name fabric1
```

When you configure the `md5-password` on a fabric that has just been created, or a fabric that has not yet been configured, there is no change in the device `app-state`. However, if the `md5-password` is set after the fabric is configured, the fabric status is set to `settings-updated` along with the field `BGP-MD5`, indicating that settings have been updated. This indicates that the fabric properties have been modified and the fabric has to be reconfigured to apply the new settings. As part of `fabric configure`, when the devices are successfully configured, the fabric status will go back to `configure-success`.

Example

```
Fabric Name: fabric1, Fabric Description: , Fabric Stage: 3, Fabric Type: clos, Fabric
Status: settings-updated
Updated Fabric Settings: BGP-MD5
+-----+-----+-----+-----+-----+-----+
+-----+-----+-----+-----+-----+-----+
| IP ADDRESS | POD | HOST NAME | ASN | ROLE | DEVICE STATE | APP STATE |
CONFIG GEN REASON | PENDING CONFIGS | VTLB ID | LB ID |
+-----+-----+-----+-----+-----+-----+
+-----+-----+-----+-----+-----+-----+
| 10.17.112.223 | | spine1 | 64512 | spine | provisioned | cfg in-sync |
MD5 | MD5-U | | NA | 1 | |
| 10.17.112.224 | | spine2 | 64512 | spine | provisioned | cfg in-sync |
MD5 | MD5-U | | NA | 1 | |
| 10.17.112.221 | | leaf1 | 65002 | leaf | provisioned | cfg in-sync |
MD5 | MD5-U | 2 | 1 | |
| 10.17.112.222 | | leaf2 | 65002 | leaf | provisioned | cfg in-sync |
MD5 | MD5-U | 2 | 1 | |
| 10.17.112.225 | | leaf3 | 65000 | leaf | provisioned | cfg in-sync |
MD5 | MD5-U | 2 | 1 | |
| 10.17.112.226 | | leaf4 | 65000 | leaf | provisioned | cfg in-sync |
MD5 | MD5-U | 2 | 1 | |
+-----+-----+-----+-----+-----+-----+
+-----+-----+-----+-----+-----+-----+
FABRIC SETTING:
BGPLL - BGP Dynamic Peer Listen Limit, BGP-MD5 - BGP MD5 Password
CONFIG GEN REASON:
LD - Link Delete, LA - Link Add, IU - Interface Update, PLC - IPPrefixList Create, PLD
- IPPrefixList Delete, PLU - IPPrefixList Update
MD/MU - MCT Delete/Update, OD - Overlay Gateway Delete, OU - Overlay Gateway Update,
ED - Evpn Delete, PC - RouterPim Create, PD - RouterPim Delete, BGP - BGP Config
DD - Dependent Device Update, DA - Device Add, DR - Device ReAdd, ASN - Asn Update, PU
- RouterPim Update, SYS - System Properties Update, NA - Not Applicable
PENDING CONFIGS:
MCT - MCT Cluster, O - Overlay Gateway, SYSP - System Properties, INTIP - Interface
IP, BGP - BGP Config
C/D/U - Create/Delete/Update, PA/PD - Port Add/Port Delete
```

For App or Device Error/Failure reason, run "efa fabric error show" for details
 For config refresh reason, run "efa fabric debug config-gen-reason" for details



Note

When the MD5 password is updated, for the new configuration to take effect, the neighbor sessions have to be cleared, resulting in a network outage until the new sessions are established. Since the configuration of the MD5 password toggles the network, a new warning message with a confirmation is provided indicating the impact of the `md5-password` setting on an active fabric, before it is applied. This warning message is displayed only when there is a need to reconfigure the fabric, that is, the password is set after the fabric is configured.

```
efa fabric setting modify --name fabric1 --md5-password-enable yes
```

```
Please supply a password for BGP MD5 authentication on fabric links:
```

```
WARNING: configuring/clearing md5-password on an active fabric will result in  
BGP neighborsessions  
going down for a brief period when the fabric is reconfigured.
```

```
Please confirm if you want to continue with the fabric setting update [y/n]?
```

Employ a Phased Approach

Reduce network impact by applying the password and clearing sessions on fabric links in a phased manner.

When an updated MD5 password is being applied on fabric links during fabric configuration, you can reduce network impact by applying the password and clearing sessions on fabric links in a phased manner. First gather a list of neighbor sessions on each device. Then, one device at a time, apply `md5-password` and clear a single peering. Application of password and clearing of the neighbor session is done on both ends of the peering session simultaneously. When the session is established, the MD5 password is applied on the next peering session. When all the neighbor sessions on the device are updated, it will move to the next device.

After you clear the session, it takes 10 seconds for the new session to be established. Before applying the new MD5 password, the session states are determined. After applying the password and clearing the session on both ends of the peering session, the session state is checked again. Only when the state matches with the previous (pre-password update) state on the session, or is better than the previous state (for example, previously the session was not established and the current state is established), it will move to updating the next neighbor session.



Note

The phased application of clearing the session and checking the state is performed **only** when the password is updated on a configured fabric and is not applicable during the configuration of a new fabric. Because the neighbor sessions are created for the first time during configuration of a new fabric with the MD5 password, there is no need to clear sessions.

Configure BGP MD5 Password: Failures When Clearing the Neighbor Session

Failures can occur during the process of clearing the neighbor sessions when the MD5 password is updated.

After clearing the session, if the session state is not established, or is not in the same state as it was prior to clearing the session within the wait time of 10 seconds, then the wait time is extended for an additional 10 seconds. If the session state is not established after the expiration of the second wait time, it is marked as a failure. The execution continues with the clearing of the remaining neighbor sessions.

When all the neighbor sessions are cleared, any sessions that have been marked as failure are presented under fabric errors at the end of the fabric configure operation, as part of the existing fabric error command `efa fabric error show`. The failure information will include the details of the neighbor session that could not be established.



Note

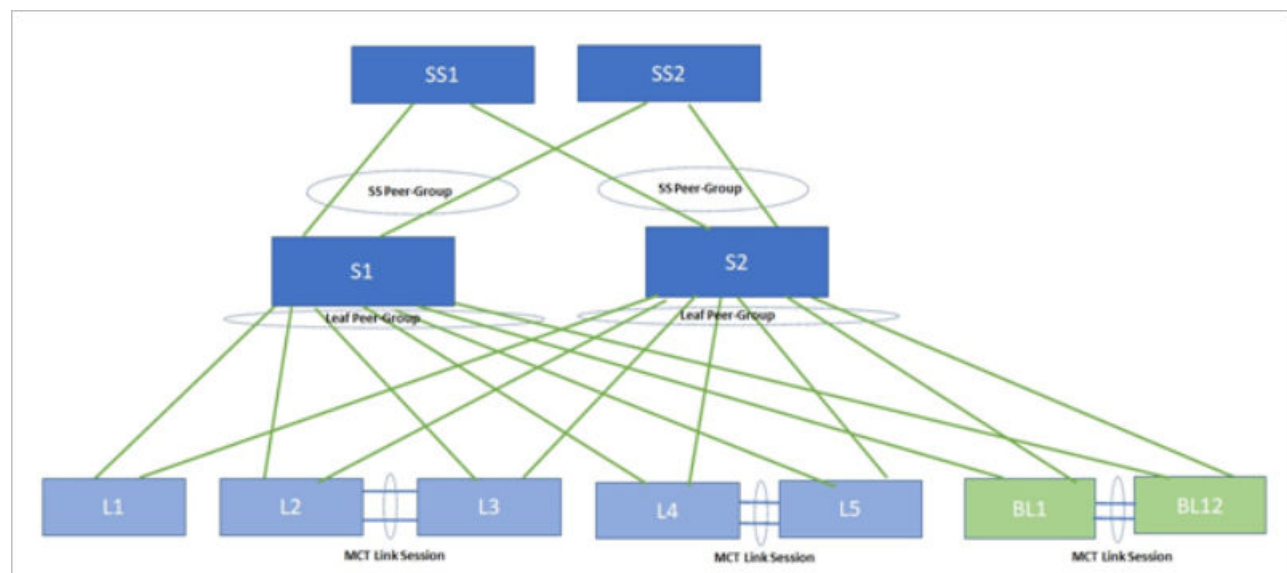
If there are any sessions that could not be established during the clear operation, the fabric configuration operation displays an error indicating a failure. However, if the fabric configuration has been successfully pushed to the devices, the devices are set to `cfg-in-sync`, even though the clear operation failed.

Configure BGP MD5 Password: Clos Topology (3-Stage and 5-Stage)

In a Clos topology, session clearing is done on each device on the fabric.

In a Clos topology, the clearing of the sessions is done by walking through each device on the fabric and the sessions cleared, depending on the role of the device. The clearing is done in the following manner:

- **Spine:** Neighbor sessions at the peer-group level are cleared. On the spine, there are two peer-groups, one for the links to the leaves and the other for the links to the super-spines (in the case of 5-stage Clos). Clearing of the sessions is done on both peer-groups simultaneously.
- **Leaf:** The neighbor session on the MCT link is cleared.
- **Super-spine:** There is no need to clear any session on the super-spine because all the sessions on the super-spines are covered at the spine device.



Configure BGP MD5 Password: Non-Clos Topology and Fabric Events

In a non-Clos topology, the process of clearing the session is performed on one leaf device at a time, until all sessions are cleared.

Since the MD5 password on fabric links is applied as part of fabric configure, there are no additional events that are generated as a result of configuring MD5 password. The Fabric Deployed event that is currently generated by Fabric service when a fabric is configured should be used as a trigger by other services that are dependent on the Fabric MD5 password for any of their operations.

Fabric Events and the MD5 Password

Because the MD5 password on fabric links are applied as part of fabric configuration, no additional events are generated as a result of configuring the MD5 password.

The Fabric Deployed event that is generated by the fabric service when a fabric is configured should be used as a trigger by other services that are dependent on the fabric MD5 password for any of their operations.

Verify the BGP MD5 Password

Use the **efa fabric setting show** command to verify that the MD5 password is configured on the fabric.

About This Task

The password is not displayed in clear text because of security concerns. The password is displayed as a hidden string (*****) or as an encrypted string. The password is displayed as a hidden string if the fabric is not configured after the MD5 password is set. And if the fabric is configured after the password is set, then the show command displays the password as an encrypted string. This encrypted string matches the password string displayed on the SLX devices.

The same applies to the **efa show-running-config** command. Until the fabric is configured, the MD5 password is displayed as *****. After the fabric is configured, the encrypted string is displayed.

Procedure

Run the **efa fabric setting show --name <fabric-name> --advanced** command.

```

efa fabric setting show --name fabric1 --advanced
+-----+-----+
| NAME | VALUE |
+-----+-----+
| Fabric Name | fabric1 |
+-----+-----+
| Link IP Range | 10.10.10.0/23 |
+-----+-----+
| Loopback IP Range | 172.31.254.0/24 |
+-----+-----+
| Loopback Port Number | 1 |
+-----+-----+
| VTEP Loopback Port Number | 2 |
+-----+-----+
| Spine ASN Block | 64512-64768 |
+-----+-----+
| SuperSpine ASN Block | 64769 |
+-----+-----+
| Leaf ASN Block | 65000-65534 |

```



```
+-----+
| Border Leaf ASN Block | 66000-66100 |
+-----+
| P2P IP Type | numbered |
+-----+
| Any cast MAC | 0201.0101.0101 |
+-----+
| IPV6 Any cast MAC | 0201.0101.0102 |
+-----+
| MAC Aging Timeout | 1800 |
+-----+
| MAC Aging Conversational | 300 |
| Timeout | |
+-----+
| MAC Move Limit | 20 |
+-----+
| Duplicate MAC Timer | 5 |
+-----+
| Duplicate MAC Timer MAX Count | 3 |
+-----+
| BFD Enable | Yes |
+-----+
| BFD Tx | 300 |
+-----+
| BFD Rx | 300 |
+-----+
| BFD Multiplier | 3 |
+-----+
| BGP MultiHop | 2 |
+-----+
| MaxPaths | 8 |
+-----+
| AllowAsIn | 0 |
+-----+
| MTU | 9216 |
+-----+
| IPMTU | 9100 |
+-----+
| MCT Link IP Range | 10.20.20.0/24 |
+-----+
| MCT PortChannel | 64 |
+-----+
| LACP Timeout | long |
+-----+
| Control Vlan | 4090 |
+-----+
| Control VE | 4090 |
+-----+
| Leaf PeerGroup | spine-group |
+-----+
| Spine PeerGroup | leaf-group |
+-----+
| SuperSpine PeerGroup | spine-group |
+-----+
| Configure Overlay Gateway | Yes |
+-----+
| VNI Auto Map | Yes |
+-----+
| Backup Routing Enable | No |
+-----+
| Backup Routing IPv4 Range | 10.40.40.0/24 |
+-----+
| Backup Routing IPv6 Range | fd40:4040:4040:1::/120 |
+-----+
```

```

| Optimized Replication Enable | No |
+-----+-----+
| MDT Group IPv4 Range | 239.0.0.0/8 |
+-----+-----+
| Default MDT Group IPv4 address | 239.1.1.1 |
+-----+-----+
| MD5 Password Enable | Yes |
+-----+-----+
| MD5 Password | $9$GiXG/W7938rCj4lzgf14NQ== |
+-----+-----+

```

Configure BGP MD5 Password: Switch Configuration

The following provides a sample EFA MD5 password configuration and the corresponding switch configuration on one of the fabric devices after the fabric is configured.

About This Task

Procedure

1. Run the **efa fabric setting update** command.

```

efa fabric setting update --name fabric1 --md5-password-enable yes
Please supply a password for BGP MD5 authentication on fabric links:
efa fabric configure --name fabric1

```

2. Complete the following configuration on SLX:

```

router bgp
  local-as 65000
  capability as4-enable
  fast-external-fallover
  neighbor spine-group peer-group
  neighbor spine-group remote-as 64512
  neighbor spine-group description To Spine
  neighbor spine-group password $9$GiXG/W7938rCj4lzgf14NQ==
  neighbor 10.10.10.2 peer-group spine-group
  neighbor 10.20.20.2 remote-as 65000
  neighbor 10.20.20.2 next-hop-self
  neighbor 10.20.20.2 password $9$GiXG/W7938rCj4lzgf14NQ==
  address-family ipv4 unicast
  network 172.31.254.2/32
  maximum-paths 8
  graceful-restart
  !
  address-family ipv6 unicast
  !
  address-family l2vpn evpn
  graceful-restart
  neighbor spine-group encapsulation vxlan
  neighbor spine-group next-hop-unchanged
  neighbor spine-group enable-peer-as-check
  neighbor spine-group activate
  !
  !

```

The BGP MD5 Password, Drift and Reconcile, and Idempotency

Drift is identified if you modify the MD5 password through SLX, the CLI, or other management tool.

A reconcile operation pushes the intended configuration to SLX, thereby synchronizing the SLX configuration with EFA.



Note

A reconcile operation configures the MD5 password on the device back to its original value (pre-drift) but does not clear the session. Also, the state is not verified after the password is configured.

Field	Identity Drift	Reconcile Configuration	Idempotency
md5-password	Yes	Yes	Yes*

* There are some caveats to idempotency. The fabric service does not store the plain text password you provide after the fabric has been configured. It stores the encrypted string of the user-provided password, matching with the encrypted string available on the SLX device. So, setting the same original MD5 password after the fabric is configured results in devices going into `cfg-refreshed` state. For the operation to be idempotent, after the fabric is configured, the encrypted string should be provided as the `md5-password` and not the original plain text password.

Devices in the fabric are in `cfg-refreshed` state when the MD5 password has been updated but the fabric is not yet reconfigured. In such a scenario, the previous MD5 password is used for drift detection until the fabric is configured with the new password.

Configure BGP MD5 Password: Fabric Clone

When a fabric is cloned using the command `efa fabric clone --source <old-fabric-name> -- destination <new-fabric-name>`, the MD5 password configuration from the source fabric is used in the new cloned fabric.

Configure BGP MD5 Password: Rules for Clearing BGP Sessions

MD5 password configuration and clearing BGP neighbor sessions.

The following table lists different scenarios of MD5 password configuration, under which a BGP neighbor session is required to be cleared.

Scenarios	Clearing BGP neighbor session required?
Creating new BGP peer-group with MD5 password	Not required
Creating new BGP neighbor with MD5 password	Not required
Updating existing BGP peer-group with MD5 password	Yes
Updating existing BGP neighbor with MD5 password	Yes
Removing MD5 password from a BGP peer-group	Yes
Removing MD5 password from a BGP neighbor	Yes

Scenarios	Clearing BGP neighbor session required?
Modifying MD5 password for a BGP peer-group	Yes
Modifying MD5 password for a BGP neighbor	Yes
Reload	Not required
Copy <ftp://backup-config> startup-config and reload	Not required
Copy <ftp://backup-config> running-config	Not required

SLX Commands to clear the BGP neighbor sessions

```
clear ip bgp neighbor <neighbor ip> vrf <vrf name>
```

```
clear ip bgp neighbor <peer-group> vrf <vrf name>
```

The `vrf` used on the EFA fabric links is `default-vrf`.

BGP MD5 Authentication on Edge Links

This feature authenticates all the BGP peer and peer-group used for edge connectivity. You can provide an MD5 password per BGP peer and peer-group created for external connectivity.



Important

BGP MD5 authentication for tenant dynamic peers is not yet supported.

Configure BGP MD5 Authentication for Tenant BGP Peer

About This Task

Provide `md5-password` during BGP peer create or update operations.

Procedure

1. Run the **`efa tenant service bgp peer create`** command to create the peer.

```
efa tenant service bgp peer create
  --name <bgp-peer-name> --tenant <tenant-name>
  --ipv4-uc-nbr <device-ip,vrf-name:neighbor-ip,remote-asn>
  --ipv4-uc-nbr-bfd <device-ip,vrf-name:neighbor-ip,true|false>
  --ipv4-uc-nbr-md5-password <device-ip,vrf-name:neighborip,
  ipv4-md5-password>
```

2. Run the **`efa tenant service bgp peer update`** command to update the peer.

```
efa tenant service bgp peer update
  --name <bgp-peer-name> --tenant <tenant-name>
  --operation peer-add
  --ipv4-uc-nbr <device-ip,vrf-name:neighbor-ip,remote-asn>
  --ipv4-uc-nbr-bfd <device-ip,vrf-name:neighbor-ip,true|false>
  --ipv4-uc-nbr-md5-password <device-ip,vrf-name:neighborip,
  ipv4-md5-password>
```

Example

```
efa tenant service bgp peer create
--name tenlbgppeer1 --tenant ten1
--ipv4-uc-nbr 10.20.246.15,ten1vrf1:10.20.30.40,50000
--ipv4-uc-nbr-bfd 10.20.246.15,ten1vrf1:10.20.30.40,true
--ipv4-uc-nbr-md5-password 10.20.246.15,ten1vrf1:10.20.30.40,password
--ipv4-uc-nbr 10.20.246.16,ten1vrf1:10.20.30.40,50000
--ipv4-uc-nbr-bfd 10.20.246.16,ten1vrf1:10.20.30.40,true
--ipv4-uc-nbr-md5-password 10.20.246.16,ten1vrf1:10.20.30.40,password
efa tenant service bgp peer update
--name tenlbgppeer1 --tenant ten1
--operation peer-add
--ipv4-uc-nbr 10.20.246.15,ten1vrf1:10.20.30.50,50000
--ipv4-uc-nbr-bfd 10.20.246.15,ten1vrf1:10.20.30.50,true
--ipv4-uc-nbr-md5-password 10.20.246.15,ten1vrf1:10.20.30.50,password1
--ipv4-uc-nbr 10.20.246.16,ten1vrf1:10.20.30.50,50000
```

```
--ipv4-uc-nbr-bfd 10.20.246.16,ten1vrfl:10.20.30.50,true
--ipv4-uc-nbr-md5-password 10.20.246.16,ten1vrfl:10.20.30.50,password1
```

<pre>efa tenant service bgp peer show -- detail ===== Name : ten1bgppeer1 Tenant : ten1 State : bs-state-created Description : Static Peer ----- Device IP : 10.20.246.15 VRF : ten1vrfl AFI : ipv4 SAFI : unicast Remote IP : 10.20.30.40 Remote ASN : 50000 Next Hop Self : false Update Source IP : BFD Enabled : true BFD Interval : 0 BFD Rx : 0 BFD Multiplier : 0 MD5 Password : \$9\$MCgKGaNT6OASX68/7TC6Lw== Dev State : provisioned App State : cfg-in- sync Device IP : 10.20.246.15 VRF : ten1vrfl AFI : ipv4 SAFI : unicast Remote IP : 10.20.30.50 Remote ASN : 50000 Next Hop Self : false Update Source IP : BFD Enabled : true BFD Interval : 0 BFD Rx : 0 BFD Multiplier : 0 MD5 Password : \$9\$ufD04Gw +49ex4H8UtvifqA== Dev State : provisioned App State : cfg-in- sync</pre>	<pre>Device IP : 10.20.246.16 VRF : ten1vrfl AFI : ipv4 SAFI : unicast Remote IP : 10.20.30.40 Remote ASN : 50000 Next Hop Self : false Update Source IP : BFD Enabled : true BFD Interval : 0 BFD Rx : 0 BFD Multiplier : 0 MD5 Password : \$9\$MCgKGaNT6OASX68/7TC6Lw== Dev State : provisioned App State : cfg-in- sync Device IP : 10.20.246.16 VRF : ten1vrfl AFI : ipv4 SAFI : unicast Remote IP : 10.20.30.50 Remote ASN : 50000 Next Hop Self : false Update Source IP : BFD Enabled : true BFD Interval : 0 BFD Rx : 0 BFD Multiplier : 0 MD5 Password : \$9\$ufD04Gw +49ex4H8UtvifqA== Dev State : provisioned App State : cfg-in- sync Dynamic Peer ----- 0 Records 0 Records ===== =====</pre>
--	--

3. Complete the configuration on SLX as provided in the following example.

<pre>L1# show running-config router bgp router bgp local-as 4200000000 capability as4-enable fast-external-fallover neighbor 10.20.20.4 remote-as 4200000000 neighbor 10.20.20.4 next-hop-self address-family ipv4 unicast network 172.31.254.46/32 network 172.31.254.123/32 maximum-paths 8 graceful-restart ! address-family ipv4 unicast vrf tenlvrf1 redistribute connected neighbor 10.20.30.40 remote-as 50000 neighbor 10.20.30.40 password \$9\$MCgKGaT6OASX68/7TC6Lw== neighbor 10.20.30.40 bfd neighbor 10.20.30.50 remote-as 50000 neighbor 10.20.30.50 password \$9\$ufD04Gw+49ex4H8UtvifqA== neighbor 10.20.30.50 bfd maximum-paths 8 ! address-family ipv6 unicast ! address-family ipv6 unicast vrf tenlvrf1 redistribute connected maximum-paths 8 ! address-family l2vpn evpn graceful-restart ! !</pre>	<pre>L2# show running-config router bgp router bgp local-as 4200000000 capability as4-enable fast-external-fallover neighbor 10.20.20.5 remote-as 4200000000 neighbor 10.20.20.5 next-hop-self address-family ipv4 unicast network 172.31.254.46/32 network 172.31.254.176/32 maximum-paths 8 graceful-restart ! address-family ipv4 unicast vrf tenlvrf1 redistribute connected neighbor 10.20.30.40 remote-as 50000 neighbor 10.20.30.40 password \$9\$MCgKGaT6OASX68/7TC6Lw== neighbor 10.20.30.40 bfd neighbor 10.20.30.50 remote-as 50000 neighbor 10.20.30.50 password \$9\$ufD04Gw+49ex4H8UtvifqA== neighbor 10.20.30.50 bfd maximum-paths 8 ! address-family ipv6 unicast ! address-family ipv6 unicast vrf tenlvrf1</pre>
--	--



Note

The MD5 password cannot be set or unset on an existing BGP peer present within a peer instance. You need to remove the BGP peer from the BGP peer instance and then add back the BGP peer to the peer instance with the desired MD5 password configuration.

Configure BGP MD5 Authentication for Tenant BGP Peer-group

You can provide an MD5 password during BGP peer-group create or update operations.

Procedure

1. Use the **efa tenant service bgp peer-group create** command to create the peer group.

```
efa tenant service bgp peer-group create
--name <bgp-pg-name> --tenant <tenant-name>
--pg-name <device-ip:pg-name> --pg-asn <device-ip,pg-name:remote-asn>
--pg-bfd-enable <device-ip,pg-name:true|false>
--pg-md5-password <device-ip,pg-name:md5-password>
```

2. Use the **efa tenant service bgp peer-group update** command to update the peer group.

```
efa tenant service bgp peer-group update
--name <bgp-pg-name> --tenant <tenant-name>
--operation peer-group-add
--pg-name <device-ip:pg-name> --pg-asn <device-ip,pg-name:remote-asn>
--pg-bfd-enable <device-ip,pg-name:true|false>
--pg-md5-password <device-ip,pg-name:md5-password>
```

Example

```
efa tenant service bgp peer-group create
--name ten1bgppg1 --tenant ten1
--pg-name 10.20.246.15:pg1 --pg-asn 10.20.246.15,pg1:55001
--pg-bfd-enable 10.20.246.15,pg1:true
--pg-md5-password 10.20.246.15,pg1:password
--pg-name 10.20.246.16:pg1 --pg-asn 10.20.246.16,pg1:55001
--pg-bfd-enable 10.20.246.16,pg1:true
--pg-md5-password 10.20.246.16,pg1:password

efa tenant service bgp peer-group update
--name ten1bgppg1 --tenant ten1
--operation peer-group-add
--pg-name 10.20.246.15:pg2 --pg-asn 10.20.246.15,pg2:55002
--pg-bfd-enable 10.20.246.15,pg2:true
--pg-md5-password 10.20.246.15,pg2:password1
--pg-name 10.20.246.16:pg2 --pg-asn 10.20.246.16,pg2:55002
--pg-bfd-enable 10.20.246.16,pg2:true
--pg-md5-password 10.20.246.16,pg2:password1

efa tenant service bgp peer-group show --detail
=====
=====
Name           : ten1bgppg1
Tenant         : ten1
State          : bgp-pg-state-created

Peer Group
-----
Device IP      : 10.20.246.15
Peer Group    : pg1
Remote ASN     : 55001
Next Hop Self  : false
BFD Enabled   : true
BFD Interval  :
BFD Rx        :
BFD Multiplier :
MD5 Password  : $9$MCgKGaNT6OASX68/7TC6Lw==
Dev State     : provisioned
App State     : cfg-in-sync
```



```
Device IP      : 10.20.246.15
Peer Group    : pg2
Remote ASN    : 55002
Next Hop Self : false
BFD Enabled   : true
BFD Interval  :
BFD Rx       :
BFD Multiplier :
MD5 Password : $9$ufD04Gw+49ex4H8UtvifqA==
Dev State     : provisioned
App State     : cfg-in-sync

Device IP      : 10.20.246.16
Peer Group    : pg1
Remote ASN    : 55001
Next Hop Self : false
BFD Enabled   : true
BFD Interval  :
BFD Rx       :
BFD Multiplier :
MD5 Password : $9$MCgKGaNT6OASX68/7TC6Lw==
Dev State     : provisioned
App State     : cfg-in-sync

Device IP      : 10.20.246.16
Peer Group    : pg2
Remote ASN    : 55002
Next Hop Self : false
BFD Enabled   : true
BFD Interval  :
BFD Rx       :
BFD Multiplier :
MD5 Password : $9$ufD04Gw+49ex4H8UtvifqA==
Dev State     : provisioned
App State     : cfg-in-sync
```

```
=====
=====
```

3. Complete the following configuration on SLX.

<pre>L1# show running-config router bgp router bgp local-as 4200000000 capability as4-enable fast-external-fallover neighbor pg1 peer-group neighbor pg1 remote-as 55001 neighbor pg1 password \$9\$MCgKGaT6OASX68/7TC6Lw== neighbor pg1 bfd neighbor pg2 peer-group neighbor pg2 remote-as 55002 neighbor pg2 password \$9\$ufD04Gw +49ex4H8UtvifqA== neighbor pg2 bfd neighbor 10.20.20.4 remote-as 4200000000 neighbor 10.20.20.4 next-hop-self address-family ipv4 unicast network 172.31.254.46/32 network 172.31.254.123/32 maximum-paths 8 graceful-restart ! address-family ipv4 unicast vrf tenlvrf1 redistribute connected maximum-paths 8 ! address-family ipv6 unicast ! address-family ipv6 unicast vrf tenlvrf1 redistribute connected maximum-paths 8 ! address-family l2vpn evpn graceful-restart ! !</pre>	<pre>L2# show running-config router bgp router bgp local-as 4200000000 capability as4-enable fast-external-fallover neighbor pg1 peer-group neighbor pg1 remote-as 55001 neighbor pg1 password \$9\$MCgKGaT6OASX68/7TC6Lw== neighbor pg1 bfd neighbor pg2 peer-group neighbor pg2 remote-as 55002 neighbor pg2 password \$9\$ufD04Gw +49ex4H8UtvifqA== neighbor pg2 bfd neighbor 10.20.20.5 remote-as 4200000000 neighbor 10.20.20.5 next-hop-self address-family ipv4 unicast network 172.31.254.46/32 network 172.31.254.176/32 maximum-paths 8 graceful-restart ! address-family ipv4 unicast vrf tenlvrf1 redistribute connected maximum-paths 8 ! address-family ipv6 unicast ! address-family ipv6 unicast vrf tenlvrf1 redistribute connected maximum-paths 8 ! address-family l2vpn evpn graceful-restart ! !</pre>
---	---

**Note**

The MD5 password cannot be set or unset on an existing BGP peer-group present within a peer-group instance. You need to remove the BGP peer-group from the BGP peer-group instance and then add back the BGP peer-group to the peer-group instance with the desired MD5 password configuration.

Configure BGP MD5 Authentication for Tenant BGP Peer and Peer-group Securely

You can securely provide MD5 passwords during BGP peer-group create or update operations.

About This Task

[Configure BGP MD5 Authentication for Tenant BGP Peer](#) on page 28 and [Configure BGP MD5 Authentication for Tenant BGP Peer-group](#) on page 32 present instructions for providing an md5-password per BGP peer or peer-group during the BGP peer or peer-group create and update operations.

This topic provides an additional method for doing so - in a secure manner - using the `--md5-password-prompt-enable=true` option in the `efa tenant service bgp peer create` and `efa tenant service bgp peer-group create` commands.

You are prompted to supply a password the same number of times as the number of BGP peer or peer-group inputs you specify in the command. Then, you can choose to type in a password, in which case, it is not shown. Alternatively, you can skip the input of the password altogether by pressing Enter.



Note

You can choose to provide the BGP peer or peer-group md5-password either in a secure manner (using the prompt) or in an unsecure manner, as described in the topics referenced above, but not both.

Procedure

1. Run the `efa tenant service bgp peer create` command to create or update a BGP peer.

For example:

```
efa tenant service bgp peer create --name bgp173-2501 --tenant tenant11
--ipv4-uc-nbr 10.20.246.6,v1:25.1.1.3,5901
--ipv4-uc-nbr-bfd 10.20.246.6,v1:25.1.1.3,true
--ipv6-uc-nbr 10.20.246.5,v1:25:1::3,5901
--ipv6-uc-nbr-bfd 10.20.246.5,v1:25:1::3,true
--md5-password-prompt-enable=true
```

The following output is displayed:

```
Enter Md5 Password for 10.20.246.6::v1::25.1.1.3:
```

For the first prompt, let's say you type in a password. It is not displayed on the screen. You are prompted for the second password:

```
Enter Md5 Password for 10.20.246.5::v1::25:1::3:
```

Let's say that this time, you do not type in a password at all but press Enter.

The following output is displayed:

```
BgpService created successfully.
```

2. Run the `efa tenant service bgp peer-group create` command to create or update a BGP peer-group.

For example:

```
efa tenant service bgp peer-group create --tenant "tenant11" --name "v1-PeerGrp"
--pg-name 10.20.246.5:v1-PeerGrp --pg-asn 10.20.246.5,v1-
PeerGrp:5200
--pg-bfd-enable 10.20.246.5,v1-PeerGrp:true
--pg-name 10.20.246.5:v3-PeerGrp --pg-asn 10.20.246.5,v3-
PeerGrp:5201
--pg-bfd-enable 10.20.246.5,v3-PeerGrp:true
--pg-name 10.20.246.6:v1-PeerGrp --pg-asn 10.20.246.6,v1-
PeerGrp:5200
--pg-bfd-enable 10.20.246.6,v1-PeerGrp:true
--pg-md5-password-prompt-enable=true
```

The following output is displayed:

```
Enter Md5 Password for 10.20.246.5::v1-PeerGrp:
```

For the first prompt, let's say you type in a password. It is not displayed on the screen.

You are prompted for the second password:

```
Enter Md5 Password for 10.20.246.5::v3-PeerGrp:
```

Let's say you type in a password this time, too. It is not displayed on the screen.

Now you are prompted a third time:

```
Enter Md5 Password for 10.20.246.6::v1-PeerGrp:
```

Let's say that this time, you do not type in a password at all but press Enter.

The following output is displayed:

```
BgpService created successfully.
```

Configure BGP MD5 Authentication for Backup Routing Neighbors

About This Task

The BGP MD5 password for the backup routing neighbors is the same as the one set at the fabric setting level. The BGP MD5 password for the backup routing neighbors is configured during the configuration of VRF on SLX.

If the MD5 password setting is updated or set on a provisioned fabric followed by **efa fabric configure**, then the modified backup routing neighbor configuration is applied on all the tenant VRF backup routing BGP neighbors during **efa fabric configure**.

Procedure

1. Run the series of commands as shown in the following example.

```
efa fabric show --name fabric1
Fabric Name: fabric1, Fabric Description: , Fabric Type: non-clos
+-----+-----+-----+-----+-----+-----+-----+
+-----+-----+-----+-----+-----+-----+-----+
| IP ADDRESS | RACK | HOST NAME | ASN | ROLE | DEVICE STATE | APP STATE |
| CONFIG GEN REASON | PENDING CONFIGS | VTLB ID | LB ID |
+-----+-----+-----+-----+-----+-----+-----+
+-----+-----+-----+-----+-----+-----+
| 10.20.246.15 | rack1 | Avalanche-01 | 4200000000 | leaf | provisioned | cfg in-sync |
| NA | NA | | 2 | 1 |
| 10.20.246.16 | rack1 | Avalanche-02 | 4200000000 | leaf | provisioned | cfg in-sync |
| NA | NA | | 2 | 1 |
+-----+-----+-----+-----+-----+-----+-----+
+-----+-----+-----+-----+-----+-----+

efa fabric setting show --name fabric1 --advanced | grep -i "backup routing"
| Backup Routing Enable | Yes |
| Backup Routing IPv4 Range | 10.40.40.0/24 |
| Backup Routing IPv6 Range | fd40:4040:4040:1::/120 |
```

efa tenant show

```

+-----+-----+-----+-----+-----+-----+
+-----+
| Name | Type | VLAN Range | L2VNI Range | L3VNI Range | VRF Count | Enable BD |
|      | Ports |           |             |             |           |           |
+-----+-----+-----+-----+-----+-----+
+-----+
| ten1 | private | 11-20 | 20001-20020 | 21001-210020 | 10 | false |
| 10.20.246.15[0/1-10] |
|      |         |       |             |             |     |       |
| 10.20.246.16[0/1-10] |
+-----+-----+-----+-----+-----+-----+
+-----+

```

efa tenant vrf show

```

+-----+-----+-----+-----+-----+-----+
+-----+-----+-----+-----+-----+
| Name | Tenant | Routing Type | Centralized Routers | Redistribute | Max Path |
| Local Asn | Enable GR | State | Dev State | App State |
+-----+-----+-----+-----+-----+-----+
+-----+-----+-----+-----+-----+
| ten1vrf1 | ten1 | distributed | | connected | 8 |
| | false | vrf-create | not-provisioned | cfg-ready |
+-----+-----+-----+-----+-----+-----+
+-----+-----+-----+-----+-----+

```

efa fabric setting show --name fabric1 -advanced | grep -i MD5

```

| MD5 Password Enable | Yes |
| MD5 Password | $9$jrujIQqNxWkAyUOoI4cMtzhc4oP2VGREKwLOsSKH8bw= |

```

efa tenant epg show --name tenlepg1 --tenant ten1 -detail

```

=====
Name          : tenlepg1
Tenant        : ten1
Type          : extension
State         :
Description   :
Ports         : 10.20.246.15[0/1]
POs           :
Port Property : SwitchPort Mode           : trunk
               : Native Vlan Tagging                 : false
               : Single-Homed BFD Session Type       : auto
NW Policy     : Ctag Range                   : 11
               : VRF                                   : ten1vrf1
               : L3Vni                                : 21001
+-----+-----+-----+-----+-----+-----+
+-----+-----+-----+-----+-----+
+-----+-----+-----+-----+-----+
| Ctag | Ctag | L2Vni | BD Name | Anycast IPv4 | Anycast IPv6 |
| Local IP | IP MTU | IPv6 ND | IPv6 ND | IPv6 ND | Dev |
State | App State |
| Description | | | | |
[Device-IP->Local-IP] | Mtu | Managed Config | Other Config |
| | | | |
+-----+-----+-----+-----+-----+-----+
+-----+-----+-----+-----+-----+
+-----+-----+-----+-----+-----+
| 11 | Tenant L3 Extended VLAN | 20001 | | 10.0.11.1/24 |
| | | | | false | false |
provisioned | cfg-in-sync |
+-----+-----+-----+-----+-----+-----+

```

```

+-----+-----+-----+-----+-----+
+-----+-----+

```

2. Complete the configuration on SLX as provided in the following example.

<pre> L1# show running-config router bgp router bgp local-as 4200000000 capability as4-enable fast-external-fallover neighbor 10.20.20.3 remote-as 4200000000 neighbor 10.20.20.3 next-hop-self address-family ipv4 unicast network 172.31.254.71/32 network 172.31.254.151/32 maximum-paths 8 graceful-restart ! address-family ipv4 unicast vrf tenlvrf1 redistribute connected neighbor 10.40.40.252 remote-as 4200000000 neighbor 10.40.40.252 next-hop- self neighbor 10.40.40.252 password \$9\$jrujIQqNxWkAyUOoI4cMtzhc4oP2VGRE KwL0sSKH8bw= maximum-paths 8 ! address-family ipv6 unicast ! address-family ipv6 unicast vrf tenlvrf1 redistribute connected neighbor fd40:4040:4040:1::fe remote-as 4200000000 neighbor fd40:4040:4040:1::fe next-hop-self neighbor fd40:4040:4040:1::fe password \$9\$jrujIQqNxWkAyUOoI4cMtzhc4oP2VGRE KwL0sSKH8bw= neighbor fd40:4040:4040:1::fe activate maximum-paths 8 ! address-family l2vpn evpn graceful-restart ! ! </pre>	<pre> L2# show running-config router bgp router bgp local-as 4200000000 capability as4-enable fast-external-fallover neighbor 10.20.20.2 remote-as 4200000000 neighbor 10.20.20.2 next-hop-self address-family ipv4 unicast network 172.31.254.71/32 network 172.31.254.195/32 maximum-paths 8 graceful-restart ! address-family ipv4 unicast vrf tenlvrf1 redistribute connected neighbor 10.40.40.253 remote-as 4200000000 neighbor 10.40.40.253 next-hop- self neighbor 10.40.40.253 password \$9\$jrujIQqNxWkAyUOoI4cMtzhc4oP2VGRE KwL0sSKH8bw= maximum-paths 8 ! address-family ipv6 unicast ! address-family ipv6 unicast vrf tenlvrf1 redistribute connected neighbor fd40:4040:4040:1::ff remote-as 4200000000 neighbor fd40:4040:4040:1::ff next-hop-self neighbor fd40:4040:4040:1::ff password \$9\$jrujIQqNxWkAyUOoI4cMtzhc4oP2VGRE KwL0sSKH8bw= neighbor fd40:4040:4040:1::ff activate maximum-paths 8 ! address-family l2vpn evpn graceful-restart ! ! </pre>
--	--



Security Hardening

[Overview](#) on page 39

[System Hardening for CIS-CAT Assessments](#) on page 40

[The iptables Policy](#) on page 43

[Secure the Grub Boot Loader](#) on page 45

[Install the Linux Audit System](#) on page 46

[Install and Use OSSEC](#) on page 47

[Secure NTP and EFA](#) on page 49

[Secure DNS and EFA](#) on page 50

[Detect Rootkits with rkhunter](#) on page 51

This document provides guidance for the installation and usage of open source security tools to achieve a hardened operational security stance for EFA.

Overview

The following provides an overview of techniques for hardening security in EFA.

Some basic knowledge of security principles and operations of the Linux operating system and associated technologies is assumed.

Note: This document assumes that EFA has been installed in secure mode and is operational. For details on how to achieve this, refer to the [Extreme Fabric Automation Administration Guide, 2.7.0](#), which includes details on EFA security options and commands.

Security hardening topics included in this document are:

- **CIS-CAT security hardening:** Details of a custom python script from Extreme Networks that hardens the underlying operating system.
- **Iptables firewall:** Securing the EFA networking stance.
- **Grub boot loader security:** How to set a hardened security posture for Grub.
- **System auditing with `auditd`:** Instructions for monitoring various aspect of system runtime activities.
- **OSSEC HIDS installation and usage:** A broad set of indicators relevant for host intrusion detection.
- **Authenticated NTP:** How to ensure that NTP communications are authenticated.

- **Secure DNS:** Details about encrypted DNS communications.
- **Detecting rootkits with `rkhunter`:** Specialized run time checks for various types of Linux rootkits.

System Hardening for CIS-CAT Assessments

Harden your system to enable CIS-CAT assessments.

Before You Begin

TPVM provides a security hardening capability in the form of a python script located at `/opt/security/extr-granite.py` (for TPVM installations valid from version 4.5.0). The goal of this script is to modify various system security settings to achieve a more secure state under the examination of the CIS-CAT host scanner. Specifically, EFA running on TPVM is deployed into the Ubuntu server environment, and it is this environment that is hardened by `extr-granite.py`.

To perform the steps that follow, make sure you have Java Runtime Environment (JRE) installed. Also, ensure that the CIS-CAT scanner has been copied to the TPVM at `/root/cis-cat/Assessor-CLI`.

About This Task

The procedure itself involves the running of the `extr-granite.py` script. Notable aspects of the script are:

- The `extr-granite.py` script keeps a dedicated Git repository for all changes it makes to the underlying file system. This repository is located at `/opt/extr-granite-hardening/OS-files-git`.
- Git tags are used by `extr-granite.py` for every run cycle. This allows easy comparison of what the script changes on the host operating system from one run to the next.
- A comprehensive log file is kept at `/opt/extr-granite-hardening/hardening-log`.

One powerful feature offered by the hardening script is that it can be run multiple times (and even done at every boot via the init scripts), and every run receives its own Git tag in the `OS-files-git` repository. This make it easy to see every change the script has made, going back to the initial import.

Perform the following steps to run the security script and verify that it is working:

Procedure

1. Install CIS-CAT on the host and produce a scan result.

Here is an example of the abbreviated output. There is a CIS-CAT score of 56.11% on the TPVM.

```
root@tpvm:~# apt-get install default-jre
root@tpvm:~# cd /root/cis-cat/Assessor-CLI
root@tpvm:~# chmod u+x Assessor-CLI.sh
root@tpvm:~# ./Assessor-CLI.sh -b
benchmarks/CIS_Ubuntu_Linux_18.04_LTS_Benchmark_v2.0.1-xccdf.xml -html
-----
**** Assessment Results Summary ****
-----
Total # of Results: 241
Total Scored Results: 180
Total Pass: 101 Total
Fail: 78
Total Error: 1
Total Unknown: 0
```



```
Total Not Applicable: 0
Total Not Checked: 20
Total Not Selected: 37
Total Informational: 4
```

```
-----
***** Assessment Scoring *****
-----
```

```
Score Earned: 101.0
Maximum Available: 180.0
Total: 56.11%
-----
```

2. Run `/opt/security/extr-granite.py`.

Here is an example of an abbreviated output.

```
root@node-1:/opt/security# ./extr-granite.py
Initialized empty Git repository in /opt/extr-granite-hardening/OS-files-git/.git/
[master (root-commit) 1e2796b] initial import
1 file changed, 1 insertion(+)
create mode 100644 README
[+] ./extr-granite.py version: 0.89 - Initialized
sending incremental file list
/usr/
/usr/sbin/
/usr/sbin/grub-mkconfig

sent 3,467 bytes received 47 bytes 7,028.00 bytes/sec
total size is 8,219 speedup is 2.34
[master 8b2b4bd] initial import: /usr/sbin/grub-mkconfig, for CIS-CAT test: 1.4.1
Ensure permissions on bootloader config are not overridden
1 file changed, 311 insertions(+)
create mode 100755 usr/sbin/grub-mkconfig
sending incremental file list
/etc/
/etc/sysctl.conf

sent 1,163 bytes received 39 bytes 2,404.00 bytes/sec
total size is 2,683 speedup is 2.23
[master 2f4ad6c] initial import: /etc/sysctl.conf, for CIS-CAT test: 1.5.2 Ensure
address space layout randomization (ASLR) is enabled
1 file changed, 77 insertions(+)
create mode 100644 etc/sysctl.conf
kernel.randomize_va_space = 2
sending incremental file list
/etc/security/
/etc/security/limits.conf

.....
.....
.....

sent 1,168 bytes received 44 bytes 2,424.00 bytes/sec
total size is 2,306 speedup is 1.90
[master 7ef96f4] file: /etc/pam.d/su, CIS-CAT test: 5.7 Ensure access to the su
command is restricted
1 file changed, 1 insertion(+)
sending incremental file list
/etc/group

sent 544 bytes received 36 bytes 1,160.00 bytes/sec
total size is 832 speedup is 1.43
[master 374f6c9] file: /etc/group, CIS-CAT test: 5.7 Ensure access to the su command
is restricted (2)
1 file changed, 1 insertion(+), 1 deletion(-)
sending incremental file list
```

```

/etc/profile

sent 436 bytes received 36 bytes  944.00 bytes/sec
total size is 619 speedup is 1.31
[master b305cfe] file: /etc/profile, CIS-CAT test: 5.5.5 Ensure default user shell
timeout is 900 seconds or less
1 file changed, 3 insertions(+)

[+] Total checks run: 55

root@node-1:/opt/security#

```

The following example shows all changes from the initial import to the tag of the first run:

```
extr-granite-run1
```

Further, the example does the same `git diff`, except it shows just the changes that were made to the original `/etc/ssh/sshd_config` file.

```

root@tpvm:/opt/extr-granite-hardening/OS-files-git# git tag -l
extr-granite-initial-import
extr-granite-run1
root@tpvm:/opt/extr-granite-hardening/OS-files-git# git diff extr-granite-initial-
import extr-granite-run1
diff --git a/etc/issue b/etc/issue
index 80ae21e..5192c40 100644
--- a/etc/issue
+++ b/etc/issue
@@ -1,2 +1 @@
-Ubuntu 18.04.5 LTS \n \l
-
+Extreme Networks, Inc. EFA product. Authorized users only. All activity may be
monitored and reported.
diff --git a/etc/issue.net b/etc/issue.net
index 5e9e2fa..5192c40 100644
--- a/etc/issue.net
+++ b/etc/issue.net
@@ -1 +1 @@
-Ubuntu 18.04.5 LTS
+Extreme Networks, Inc. EFA product. Authorized users only. All activity may be
monitored and reported.
diff --git a/etc/modprobe.d/cramfs.conf b/etc/modprobe.d/cramfs.conf
new file mode 100644
index 0000000..b77c93a
--- /dev/null
+++ b/etc/modprobe.d/cramfs.conf
@@ -0,0 +1 @@
+install cramfs /bin/true
diff --git a/etc/modprobe.d/freevxfs.conf b/etc/modprobe.d/freevxfs.conf
new file mode 100644
index 0000000..72d4aec
--- /dev/null
+++ b/etc/modprobe.d/freevxfs.conf
@@ -0,0 +1 @@
+install freevxfs /bin/true
root@tpvm:/opt/extr-granite-hardening/OS-files-git# ls etc/ssh/sshd_config
etc/ssh/sshd_config
root@tpvm:/opt/extr-granite-hardening/OS-files-git# git help diff^C
root@tpvm:/opt/extr-granite-hardening/OS-files-git# git diff extr-granite-initial-
import extr-granite-run1 ./etc/ssh/sshd_config
diff --git a/etc/ssh/sshd_config b/etc/ssh/sshd_config
index 3f0e52e..f640120 100644
--- a/etc/ssh/sshd_config

```

```

+++ b/etc/ssh/sshd_config
@@ -121,4 +121,11 @@ Subsystem sftp      /usr/lib/openssh/sftp-server
#       PermitTTY no
#       ForceCommand cvs server
PasswordAuthentication yes
-MaxStartups 30:30:100
+MaxStartups 10:30:60
+MaxAuthTries 4
+MACs hmac-sha2-512-etm@openssh.com,hmac-sha2-256-etm@openssh.com,hmac-sha2-512,hmac-
sha2-256
+KexAlgorithms curve25519-sha256,curve25519-sha256@libssh.org,diffie-hellman-group14-
sha256,diffie-hellman-group16-sha512,diffie-hellman-group18-sha512,ecdh-sha2-
nistp521,ecdh-sha2-nistp384,ecdh-sha2-nistp256,diffie-hellman-group-exchange-sha256
+ClientAliveInterval 300
+ClientAliveCountMax 0
+LoginGraceTime 60
+Banner /etc/issue.net

```



Note

After running the security hardening script, reconnect to the existing TPVM SSH sessions (active/standby/vip) for the new ssh parameters to take effect.

3. Rerun the CIS-CAT auditor and verify that the scan results produce a score greater than 80%.

In the following example, the results have improved from 56.11% originally to 81.67%.

Note that the hardening results are against version 4.2.0 of the CIS-CAT Assessor. The score achieved by the hardening script is reduced against the latest version (4.7.0). Improvements are expected in forthcoming releases.

```

root@tpvm:~# ./Assessor-CLI.sh -b
benchmarks/CIS_Ubuntu_Linux_18.04_LTS_Benchmark_v2.0.1-xccdf.xml -html
-----
**** Assessment Results Summary ****
-----
      Total # of Results: 241
Total Scored Results: 180
      Total Pass: 147
      Total Fail: 32
      Total Error: 1
      Total Unknown: 0
Total Not Applicable: 0
      Total Not Checked: 20
      Total Not Selected: 37
      Total Informational: 4
-----
**** Assessment Scoring ****
-----
      Score Earned: 147.0
      Maximum Available: 180.0
      Total: 81.67%
-----

```

The iptables Policy

This topic describes the custom iptables policy that EFA builds to firewall off services.

As a core component of EFA, Kubernetes uses iptables to control the network connections between pods (and between nodes), handling many of the networking and port forwarding rules. EFA builds a custom iptables policy to firewall off services (such as the MySQL database) on the EFA management

interface. The EFA iptables policy is instantiated by default at installation time and is enabled at boot, through the system service.

To see the status of the EFA iptables policy, in addition to the policy itself, run the following commands:

```
$ ssh -l extreme 192.168.10.109
Password:
extreme@tpvm:~$ sudo su -
root@tpvm:~#
root@tpvm:~# systemctl status efa-iptables
  efa-iptables.service - iptables rules for EFA
    Loaded: loaded (/lib/systemd/system/efa-iptables.service; enabled; vendor preset:
    enabled)
    Active: active (exited) since Fri 2020-10-09 20:48:03 UTC; 1 day 17h ago
    Main PID: 19384 (code=exited, status=0/SUCCESS)
    Tasks: 0 (limit: 4638)
    CGroup: /system.slice/efa-iptables.service

Oct 09 20:48:03 tpvm systemd[1]: Starting iptables rules for EFA...
Oct 09 20:48:03 tpvm systemd[1]: Started iptables rules for EFA.
root@tpvm:~# iptables -v -nL EFA_INPUT
Chain EFA_INPUT (1 references)
pkts bytes target      prot opt in      out     source        destination
  0     0 DROP          tcp  --  eth0    *        0.0.0.0/0     0.0.0.0/0
multiport dports 1024:6513,6515:8077,8079:65535 ctstate NEW
  0     0 DROP          udp  --  eth0    *        0.0.0.0/0     0.0.0.0/0
multiport dports 1024:65535 ctstate NEW
```

With iptables policy active, it should not be possible to connect to the MySQL database on TCP port 3306 on the management interface from an external host. Use Nmap to verify that port 3306 has, indeed, been firewalled off:

```
# nmap -n -p 3306 -sV 192.168.10.109

Starting Nmap 7.60 ( https://nmap.org ) at 2020-10-11 14:42 UTC
Nmap scan report for 192.168.10.109
Host is up (0.0039s latency).

PORT      STATE SERVICE VERSION
3306/tcp  filtered mysql

Service detection performed. Please report any incorrect results at https://nmap.org/submit/ .
Nmap done: 1 IP address (1 host up) scanned in 0.81 seconds

Service detection performed. Please report any incorrect results at https://nmap.org/submit/ .
Nmap done: 1 IP address (1 host up) scanned in 0.85 seconds
```

Currently visible services include the following:

- 22 tcp - SSH port (EFA management)
- 80 tcp - HTTP port (EFA web interface)
- 111 tcp - NFS port (not used by EFA)
- 443 tcp - HTTPs port (EFA northbound REST interface)
- 514 tcp - Syslog port (syslog receiver)
- 6514 tcp - Secure syslog port (TLS-enabled syslog receiver)
- 8078 tcp - EFA host monitoring service

- 30148 tcp - Not used by EFA and can be disabled
- 30672 tcp - RabbitMQ port for EFA
- 31672 tcp - RabbitMQ management port for EFA
- 31730 tcp - Not used by EFA and can be turned off
- 111 udp - rpcbind NFS port (not used by EFA)

Forthcoming releases will further tighten the iptables policy so that even Kubernetes services are not visible over the network on the EFA management interface. At a minimum, ports 111 (both TCP and UDP), 30148, 30672 (RabbitMQ), 31672 (RabbitMQ), and 31730 will be firewalled off by default.

Secure the Grub Boot Loader

This topic provides instructions for adding the Grub boot loader to the security posture.

About This Task

Securing the Grub boot loader is an important addition to the security posture for the operating system where EFA is deployed. There are two general phases for securing the boot loader:

- Set a password in the Grub configuration to harden against modifications to the Linux kernel boot-time command line.
- Set a password for the 'root' user to protect against attempts to acquire single-user mode at boot.

Procedure

1. Set a password in the Grub configuration:
 - a. Acquire root and then run the `grub-mkpasswd-pbkdf2` command (full output is shown below).
 - b. Append the password hash and the string `set superusers="root"` to the file `/etc/grub.d/40_custom`.
 - c. Add `--unrestricted` to the `"CLASS="` definition line in `/etc/grub.d/10_linux`.
 - d. Run the command `update-grub`.

```
root@tpvm:~# grub-mkpasswd-pbkdf2
Enter password:
Reenter password:
PBKDF2 hash of your password is
grub.pbkdf2.sha512.10000.72C8CE3112C007A315A94DD7A63B58392DD00653ACAF8795C8528D83967
FA24105B0B53D0092522460532AF05C60EE3E0C7EAC95213E865DF31580A341188ABC.843EF94A9C8EE8
AC1776F5B88261D1B6DE437A70AEABE3C814764596F696EE5F7FDF912E63B4D47AE3E7BB468A6B639F00
051D142698142EF158E6C141CF38B7
root@tpvm:~# cat >> /etc/grub.d/40_custom
set superusers="root"
password_pbkdf2 root
grub.pbkdf2.sha512.10000.A577D1C8F13C93B82EA5E25E834D5BD88ECB94A5B42F2DABE4FB7A235F3
A25A12E6542CB5DA9620B2E0342FE28A4F066BE1B99F2EFBE8C0688FBE11FDB3138DD.2C7C81C7FA0404
C768DDCE097B3AA8DD08C042B4FDBA089C0837F91B6C8864EE83B19C8C6D4C5C126E76FA20BE93789920
913B12CAC841CA65EA3BAD5921F8D5
root@tpvm:~# <edit /etc/grub.d/10_linux to make the CLASS line look like the
following>
root@tpvm:~# grep CLASS /etc/grub.d/10_linux | head -n 1
CLASS="--class gnu-linux --class gnu --class os --unrestricted"
root@tpvm:~ # update-grub
Sourcing file `/etc/default/grub'
Generating grub configuration file ...
Found linux image: /boot/vmlinuz-5.4.0-48-generic
```

```

Found initrd image: /boot/initrd.img-5.4.0-48-generic
Found linux image: /boot/vmlinuz-5.3.0-40-generic
Found initrd image: /boot/initrd.img-5.3.0-40-generic
Found linux image: /boot/vmlinuz-4.15.0-118-generic
Found initrd image: /boot/initrd.img-4.15.0-118-generic
Found linux image: /boot/vmlinuz-4.15.0-88-generic
Found initrd image: /boot/initrd.img-4.15.0-88-generic
done

```

2. Set a password for the 'root' user by running the following commands:

```

root@tpvm~:# passwd
Enter new UNIX password:
Retype new UNIX password:
passwd: password updated successfully

```

Install the Linux Audit System

Provides instructions for installing the Linux Audit System.

About This Task

EFA is based on the Ubuntu operating system, which by default, does not come with auditd (Linux Audit System) installed. Perform the instructions that follow to install it in EFA. During the installation process, auditd is enabled by default via systemctl and starts writing audit trail log data to the file `/var/log/audit/audit.log`.

To install, run the following command:

```
# apt-get install auditd
```

This is example output:

```

type=DAEMON_START msg=audit(1591152521.117:3494): op=start ver=2.8.2 format=raw
kernel=5.3.0-53-generic auid=4294967295 pid=27162 uid=0 ses=4294967295 subj=unconfined
res=success
type=CONFIG_CHANGE msg=audit(1591152521.153:23): op=set audit_backlog_limit=8192 old=64
auid=4294967295 ses=4294967295 res=1
type=CONFIG_CHANGE msg=audit(1591152521.157:24): op=set audit_failure=1 old=1
auid=4294967295 ses=4294967295 res=1
type=CONFIG_CHANGE msg=audit(1591152521.157:25): op=set audit_backlog_wait_time=0
old=15000 auid=4294967295 ses=4294967295 res=1
type=SERVICE_START msg=audit(1591152521.157:26): pid=1 uid=0 auid=4294967295
ses=4294967295 msg='unit=auditd comm="systemd" exe="/lib/systemd/systemd" hostname=?
addr=? terminal=? res=success'

```

You can produce more interesting data from the `audit.log` file when you authenticate to the EFA host. In this case, the audit trail data for the user 'extreme' authenticating via SSH is displayed:

```

type=USER_LOGIN msg=audit(1591975761.779:39): pid=7894 uid=0 auid=4294967295
ses=4294967295 msg='op=login acct="extreme" exe="/usr/sbin/sshd" hostname=?
addr=192.168.10.12 terminal=sshd res=failed'
type=USER_AUTH msg=audit(1591975763.219:40): pid=7894 uid=0 auid=4294967295
ses=4294967295 msg='op=PAM:authentication acct="extreme" exe="/usr/sbin/sshd"
hostname=192.168.10.12 addr=192.168.10.12 terminal=ssh res=success'
type=USER_ACCT msg=audit(1591975763.219:41): pid=7894 uid=0 auid=4294967295
ses=4294967295 msg='op=PAM:accounting acct="extreme" exe="/usr/sbin/sshd"
hostname=192.168.10.12 addr=192.168.10.12 terminal=ssh res=success'
type=CRED_ACQ msg=audit(1591975763.223:42): pid=7894 uid=0 auid=4294967295 ses=4294967295
msg='op=PAM:setcred acct="extreme" exe="/usr/sbin/sshd" hostname=192.168.10.12
addr=192.168.10.12 terminal=ssh res=success'
type=LOGIN msg=audit(1591975763.223:43): pid=7894 uid=0 old-auid=4294967295 auid=1000

```

```

tty=(none) old-ses=4294967295 ses=127 res=1
type=USER_START msg=audit(1591975764.487:44): pid=7894 uid=0 auid=1000 ses=127
msg='op=PAM:session_open acct="extreme" exe="/usr/sbin/sshd" hostname=192.168.10.12
addr=192.168.10.12 terminal=ssh res=success'
type=CRED_ACQ msg=audit(1591975764.491:45): pid=8113 uid=0 auid=1000 ses=127
msg='op=PAM:setcred acct="extreme" exe="/usr/sbin/sshd" hostname=192.168.10.12
addr=192.168.10.12 terminal=ssh res=success'
type=USER_LOGIN msg=audit(1591975764.547:46): pid=7894 uid=0 auid=1000 ses=127
msg='op=login id=1000 exe="/usr/sbin/sshd" hostname=192.168.10.12 addr=192.168.10.12
terminal=/dev/pts/1 res=success'

```

Further, when the 'extreme' user authenticates to the 'efa' command line with `efa login`, the following audit trail message is generated (in case of success):

```

type=USER_AUTH msg=audit(1591975780.823:47): pid=21139 uid=0 auid=4294967295
ses=4294967295 msg='op=PAM:authentication acct="extreme" exe="/apps/bin/hostauth"
hostname=? addr=? terminal=? res=success'

```

And here is the output in case of failures:

```

type=USER_AUTH msg=audit(1591976323.760:58): pid=21139 uid=0 auid=4294967295
ses=4294967295 msg='op=PAM:authentication acct="extreme" exe="/apps/bin/hostauth"
hostname=? addr=? terminal=? res=failed'

```

Install and Use OSSEC

OSSEC is a multiplatform, open source, and free Host Intrusion Detection System (HIDS). The OSSEC HIDS project is the current open source standard-bearer for host-based intrusion detection on Linux.

This topic shows you how to install and use OSSEC HIDS with EFA.

Before You Begin

Installation of OSSEC requires a few packages to be installed on EFA to achieve a functioning OSSEC HIDS instance. Run the following to install the prerequisite packages:

```

# apt-get install
    libpcre2-dev libssl-dev zlib1g-dev libevent-dev build-essentials

```

Procedure

1. Download the latest release of OSSEC HIDS from GitHub. Then extract, install, and start it, using the following commands:

```

# wget https://github.com/ossec/ossec-hids/archive/3.6.0.tar.gz
# tar xvfz 3.6.0.tar.gz
# cd ossec-hids-3.6.0
# ./install.sh
# systemctl start ossec

```

At this point, OSSEC HIDS is running on EFA, and various real-time monitoring tasks are run. Log and alert information is produced in the `/var/ossec/logs` directory. Some example alerts are illustrated in the scenarios below.

- In the first scenario, a process starts listening on a previously closed TCP port on port 12345. You can easily achieve this with a call to `netcat nc -l -p 12345`.

Here is what OSSEC HIDS reports in the `/var/ossec/logs/alerts/alerts.log` file. The service is bolded:

```
** Alert 1592072520.408: mail - ossec,
2020 Jun 13 18:22:00 tpvm->netstat -tan |grep LISTEN |egrep -v '(127.0.0.1| \\\1)' |
sort
Rule: 533 (level 7) -> 'Listened ports status (netstat) changed (new port opened or
closed).'
```

ossec: output: 'netstat -tan grep LISTEN egrep -v '(127.0.0.1 \\\1)' sort':					
tcp	0	0	0.0.0.0:12345	0.0.0.0:*	LISTEN
tcp	0	0	0.0.0.0:22	0.0.0.0:*	LISTEN
tcp	0	0	0.0.0.0:30085	0.0.0.0:*	LISTEN
tcp	0	0	0.0.0.0:31672	0.0.0.0:*	LISTEN
tcp6	0	0	:::10250	:::*	LISTEN
tcp6	0	0	:::12865	:::*	LISTEN
tcp6	0	0	:::22	:::*	LISTEN
tcp6	0	0	:::514	:::*	LISTEN
tcp6	0	0	:::6443	:::*	LISTEN

```
Previous output:
ossec: output: 'netstat -tan |grep LISTEN |egrep -v '(127.0.0.1| \\\1)' | sort':
tcp      0      0 0.0.0.0:22          0.0.0.0:*          LISTEN
tcp      0      0 0.0.0.0:30085       0.0.0.0:*          LISTEN
tcp      0      0 0.0.0.0:30335       0.0.0.0:*          LISTEN
tcp      0      0 0.0.0.0:31672       0.0.0.0:*          LISTEN
tcp6     0      0 :::10250            :::*                LISTEN
tcp6     0      0 :::12865            :::*                LISTEN
tcp6     0      0 :::22               :::*                LISTEN
tcp6     0      0 :::514              :::*                LISTEN
tcp6     0      0 :::6443             :::*                LISTEN
```



Note

In general, any changes to the set of listening services is important to understand from a security perspective, since new services can potentially represent a system compromise.

- In the second scenario, a package update is made to the underlying Ubuntu OS. Perform this by running `apt-get upgrade`.

Here is what OSSEC HIDS reports again in the `/var/ossec/logs/alerts/alerts.log` file.

```
** Alert 1592073596.29116: mail - syslog,dpkg,config_changed,
2020 Jun 13 18:39:56 tpvm->/var/log/dpkg.log
Rule: 2902 (level 7) -> 'New dpkg (Debian Package) installed.'
2020-06-13 18:39:56 status installed initramfs-tools:all 0.130ubuntu3.9

** Alert 1592073596.29360: mail - syslog,dpkg,config_changed,
2020 Jun 13 18:39:56 tpvm->/var/log/dpkg.log
Rule: 2902 (level 7) -> 'New dpkg (Debian Package) installed.'
2020-06-13 18:39:56 status installed libc-bin:amd64 2.27-3ubuntu1

** Alert 1592073599.29598: mail - syslog,dpkg,config_changed,
2020 Jun 13 18:39:59 tpvm->/var/log/dpkg.log
```



```
Rule: 2902 (level 7) -> 'New dpkg (Debian Package) installed.'  
2020-06-13 18:39:57 status installed systemd:amd64 237-3ubuntu10.41
```



Note

Although upgrading OS packages is a routine maintenance activity, it is an important verification step for OSSEC HIDS to alert on all package upgrades. Similarly, new packages are also detected by OSSEC HIDS, and in both cases, having an understanding of what packages are changing on the system provides useful security auditing data.

Secure NTP and EFA

This topic presents options for using authenticated NTP in EFA.

About This Task

For EFA, authenticated NTP can be used in two possible configurations:

- EFA acts as a client to existing NTP infrastructure.
- EFA runs its own NTP server. Steps for achieving this solution are provided below.

Also, there are many different visions for secure NTP that can range from authentication (at the low end) to leveraging the newest NTS (Network Time Secure) protocol that leverages public key cryptography via TLS (at the high end).

EFA is based on Ubuntu 18.04, and `ntpsec` from the upstream package maintainers does not allow for a comprehensive and well-supported usage of NTS. However, if this is an absolute requirement, you can achieve NTS support on EFA by manually compiling and deploying a recent version of `ntpsec`, **but this is not a supported solution**. This comes at the cost of not using the sanctioned Ubuntu packaging system for package updates, and this tradeoff may not be worth it within certain operational environments.

Enforcing authentication where EFA acts as a client to existing NTP infrastructure

If EFA is to act only as a client for authenticated NTP, then upstream NTP servers where EFA is pointed also need to support authenticated NTP. You can use public NTP servers for this purpose, such as those of the US National Institute of Standards and Technology (NIST). For more information, see the [NIST website](#). The following is an excerpt from the site:

"The time messages will be authenticated using symmetric-key encryption in a manner that is fully compatible with the published NTP documentation. (Autokey and asymmetric key modes will not be used.) Each registered user will be assigned a unique encryption key, which will be linked to the IP address of the user's system. A registered user will be able to communicate with the authenticated server using this assigned encryption key or using a default key of 0, which is equivalent to disabling the encryption algorithm. Users who are not registered will not be able to connect to this server, but can use any of the other NIST servers, which will not be modified."

Procedure

Enforcing authentication where EFA provides its own NTP server

1. Populate the `/etc/ntp.keys` file with a SHA1 symmetric key. This can easily be done with the following command (which uses the `/etc/shadow` file as input to produce the SHA1 hash, and this hash must be shared with all clients that need to authenticate):

```
echo "3 SHA1 `shasum /etc/shadow |cut -d ' ' -f 1`" > ntp.keys
```

2. Add the line `'keys /etc/ntp.keys'` to the `/etc/ntp.conf` file on EFA.
3. Share the SHA1 symmetric key from [step 1](#) with any clients that are authorized to acquire NTP data from EFA. This could include the population of managed SLX devices for example.
4. Restart the NTP daemon on EFA, and verify that no start-up error for the `ntpd` process is logged to `/var/log/syslog`. For example, if the format of the `/etc/ntp.keys` file is invalid, then the following error is displayed.

```
Jun 12 20:38:36
    tpvm ntpd[16498]: authreadkeys: rejecting file '/etc/ntp.keys' after 1 error(s)
```



Note

From the perspective of NTP clients, if coming from a Linux system the `/etc/ntp.keys` file from the server configuration steps above should be available in the file system. For example, if this file is also placed at `/etc/ntp.keys`, then the step for acquiring NTP information from EFA is as follows:

```
# sntp -k /etc/ntp.keys <efa hostname>
```

Secure DNS and EFA

This topic provides steps for achieving secure DNS communications using `dnscrypt-proxy`.

Before You Begin

Achieving secure DNS communications can be a critical aspect of a strong operational security posture. Use the open source `'dnscrypt-proxy'` package to bring encrypted DNS communications to EFA.

About This Task

Procedure

1. Install `dnscrypt-proxy` by running:

```
apt-get install
    dnscrypt-proxy
```

2. Review the various resolvers that support encrypted DNS in the file `/usr/share/dnscrypt-proxy/dnscrypt-resolvers.csv` and select one. The name of this resolver is in the first column in this file. For example, the `'adguard-dns-ns1'` resolver is displayed here:

```
adguard-dns-ns1,"Adguard DNS 1","Remove ads and protect your computer from
    malware","Anycast","",https://adguard.com/en/adguard-dns/
overview.html,1,no,yes,no,176.103.130.130:5443,2,dnscrypt.default.ns1.adguard.com,D12B:
47F2:52DC:F2C2:BBF8:9910:86EA:F79C:E449:5D8B:16C8:A0C4:322E:52CA:3F39:0873,pk.default.n
s1.adguard.com
```

3. Edit the `/etc/dnscrypt/dnscrypt-proxy.conf` file, and set

```
ResolverName adguard-dns-ns1

LocalAddress 127.0.0.1:53
```

4. Edit the `/etc/systemd/system/sockets.target.wants/dnscrypt-proxy.socket` file and make sure the `[Socket]` section looks like this:

```
[Socket]
    ListenStream=127.0.0.1:53
    ListenDatagram=127.0.0.1:53
```

5. Reload systemd, disable and re-enable dnscrypt-proxy.socket, and reboot:

```
# systemctl daemon-reload
# systemctl disable dnscrypt-proxy.service
# systemctl enable dnscrypt-proxy.service
# reboot
```

At this point, dnscrypt-proxy should be functioning normally. One way to verify this is to look for DNS requests on the EFA management interface ('eth0' for TPVM installations of EFA). There should be no traffic on UDP port 53 because encrypted DNS traffic is sent over UDP port 443.

6. To verify, run a sniffer on eth0, cause the system to issue a DNS lookup, and make sure there are no UDP packets on port 53. The output should look similar to this:

```
# tcpdump -i eth0 -l -nn port 53 or port 443
tcpdump: verbose output suppressed, use -v or -vv for full protocol decode
listening on eth0, link-type EN10MB (Ethernet), capture size 262144 bytes

15:26:54.136556 IP 192.168.10.109.40632 > 176.103.130.130.443: UDP, length 512
15:26:54.151425 IP 176.103.130.130.443 > 192.168.10.109.40632: UDP, length 304
15:26:54.152783 IP 192.168.10.109.40632 > 176.103.130.130.443: UDP, length 512
15:26:54.166523 IP 176.103.130.130.443 > 192.168.10.109.40632: UDP, length 304
```

Detect Rootkits with rkhunter

This topic provides steps for installing and using Rootkit Hunter (rkhunter).

About This Task

Rootkit Hunter (rkhunter) is a standard tool for the detection of rootkits on Linux. The Ubuntu Linux distribution maintains a package for rkhunter. The complete set of checks that rkhunter performs provides a good security baseline for finding some of the most malicious elements of the offensive security landscape. It is recommended you regularly run the `rkhunter --check` command and review the contents of the `/var/log/rkhunter.log` file.

Procedure

1. Install rkhunter by running:

```
# apt-get install rkhunter
```

2. Run a scan for rootkits:

```
# rkhunter --check
T0rn Rootkit [ Not found ]
trNkit Rootkit [ Not found ]
Trojanit Kit [ Not found ]
Tuxtendo Rootkit [ Not found ]
URK Rootkit [ Not found ]
Vampire Rootkit [ Not found ]
VcKit Rootkit [ Not found ]
Volc Rootkit [ Not found ]
Xzibit Rootkit [ Not found ]
zaRwT.KiT Rootkit [ Not found ]
ZK Rootkit [ Not found ]
```

3. To dive in to additional details of what is actually being checked on the system, refer to the `/var/log/rkhunter.log` file. For example, in the following example, the scan looked for evidence of the T0rn rootkit and specifically, the existence of the following files were checked (output abbreviated):

```
[21:28:18] Checking for T0rn Rootkit...
[21:28:18]   Checking for file '/dev/.lib/lib/lib/t0rns' [ Not found ]
```

```
[21:28:18] Checking for file '/dev/.lib/lib/lib/du' [ Not found ]
[21:28:18] Checking for file '/dev/.lib/lib/lib/ls' [ Not found ]
[21:28:18] Checking for file '/dev/.lib/lib/lib/t0rnsb' [ Not found ]
[21:28:18] Checking for file '/dev/.lib/lib/lib/ps' [ Not found ]
[21:28:18] Checking for file '/dev/.lib/lib/lib/t0rnp' [ Not found ]
[21:28:18] Checking for file '/dev/.lib/lib/lib/find' [ Not found ]
[21:28:18] Checking for file '/dev/.lib/lib/lib/ifconfig' [ Not found ]
[21:28:18] Checking for file '/dev/.lib/lib/lib/pg' [ Not found ]
[21:28:18] Checking for file '/dev/.lib/lib/lib/ssh.tgz' [ Not found ]
[21:28:18] Checking for file '/dev/.lib/lib/lib/top' [ Not found ]
[21:28:19] Checking for file '/dev/.lib/lib/lib/sz' [ Not found ]
[21:28:19] Checking for file '/dev/.lib/lib/lib/login' [ Not found ]
[21:28:19] Checking for file '/dev/.lib/lib/lib/in.fingerd' [ Not found ]
[21:28:19] Checking for file '/dev/.lib/lib/lib/li0n.sh' [ Not found ]
```