



RH354 - RHEL 8 New Features for Experienced Linux Administrators

Travis Michette

Version 1.0

Table of Contents

5.2.2. Implementing Virtual Data Optimizer	35
5.3. Administering NFS Enhancements	
5.3.1. Mounting and Unmounting NFS Shares	
5.3.2. The nfsconf Tool	38
5.4. Automounting Network-Attached Storage	39
5.4.1. Mounting NFS Shares with the Automounter	39
5.4.1.1. Creating an Automount	39
5.4.1.2. Direct Maps	40
5.4.1.3. Indirect Wildcard Maps	40
6. Managing Containers with the New Runtime	41
6.1. Deploying Containers with the New Container Runtime	41
6.1.1. The Podman Container Engine	41
7. Implementing Enhanced Networking Features	47
7.1. Managing Server Firewalls in RHEL 8	47
7.1.1. Introducing Nftables	47
7.2. Configuring Server Networking with NetworkManager	49
8. Adapting to Virtualization Improvements	50
8.1. Configuring Virtual Machines	50
Appendix A: System Security Policy and Compliance	51
A.1. Customizing SCAP Content	51
A.2. Running a SCAP Scan with Custom Content	63
A.3. Creating an Ansible Remediation Playbook Based on SCAP Scan Results	72
Appendix B: IdM Server on RHEL8	78
B.1. Installing the IdM Server on RHEL8	78
B.2. Installing the IdM Client on RHEL8	81

Before You Begin





Table 1. Classroom Machines

Machine name	IP addresses	Role
workstation.lab.example.com	172.25.250.254	Gateway system to connect student private network to classroom server (must always be running). Graphical workstation used for system administration
servera.lab.example.com	172.25.250.10	First server
serverb.lab.example.com	172.25.250.11	Second server
serverc.lab.example.com	172.25.250.11	Third server (set as not booted)

1. Previewing Red Hat Enterprise Linux 8

1.1. Red Hat Enterprise Linux 8 Overview

New Features of Interest

- Restructure of channels BaseOS and Appstream
- Install and boot from NVDIMM devices
- BOOM Boot Manager
 - https://github.com/bmr-cymru/boom
 - https://www.redhat.com/en/blog/boom-booting-rhel-lvm-snapshots
 - https://www.youtube.com/watch?v=guM_jkA6Xfg
- Secure-boot Linux guests
- Kernel Changes
 - kernel-core provides core kernel module
 - kernel-modules and kernel-modules-extra contains kernel modules
 - kernel is now a meta which installs both kernel-core and kernel-modules
- · kdump supported earlier in the boot process
- Process scheduler enhancements with new class SCHED_DEADLINE
- Firewall changes
 - firewalld still daemon with firewall-cmd to manage firewall
 - iptables removed in favor of nft directly
 - iptables compatibility tools available for interacting with nft
- NetworkManager
 - nmcli preferred way to manage the network
 - · ifup/ifdown requires NetworkManager
- NTP time is now using chronyd as the default NTP implementation as ntpd is no longer available
- yum us now YUM4 which is dnf
 - Modules are now supported
 - Modules allow installation of packages independent of the Operating System version
 - Modules replace the Red Hat Software Collections Suite of packages for RHEL8
 - Modules are tied to application stream
 - $\circ~\text{dnf}$ is re-write of Yum but allows Yum4 to facilitate all commands using Yum
- Storage Changes

- Stratis Storage Manager is introduced
 - https://stratis-storage.github.io/
- VDO Virtual Data Optimizer is introduced
 - https://www.redhat.com/en/blog/look-vdo-new-linux-compression-layer
 - https://access.redhat.com/documentation/en-us/red_hat_enterprise_linux/7/html/storage_administration_guide/vdo
- XFS Copy-on-Write is enabled by default when filesystem created
- · LUKS2 encryption is the new encryption mechanism replacing LUKS1
 - https://fedoraproject.org/wiki/Disk_Encryption_User_Guide

Security Features

- New crypto policies
- · New sosreport tool which can support profiles
- Updated audit and rsyslog subsystems with new features
- tcp_wrapper support removed
- Graphical Environment
 - Wayland is default display server Xorg still currently available
 - $\circ\,$ Gnome has been updated and KDE has been removed
- Virtualization
 - e1000 network driver removed
 - KVM defaults to Q35 hardware model
 - virt-manager deprecated and cockpit preferred management solution

• Linux Containers

- podman runs daemonless containers
- buildah is used to build container images from scratch or a Dockerfile



Learning RHEL8 Online with Labs

Red Hat Learning Community (RHLC) RHEL8 Labs Post: https://learn.redhat.com/t5/Platform-Linux/Startwith-RHEL8-without-having-RHEL8/td-p/5813

Red Hat Labs: http://lab.redhat.com/

1.2. TMUX Usage

As part of the upgrade process and package replacement process in RHEL8, several packages have not only been deprecated, they've been completely removed. The **screen** package is one package that is no longer available for installation. Instead, a new terminal program **tmux** has been introduced to provide the **screen** functionality as well as other enhancements.

TMUX References

TMUX Usage



Red Hat Learning Community: https://learn.redhat.com/t5/Platform-Linux/Using-tmux-to-execute-commandson-servers-in-parallel/m-p/2200

Tactical TMUX: https://danielmiessler.com/study/tmux/

In the example below, we will explore the **screen** functionalities of TMUX with respect to attaching and detaching of sessions.

Installing TMUX

1. Install **tmux** with YUM

Listing 1. Installation of TMUX

```
[root@workstation *]# yum install tmux
Red Hat Enterprise Linux 8.0 AppStream (dvd) 21 MB/s | 5.3 MB 00:00
Red Hat Enterprise Linux 8.0 BaseOS (dvd) 23 MB/s | 2.2 MB 00:00
Last metadata expiration check: 0:00:01 ago on Mon 13 Apr 2020 10:55:47 AM EDT.
Package tmux-2.7-1.el8.x86_64 is already installed.
Dependencies resolved.
Nothing to do.
Complete!
```

2. Launch tmux

Listing 2. Using tmux

[student@workstation ~]\$ tmux

					root@workstati	on:~			-	•	×
File	Edit	View	Search	Terminal	Help						
[stu	dent@	works	tation	~]\$							
[0]	0:bas	h*			"wor	kstation.la	b.examp"	10:57	13-4	\pr-	20

Figure 2. tmux Terminal Window

3. Placing **tmux** Window in Background (*CTRL+B+D*)

Listing 3. Detaching a **tmux** Session

[student@workstation ~]\$ tmux
[detached (from session 0)]

4. Re-attaching to tmux

Listing 4. Source Description

```
[student@workstation ~]$ tmux list-sessions
0: 1 windows (created Mon Apr 13 11:01:53 2020) [80x23]
[student@workstation ~]$ tmux attach-session -t 0
```

Naming or Renaming Sessions

It is possible that once you are in a **tmux** session to rename the session.

Listing 5. Renaming a TMUX Session

[student@workstation ~]\$ tmux rename-session --help usage: rename-session [-t target-session] new-name

[student@workstation ~]\$ tmux rename-session -t 0 travis-demo

		root@workstation:~	-		×
File Edit View Search	Terminal Help				
capture-pane	list-keys	select-window			
choose-client	list-panes	send-keys			
choose-session	list-sessions	send-prefix			
choose-window	list-windows	server-info			
clear-history	load-buffer	set-buffer			
clock-mode	lock-client	set-environment			
command-prompt	lock-server	set-option			
confirm-before	lock-session	set-window-option			
copy-buffer	move-window	show-buffer			
copy-mode	new-session	show-environment			
delete-buffer	new-window	show-messages			
detach-client	pext-layout	show-options			
display-message	next-window	show-window-options			
display-panes 🛛 🦯	paste-buffer	source-file			
down-pane 🦊	pipe-pane	split-window			
find-window 🦊	previous-layout	start-server			
has-session	previous-window	suspend-client			
if-shell 🦯	refresh-client	swap-pane			
join-pane 🦯	rename-session	swap-window			
[student kstation	~]\$ tmux rename-sess	ionhelp			
usage: r sessio	<pre>n [-t target-session]</pre>	new-name			
[studentrkstation	~]\$ tmux rename-sess	ion -t 0 travis-demo			
[studen 🙋 workstation	~]\$				
[travis-de0:bash*		"workstation.lab.examp" 11:10	13-4	\pr-	20

Figure 3. Renamed tmux Session

[student@workstation ~]\$ tmux new-session
[detached (from session travis-demo)]

[student@workstation ~]\$ tmux list-sessions travis-demo: 1 windows (created Mon Apr 13 11:08:28 2020) [98x23]

[student@workstation ~]\$ tmux attach-session -t travis-demo

1.3. NoMachine

NoMachine can be used to remote control graphical workstations. There is a client/server model that is setup and NoMachine makes the initial connection on port 4000 using SSL. NoMachine can be found on multiple platforms using the following URL: https://www.nomachine.com/. i

Note Header

Encryption for NoMachine: https://www.nomachine.com/AR10K00705

SSL Client Authentication: https://www.nomachine.com/AR10M00866

NoMachine KB Articles for SSL: https://www.nomachine.com/articles?keys=ssl&an=ssl&form_build_id=form-111ef1c9e4bf56fd314eadead0c410b2&field_appliesto_value_many_to_one=All&custom_filter=any

2. Installing or Upgrading to Red Hat Enterprise Linux 8

2.1. Installing Red Hat Enterprise Linux 8

Notes on Installation

- BaseOS and Appstream Repositories are both required for installation of a RHEL8 system
 - Binary DVD contains both BaseOS and Appstream channels (approx. 7.5GB)
- RHEL8 Supports package Modularity
 - Modules are a set of RPMs
 - Application Streams: https://developers.redhat.com/blog/2018/11/15/rhel8-introducing-appstreams/
- System Purpose, System Roles, and Usage are all part of the new installation process in Anaconda



Red Hat Service Level Agreement

It is extremely important to note that the Service Level Agreement can cause issues with Red Hat Entitlements if not properly selected.

2.2. Upgrading Servers to Red Hat Enterprise Linux 8

2.3. FirewallD Service Definitions

FirewallD was introduced in RHEL7 as part of the SystemD transition and a new way to manage firewalls without using the underlying firewall implementation (**iptables**). FirewallD with **firewall-cmd** continues to be used in RHEL8 as the preferred method of managing and maintaining firewall rules.

FirewallD Resources



https://firewalld.org/

https://www.liquidweb.com/kb/an-introduction-to-firewalld/

https://cheatography.com/mikael-leberre/cheat-sheets/firewall-cmd/

2.3.1. The firewall-cmd Utility

The **firewall-cmd** utility is the primary method to manage and interact with firewall rules on RHEL7/8 systems. The **firewall-cmd** utility supports BASH completion and allows firewall rules to be added based on defined services or by specifying ports/protocols.



[root@servera ~]# firewall-cmd --add-port=80/tcp success [root@servera ~]# firewall-cmd --list-all public (active) target: default icmp-block-inversion: no interfaces: enpls0 sources: services: cockpit dhcpv6-client ssh ports: 80/tcp protocols: masquerade: no forward-ports: source-ports: icmp-blocks: rich rules: [root@servera ~]# firewall-cmd --remove-port=80/tcp success

Listing 7. Allowing HTTP through the Firewall by Service

```
[root@servera ~]# firewall-cmd --add-service=
Display all 154 possibilities? (y or n)
[root@servera ~]# firewall-cmd --add-service=http
success
[root@servera ~]# firewall-cmd --list-all
public (active)
 target: default
 icmp-block-inversion: no
 interfaces: enp1s0
 sources:
 services: cockpit dhcpv6-client http ssh
 ports:
 protocols:
 masquerade: no
 forward-ports:
  source-ports:
 icmp-blocks:
  rich rules:
```

firewall-cmd Usage Warning



The **firewall-cmd** utility can be used to make changes to the running firewall as shown in the examples above. This does not make changes to the firewall config file. In order to make the changes to the configuration file, it is necessary to use the **--permanent** options to have the changes written to a file.

When using **--permanent**, and you are not making changes to the current firewall runtime, it is also necessary to use: **firewall-cmd --reload** to reload or load new firewall rules from the firewall configuration file.



Important Header

It is important to note that presently the Red Hat Web Console (cockpit) only supports management of **firewalld** using defined services.

2.3.2. FirewallD Files and Locations

FirewallD has a few locations for files both for configuration and usage. As with most configuration files on a Linux system, those rely in *letcl*. The user configurable files for FirewallD also reside in *letcl*. Default configuration files for services, zones, and other FirewallD functionality resides in *lusr/lib/firewalld*. This location contains all defined services files and default configuration files for FirewallD and used by the *firewall-cmd* utility.





<pre>[root@servera ~]# tree /usr/lib/firewalld/ /usr/lib/firewalld/</pre>
├─── helpers
amanda.xml
ftp.xml
h323.xml
irc.xml
netbios-ns.xml
⊢─── pptp.xml
proto-gre.xml
zolies
drop.xml
external.xml
home.xml
internal.xml
├─── public.xml
trusted.xml
work.xml
5 directories, 222 files

2.3.3. Defining a Custom Service File

It is possible to define custom **firewalld** service files. These files can be used for your environments for custom applications or custom firewall rules. These service files can also be checked into a version control system such as **git**.



Creating a Custom Service File using Existing File as Base

It is easiest to take an existing service file, copy it to the *letc/firewalld/services* directory, rename and edit the file.

1. Copy existing FirewallD service file to /etc/firewalld/services

Listing 10. Using SSH Service File as a Starting Point

[root@servera ~]# cp /usr/lib/firewalld/services/ssh.xml /etc/firewalld/services/custom_ssh.xml

2. Edit the file and specify the new options

Listing 11. Edit the Custom SSH Service Definition

3. Listing FirewallD Services to Verify New Service

Listing 12. Getting Listing of FirewallD Services

[root@servera ~]# firewall-cmd --get-services | grep custom_ssh



FirewallD Delays in Discovering a Service

Depending on system speed and refreshing of FirewallD daemon, the new service might not be picked up immediately. It will be discovered or if you are in a hurry, you can run **firewall-cmd --reload** command to immediately have the service discovered and available.

2.3.4. FirewallD Configuration Files

As stated above, the main FirewallD configuration files are located in *letc/firewalld*. To specifically change or view the configuration file on the system, you generally want to look at the firewalls in the **Firewall Zone**. This is found by opening the corresponding zone file in *letc/firewalld/zones* directory.

1. Viewing Firewall configuration for the Public Zone.

Listing 13. FirewallD Configuration for Public

2. Adding a custom service

Listing 14. Adding our new Service

[root@servera ~]# firewall-cmd --add-service=custom_ssh --permanent
success

3. Verifying Firewall configuration for the Public Zone.

Listing 15. FirewallD Configuration for Public with Custom Service

3. Provisioning and Configuring Servers

3.1. Performing Package Management using Yum

3.2. Administering Servers with Cockpit

3.2.1. Extending Cockpit

3.2.2. Installing Additional Cockpit Packages

Listing 16. Install all Base Cockpit Packages

[root@ser	[root@servera ~]# yum install cockpit*											
			[Dashboard - serve	ra.lab.example.c	om - Mozilla F	irefox			-		×
🐣 Dashb	Servera.lab.e × +											
	C' û oard - serv	era.lab.exam	A https://ser	vera:9090/dashbo	ard		⊘	☆	liiN		٢	≡
RED HAT	ENTERPRI	SE LINUX							🔓 Privileg		🛓 ro	oot ~
	CPU	Memory	Network	Disk I/O								-
_	100%											
	80%											
$\langle \rangle$	60%											
	40%											
	20%											
	0%											
	0%		12:30	12	:31	12:32		12:33			12:34	
Servers 🖌												
servera.lab.example.com Red Hat Enterprise Linux 8.0 (Ootpa)												
	_											

Figure 4. Cockpit Dashboard

Important Header



Using the installation method above will install all Cockpit packages and plugins that are available in currently subscribed channels. However, not all components will work until additional back-end components are installed. For example, the **Composer** cockpit plugins need composer and other items installed in order to be able to be fully utilized. This installation gives the **Cockpit Dashboard Plugin** which allows connecting to multiple Web Consoles.

3.2.2.1. Cockpit to Manage Multiple Systems

It is possible to use a single **cockpit** Interface to manage multiple servers.

1. Ensure cockpit socket/service is running and configured on all systems.

Listing 17. Test and Enable Cockpit

```
[student@workstation ~]$ ssh root@servera
Activate the web console with: systemctl enable --now cockpit.socket
[root@servera ~]# systemctl enable --now cockpit.socket
Created symlink /etc/systemd/system/sockets.target.wants/cockpit.socket → /usr/lib/systemd/system/cockpit.socket.
```

2. Connect to the Cockpit Web Console



Figure 5. Red Hat Web Console (cockpit)

- 3. Add another Web Console from the Cockpit Dashboard
 - a. Navigate to the Dashboard
 - b. Click the "+" and complete the information

				Dashboard ·	- servera.lab.exa	mple.com - Mozi	lla Firefox			-		×
ڂ Dashb	ooard - servera	a.lab.e: 🗙	+									
\leftarrow	C 🕜	(https://s	ervera:9090/c	lashboard			🗵 ☆	lii\	•	٢	≣
RED HAT	RED HAT ENTERPRISE LINUX								🔓 Privileg	ged	🛓 ro	oot ~
	CPU	Memory	Network	Disk I/O								
	100%		Add Mach	ine to Dashbo	bard			×				
æ	80%		Address	serverb.lab.exa	ample.com							
$\langle \rangle \rangle$	60%		Color					•				
	40%							Cancel Add				
	20%											
	0% 12:31		12:3	2	12:33		12:34	12:35			V	
	Servers									•	+	
	🗐 s	ervera.lab.e	example.com					Red Hat Enter	prise Linux 8		tpa)	

Figure 6. Adding Additional machines

- 4. Verify the System Fingerprint and click **Connect** image::Chapter3-87075.png[title="Fingerprint Verification", align="center"]
- 5. It is now possible to switch systems from the System Drop-down menu

			System - serverb.lab.exan	nple.com - Mozilla Firefox			-	• ×	
🔸 Syster	System - serverb.lab.exar × +								
$\langle \boldsymbol{\leftarrow} \rangle$	୯ ଜ	🛈 🚯 https://ser	vera:9090/@serverb.lab.exa	ample.com/system	🗵 ☆	lii\ C			
RED H	ITERPRISE LINUX					🔓 Privileged	÷	root ~	
	serverb.lab.e.		Hardware Machine ID	Red Hat KVM 88394775ca9147439ae63aab	9c500045				
8	<u>م</u>		Operating System	Red Hat Enterprise Linux 8.0 (Ootpa)				
	Logs server	ra.lab.example.com	Secure Shell Keys	Show fingerprints					
	Networki Accounts	rb.lab.example.com	Domain	Join Domain					
5 L	Services		System Time Power Options	2020-04-13 12:38 🚯					
	Applications		Performance Profile	virtual-guest					
	Diagnostic Reports	%	of 1 CPU core						
	Kernel Dump	100							
	SELinux	50							
	Software Updates	0	12:34 12:35	12:36	12:37	12:38			

Figure 7. Switching Systems

3.3. Building System Images with Composer

3.3.1. Composer with Cockpit

In order to use the Lorax composer with **Cockpit** the cockpit packages and the composer packages must be installed.

1. Install the Lorax composer packages

Listing 18. Installing Composer Packages

[root@servera ~]# yum install lorax-composer composer-cli cockpit-composer

2. Enable the Composer Services

Listing 19. Enable Lorax Composer Socket

[root@servera ~]# systemctl enable --now lorax-composer.socket

3. Open the Red Hat Web Console and Navigate to Image Builder

Image Builder - servera.lab.example.com - Mozilla Firef	ox ×
Image Builder - servera.la × +	
$(\leftarrow) \rightarrow \mathbb{C}^{\prime}$ $\textcircled{Open a new tab (Ctrl+T)}$	… ▽ ☆
RED HAT ENTERPRISE LINUX	🔒 Privileged 💄 root 🗸
Filter by Name	Create Blueprint
example-attra	Edit Blueprint Create Image
A general purpose development image	Edit Blueprint Create Image
example-http-server	Edit Blueprint Create Image
An example http server with PHP and MySQL support.	

Figure 8. Image Builder - Web Console

3.4. Automating with RHEL System Roles

3.4.1. Ansible Basics

In order to use Ansible, the following items are required:

- ansible is installed on the Control Node
- An ansible.cfg file exists and points to a working inventory
- An inventory file exists containing managed nodes

3.4.1.1. Ansible Concepts and Architecture

Ansible Machine Types

- Control Node: Location where Ansible is installed and used to run playbooks and execute Ansible ad-hoc commands
- · Managed Host: Network device that is managed by an Ansible control node

The required Ansible components are the following:

- ansible.cfg: Ansible configuration file with directives on how Ansible should work
- · inventory: Listing and organization of ansible managed nodes/hosts
- module: Small piece of code to perform a variety of tasks (generally written in Python or Powershell)
- plugins: Code that can be used to extend and adapt Ansible to new uses. Capable of manipulating data and extracting information from output. (Used more in the next course DO447)
- Ansible Tower/API: Framework to control and manage Ansible at scale. Ansible Tower is not a core part of Ansible, but is an add-on product to more effectively utilize Ansible with teams.





3.4.1.2. Building an Ansible Inventory

An **inventory** file is the minimum needed item for Ansible to run. Inventory files provide a listing of hosts for Ansible to manage. There are multiple ways to define an inventory file.

3.4.1.2.1. Static Inventory

Static inventory files are generally defined in the INI format.

Listing 20. Simple Inventory

web1.example.com web2.example.com db1.example.com db2.example.com 192.0.2.42

[webservers]
web1.example.com
web2.example.com
192.0.2.42

3.4.1.3. Ansible Configuration Files

It is important to note, there can be multiple Ansible configuration files. Ansible will look for and process configuration files in order.

- 1. Location specified in ANSIBLE_CONFIG
- 2. The ansible.cfg file in the current working directory
- 3. The .ansible.cfg file in the user's home directory
- 4. The default /etc/ansible/ansible.cfg file



It is important to note that the first file in this order Ansible sees is what will be used. Therefore, if anything is setup for the **ansible.cfg** file in locations 1-3, it won't ever use the default file */etc/ansible/ansible.cfg*

Ansible Configuration file documentation: http://docs.ansible.com/ansible/intro configuration.html

Table 2. Ansible Settings

Setting	Command Line Option
inventory	Location of the Ansible inventory file.
remote_user	The remote user account used to establish connections to managed hosts.

Setting	Command Line Option
become	Enables or disables privilege escalation for operations on managed hosts.
become_method	Defines privilege escalation method on managed hosts.
become_user	User account to escalate privileges to on managed hosts.
become_ask_pass	Defines whether privilege escalation on managed hosts should prompt for password.

4. Adapting to Core System Changes

4.1. Displaying the Desktop with Wayland and X

4.1.1. Wayland and X11

Wayland replaced Xorg and the default X11 display server. Because of how Wayland compositing works and that Xorg no longer is providing an X11 server, there is no way to forward X11 traffic over an SSH connection. As mentioned earlier, NoMachine can provide a solution if the server has a graphical interface installed and NoMachine is installed and running on both the client and the Server.

For now, it is possible to switch the display server from Wayland to Xorg allowing the X11 server to function which would the provide the ability to forward an X11 connection.



1. Login and determine session ID using loginctl

Listing 21. Determining Session ID

```
[student@workstation ~]$ loginctl
SESSION UID USER SEAT TTY
4 0 root
6 1000 student seat0 tty2
c1 42 gdm seat0 tty1
```

2. Use loginctl

Listing 22. Determining Session Type

```
[student@workstation ~]$ loginctl show-session 6 -p Type
Type=wayland
```

3. Examine the GDM configuration

Listing 23. Examine GDM Configuration

```
[student@workstation ~]$ grep -i wayland /etc/gdm/custom.conf
#WaylandEnable=false
```

Disable Wayland and Switch to XOrg X11

You can disable Wayland and enabled Xorg X11 by modifying the line in the *letc/gdm/custom.conf* to just uncomment so it would read as:

```
[student@workstation ~]$ grep -i wayland /etc/gdm/custom.conf
WaylandEnable=false
```

Also, you will need to add **DefaultSession=gnome-xorg.desktop** to the **[daemon]** section of the file below the WaylandEnable=False.

[student@workstation ~]\$ cat /etc/gdm/custom.conf
GDM configuration storage

[daemon]
Uncoment the line below to force the login screen to use Xorg
WaylandEnable=false
DefaultSession=gnome-xorg.desktop

... output omitted ...

If any changes are made to the GDM configuration, the system must be rebooted in order to take effect.



Setting a Default Desktop Session

The **DefaultSession=gnome-xorg.desktop** section isn't specifically required for enabling X11 and disabling Wayland. It is one method that can be used to ensure that **xorg** is used as the default Gnome Session Manager.



Know Server Locations

It is important to know information about the server/desktop that is being used as the client system. The client system is responsible for having an XWindows server to allow X11 forwarding. It is possible to have a remote system with Wayland. As long as the client system supports X11 and has an X-Windows server, it is still possible to forward X11 connections over SSH.

Windows SSH Client with XWindows Support: https://mobaxterm.mobatek.net/

4.2. Managing User Authentication with Authselect

Objective

- Manage user authentication settings in PAM, NSS, and dconf using Authselect
- Explain the differences between Authselect and Authconfig

4.2.1. Introducing Authselect

Replaces **authconfig** making changes simpler for administrators. Used for configuration of passwords, certificates, smart cards, and fingerprints.

Features

- Adjusts PAM, NSS, and GNOME dconf settings.
- Ships with three ready-to-use profiles: **sssd**, **winbind**, and **nis**.
- pam_pwquality is enabled by default to enforce password quality restrictions on local users.

Comparing Authselect and Authconfig

- · Authselect uses profiles instead of modifying authentication config files directly
- Authselect only modifies files in /etc/nsswitch.conf, /etc/pam.d/, and /etc/dconf/db/distro.d/

Using Authselect

- Use the **authselect list** command to list the default and custom profiles.
- The default profiles are stored in /usr/share/authselect/ default.
- Use the authselect create-profile command to create new custom profiles.
- Custom profiles are stored in the *letc/authselect/custom/* directory.

Example 1. Authselect Demo

Listing 24. Getting Help for authselect

authoologt COMMAND	CUMMAND ADES
authselect COmmanD	CONFINITION
Available commands:	
- select	Select profile
- apply-changes	Regenerate configuration for currently selected command
- list	List available profiles
- show	Show profile information
- requirements	Print profile requirements
- current	Get identificator of currently selected profile
- check	Check if the current configuration is valid
- test	Print changes that would be otherwise written
- enable-feature	Enable feature in currently selected profile
- disable-feature	Disable feature in currently selected profile
- create-profile	Create new authselect profile
Common options:	
debug	Print error messages
trace	Print trace messages
warn	Print warning messages
Help options:	
-?,help	Show this for a command
usage	Show brief usage message for a command

Listing 25. Listing authselect Profiles

[root@servera ~]# authselect list

- nis Enable NIS for system authentication

- sssd Enable SSSD for system authentication (also for local users only)

- winbind Enable winbind for system authentication

Listing 26. Listing authselect Current Profile

[root@servera ~]# authselect current
Profile ID: sssd
Enabled features: None

Listing 27. Using authselect to Create a Profile

[root@servera ~]# authselect current
Profile ID: sssd
Enabled features: None
[root@servera ~]# authselect create-profile Travis_Custom \
> -b sssd --symlink-meta
New profile was created at /etc/authselect/custom/Travis_Custom
[root@servera ~]#

Listing 28. Modifying a Custom Profile with Additional Features





authselect keeps profiles of authentication configuration files and settings. This tool is based on profiles and allows custom profiles to be created. This should only be used for systems that are not part of a REALM with LDAP server or AD - **ipaclient** should be used or **realmd** for AD and LDAP/IPA respectively.

4.3. Configuring NTP with Chrony

Objective

• Maintain NTP time synchronization using Chrony, and configure the time zone with timedatectl.

4.3.1. Chrony Replaces ntpd

Chrony was implemented initially as part of RHEL7, but as of RHEL8, **chrony** is the only option for NTP services and **ntpd** is no longer available.

Using Chrony as NTP Replacement

- The /usr/share/doc/chrony/ntp2chrony.py script /etc/ntp.conf file to /etc/chrony.conf
- timedatectl can be used to set time aand ate and display.
- Use timedatectl set-timezone to define timezones.

Listing 29. Using the timedatectl Command



4.4. Managing Python Versions in Red Hat Enterprise Linux 8

5. Implementing Storage Using New Features

5.1. Managing Layered Storage with Stratis

Objective: Manage multiple storage layers using Stratis local storage management.

Packages

- stratis-cli: Provides the stratis command that sends user requests to stratisd service.
- stratisd: Provides the stratisd service implementing a D-Bus interface used managing and monitoring elements of Stratis.

5.1.1. Describing the Stratis Architecture

Traditional RHEL storage solutions allowed for scalable file systems, snapshots, RAID logical devices, multipathing, thin provisioning, caching, and a few other things like support for virtualization.

RHEL Storage Solutions

- dm: Device Mapper
- LVM: Logical Volume Manager
- XFS: 64-bit journaling file system (originally created by SGI)

Each of the RHEL storage solutions required management using layer specific commands requiring management of physical devices, volumes, and file systems as separate storage components.

RHEL 8 Storage Solution

• Stratis: Storage management solution that works with existing RHEL storage components.



Prior to RHEL 8, the ability to use *volume-managing file systems* was only considered development and not stable enough to become primary local storage for an enterprise-class system.

RHEL 8 introduced the **Stratis** storage management solution to work with the existing enterprise-class RHEL storage components. **Stratis** runs as a service that can manage pools of physical storage devices, create and manage volumes (transparently) for file systems being created. Since it uses existing storage drivers and tools, all advanced features available in **LVM**, **XFS**, and **device mapper** are used and supported by **Stratis**.

Volume-Managed File System

Volume-managed file systems are built inside of shared pools of disks using *thin provisioning*. **Stratis** file systems are built using a hidden LVM volume and **Stratis** manages the underlying volume and can expand it when needed.



Stratis file systems don't have fixed sized an no longer preallocate unused block space. Essentially, this can yield results similar to just-in-time configuration of file systems.

Stratis File System Facts

• In-use size of a file system is seen as amount of blocks in use by contained files.

- Space available to a file system is the amount of *unused* space in the pooled devices
- Multiple file systems can reside in same pool of disk devices (sharing available space)
- File systems can reserve pool space to guarantee availability when needed.



Stratis uses stored metadata to recognize managed pools, volumes and file systems. **Stratis** volumes should never be managed manually and only be managed using **Stratis** tools and commands. Manually managing/configuring/changing a **Stratis** file system could result in data loss and prevent the file system from being recognized.



Figure 10. Stratis Elements

Stratis can break up storage into pool tiers. The data tier and the cache tier. Typically, the cache tier should use block devices with higher (IOPS) such as SSD.

Stratis Pool Groups

· data tier: Focuses on flexibility and integrity

• cache tier: Focuses on improved performance

5.1.1.1. Describing the Simplified Storage Stack

It should be noted that **Anaconda** now leverages **Stratis** to simplify disk setup and management. Other products using **Stratis** include **Cockpit**, Red Hat Virtualization (RHV), and Red Hat Enterprise Linux Atomic Host.



Figure 11. Stratis in the Storage Management Stack



Stratis provides management to storage space and snapshots and allows easier integration with higherlevel management tools than using a CLI programmatically.

5.1.1.2. Describing Stratis Layers

Stratis has multiple layers to manage the storage system.

- · Backstore: Subsystem managing block devices
- Thinpool: Subsystem that manages pools
- **dm-thin**: The **Thinpool** subsystem uses the **dm-thin** device mapper driver to replace LVM for virtual volume sizing and management.

The Backstore subsystem has data tier maintaining on-disk metadata for block devices and detects/corrects data corruption. The

cache tier uses high-performance block devices acting as cache on top of the data tier.

The Thinpool subsystem manages thin-provisioned volumes associated with Stratis file systems.

The **dm-thin** driver creates volumes with a large virtual size, formatted with XFS, but a smaller physical size. As the physical size becomes full, **Stratis** automatically enlarges the disk.





5.1.1.3. Managing Thin-provisioned File Systems

In order to manage file systems using Stratis the stratis-cli and stratisd packages must be installed.

Example 2. Stratis Demo

1. Install Stratis

[root@servera ~]# yum install stratis-cli stratisd

2. Enable and Activate the Stratis service

```
[root@servera ~]# systemctl enable stratisd --now
[root@servera ~]# systemctl status stratisd
• stratisd.service - A daemon that manages a pool of block devices to create fl>
Loaded: loaded (/usr/lib/system/system/stratisd.service; enabled; vendor pr>
Active: active (running) since Fri 2019-06-21 14:49:32 EDT; 16s ago
Docs: man:stratisd(8)
Main PID: 23003 (stratisd)
Tasks: 1 (limit: 11405)
Memory: 884.0K
CGroup: /system.slice/stratisd.service
____23003 /usr/libexec/stratisd --debug
```

Using and Managing Stratis

· Creating Pools of one or more block devices

[root@servera ~]# stratis pool create DemoStorage /dev/vdb

· List Stratis Pools (stratis pool list)

```
[root@servera ~]# stratis pool listNameTotal Physical SizeDemoStorage5 GiB5 Z MiB
```

Add Additional Disks/Block Devices (stratis pool add-data)

[root@servera ~]# stratis pool add-data DemoStorage /dev/vdc

· View block devices in a pool (stratis blockdev list)

```
[root@servera *]# stratis blockdev list DemoStoragePool NameDevice NodePhysical SizeStateTierDemoStorage/dev/vdb5 GiBIn-useDataDemoStorage/dev/vdc5 GiBIn-useData
```

• Create a Filesystem (stratis filesystem create)

[root@servera ~]# stratis filesystem create DemoStorage DemoStorageFS1

· Create File system snapshots (stratis filesystem snapshot)

[root@servera ~]# stratis filesystem snapshot DemoStorage DemoStorageFS1 DS_FS_Snap1

· Listing Available Filesystems (stratis filesystem list)

```
[root@server ~]# stratis filesystem listPool NameNameUsedCreatedDeviceUUIDDemoStorageDemoStorageFS1546 MiBJun 21 2019 15:12/stratis/DemoStorage/DemoStorage/DemoStorageFS10155cbeda6cd4735ad540f394e74a8d3DemoStorageDS_FS_Snap1546 MiBJun 21 2019 15:14/stratis/DemoStorage/DS_FS_Snap10ad900732d8f4daba22037c0f4f1e8b0
```

Mounting Stratis Filesystems

It is best to mount the filesystem using the UUID in *letc/fstab*. Additionally, it is extremely important to use the x-systemd.requires=stratisd.service mount option as it delays the mount attempt until systemd starts stratisd.service.

· Obtaining UUID can be done with Isblk or with the stratis filesystem list

```
[root@servera ~]# lsblk --output=UUID /stratis/DemoStorage/DemoStorageFS1
UUID
0155cbed-a6cd-4735-ad54-0f394e74a8d3
```

• Updating /etc/fstab

[root@servera ~]# mkdir /MyDemoStorage ①

[root@servera ~]# echo "UUID=0155cbed-a6cd-4735-ad54-0f394e74a8d3 /MyDemoStorage xfs defaults,x-systemd.requires=stratisd.service 0
0" >> /etc/fstab ②

[root@servera ~]# mount -a ③

- 1 Create a mountpoint
- 2 Update /etc/fstab with the newly defined storage, mountpoint and parameters
- 3 Test the changes to *letc/fstab* before rebooting the system using mount -a option.

Stratis Help

The **Stratis** CLI has multiple ways of obtaining help. The traditional **man** page provides documentation and examples on using **Stratis**. In addition to the **man** pages, the **stratis** CLI can use the *--help* option at multiple levels of the command to obtain detailed help information.

Listing 30. Stratis Man Page

[root@servera stratis-cli]# man stratis

... omitted lines ...

EXAMPLES

Example 1. Creating a Stratis pool

stratis pool create mypool /dev/sdb /dev/sdc

Example 2. Creating a filesystem from a pool

stratis filesystem create mypool data1

... Output Omitted ...

Listing 31. Stratis CLI Help

```
[root@servera stratis-cli]# stratis --help ①
usage: stratis [-h] [--version] [--propagate]
               {pool,blockdev,filesystem,fs,daemon} ...
Stratis Storage Manager
... Output Omitted ...
subcommands:
 {pool,blockdev,filesystem,fs,daemon}
    pool
                       Perform General Pool Actions
                       Commands related to block devices that make up the
    blockdev
                        pool
                       Commands related to filesystems allocated from a pool
    filesystem (fs)
    daemon
                       Stratis daemon information
[root@servera stratis-cli]# stratis pool --help ②
usage: stratis pool [-h] {create,list,destroy,rename,add-data,add-cache} ...
optional arguments:
 -h, --help
                       show this help message and exit
subcommands:
 {create,list,destroy,rename,add-data,add-cache}
... Output Omitted ...
[root@servera stratis-cli]# stratis pool create --help ③
usage: stratis pool create [-h] [--redundancy {none}]
                           pool_name blockdevs [blockdevs ...]
... Output Omitted ...
```

1 Top-level stratis CLI command help

2 Pool subcommand-level CLI command help

3 Pool Create subcommand-level CLI command help


Failure to use the **x-systemd.requires=stratisd.service** mount option in *letc/fstab* for the Stratis file system will result in the machine booting to **emergency.target** on the next reboot.



Documentation: RHEL 8 - Managing File Systems - Chapter 11

https://access.redhat.com/documentation/en-us/red_hat_enterprise_linux/8/html-single/ managing_file_systems/index#managing-layered-local-storage-with-stratis_managing-file-systems

5.2. Compressing and Deduplicating Storage with VDO

Objective: Optimize storage space by using VDO to compress and deduplicate data on storage devices.

5.2.1. Describing Virtual Data Optimizer

RHEL 8 includes the Virtual Data Optimizer (VDO) driver that optimizes data footprint on block devices.

- VDO: Linux device mapper driver reducing disk space on block devices and minimizes replication of data saving disk space and can also increase data throughput. The VDO driver includes two (2) kernel modules to accomplish this task.
- · kvdo: Kernel module included as part of VDO to transparently control data compression.
- **uds**: Kernel module allowing deduplication of data.

The VDO layer is on top of existing storage devices (RAID or local disks). The block devices can also be encrypted devices. Storage layers such as LVM volumes and file systems are placed on top of a VDO device.



Figure 13. VDO-based Virtual Machines



1. Zero-Block Elimination: Filters out data blocks that contain only zero values and records information about those blocks in metadata only. Non-zero blocks are passed to next phase. This phase enables thin provisioning.

- Deduplication: Eliminates redundant data blocks. When multiple copies of same data are created, VDO detects duplicate
 data blocks and updates the metadata to use duplicate blocks as references to the original data block without creating
 redundant data blocks. This is accomplished using uds which is the universal deduplication service kernel module that
 checks redundancy of data through the metadata it maintains.
- 3. **Compression**: Last phase of VDO. The **kvdo** kernel module compresses data blocks using LZ4 compression and groups them on 4KB blocks.

5.2.2. Implementing Virtual Data Optimizer

Logical devices created with VDO are called **VDO Volumes**. A VDO Volume is similar to a disk partition in that it can be formatted and mounted like a regular filesystem. A VDO volume can also be used as an LVM physical volume.

In order to create a VDO volume, the block device must be specified along with the name of the logical device that VDO will present to the user. Optionally, the logical size of the VDO volume can be specified. The logical size of the VDO volume can be greater than the physical size of the actual block device.



VDO volumes are thinly provisioned allowing the user to only see the logical space in use and not the physical space available. If the VDO logical size isn't specified during creation, VDO assumes the actual physical size as the logical size of the volume.

If the VDO volume logical size is greater than the physical size, it should be proactively monitored using **vdostats --verbose**.

Example 3. VDO Examples

Install VDO packages (vdo kmod-kvdo)

[root@servera ~]# yum install vdo kmod-kvdo

Create a VDO Volume (vdo create)

[root@servera ~]# vdo create --name=MyVD0 --device=/dev/vdd --vdoLogicalSize=756 Creating VD0 MyVD0 Starting VD0 MyVD0 Starting compression on VD0 MyVD0 VD0 instance 0 volume is ready at /dev/mapper/MyVD0

• Format a VDO Volume

<pre>[root@servera ~]# mkfs.xfs /dev/mapper/MyVD0</pre>							
<pre>meta-data=/dev/mapper/MyVD0</pre>		isize= <mark>512</mark>	agcount=4, agsize=4915200 blks				
	=	sectsz=4096	attr <mark>=2</mark> , projid32bit=1				
	=	crc=1	finobt=1, sparse=1, rmapbt=0				
	=	reflink <mark>=1</mark>					
data	=	bsize =4096	blocks=19660800, imaxpct=25				
	=	sunit=0	swidth=0 blks				
naming	=version 2	bsize =4096	ascii-ci=0, ftype=1				
log	=internal log	bsize =4096	blocks=9600, version=2				
	=	sectsz=4096	<pre>sunit=1 blks, lazy-count=1</pre>				
realtime	e =none	extsz=4096	blocks=0, rtextents=0				

• Mounting a VDO Volume

Listing 32. Getting UUID for Block Device

[root@servera ~]# lsblk --output=UUID /dev/mapper/MyVDO
UUID
f9fa889a-1e7f-44c3-9edf-47ee986d73f0

Listing 33. Creating a Mount Point

[root@servera ~]# mkdir /MyVDOStorage

Listing 34. Modifying letclfstab

[root@servera ~]# echo "UUID=f9fa889a-1e7f-44c3-9edf-47ee986d73f0 /MyVDOStorage xfs defaults,x-systemd.requires=vdo.service 0 0" >>
/etc/fstab

Listing 35. Testing letc/fstab

[root@servera ~]# mount -a

• Analyzing a VDO Volume

```
[root@servera ~]# vdo status --name=MyVDO
VDO status:
 Date: '2019-06-21 17:11:14-04:00'
 Node: servera.lab.example.com
Kernel module:
... Output Omitted ...
    Logical size: 75G 1
    Logical threads: 1
    Max discard size: 4K
    Physical size: 5G (2)
    Physical threads: 1
    Slab size: 2G
    Storage device: /dev/vdd
    VDO statistics: ③
      /dev/mapper/MyVDO:
... Output Omitted ...
```

- 1 Shows the logical size that was defined for the VDO volume
- ② Shows the physical size of the disks backing the VDO volume
- 3 Begins the section providing VDO statistics

5.3. Administering NFS Enhancements

Objectives

- Identify NFS share information.
- Create a directory to use as a mount point.
- Mount an NFS share using the mount command or by configuring the /etc/fstab file.
- Unmount an NFS share using the umount command.
- · Configure an NFS client to use NFSv4 using the new nfsconf tool.

5.3.1. Mounting and Unmounting NFS Shares

NFS (Network File System) - default in RHEL 8 is now NFS v4.2. NFS exports shares (directories) which a client can mount using

- The mount command manually
- Automatically at boot with Ifstab
- · On demand using autofs service or systemd.automount



Only the major versions NFSv4 and NFSv3 are supported and NFSv2 support has been dropped.

NFSv4 and greater use TCP protocol



Since NFSv4 is being used, the **showmount** command can no longer be used as it can only reliably query NFSv3 servers.

Example 4. Demo of NFS

Listing 36. Prepare NFS Export

```
[root@serverb ~]# mkdir /ExportDemo
[root@serverb ~]# echo "/ExportDemo *(rw,no_root_squash)" > /etc/exports
[root@serverb ~]# exportfs -ra
```

Listing 37. Mount NFS Export

```
[root@servera ~]# mkdir /NFS_Share_Demo
[root@servera ~]# mount -t nfs serverb:/ExportDemo /NFS_Share_Demo
[root@servera ~]# touch /NFS_Share_Demo/testfile
```

5.3.2. The nfsconf Tool

RHEL 8 bring a new tool **nfsconf** to manage NFS client and server configuration files. The **nfsconf** tool manages the *letc/nfs.conf* file. Previous versions of this configuration files resided in *letc/sysconfig/nfs*, but are now deprecated.



The new nfsconf tool should be used to set, get, or unset NFS configuration parameters.

Example 5. Demo of nfsconf



nfsconf Tool Usage

nfsconf --set section key value nfsconf --get section key

Listing 38. Changing letc/nfs.conf Values with nfsconf

```
[root@servera ~]# nfsconf --set nfsd vers4.2 n
[root@servera ~]# grep vers4.2 /etc/nfs.conf
vers4.2 = n
# vers4.2=y
[root@servera ~]#
```

Listing 39. Retrieving letclnfs.conf Values with nfsconf

[root@servera ~]# nfsconf --get nfsd vers4.2
n

5.4. Automounting Network-Attached Storage

Objectives

- Describe the benefits of using the automounter.
- Automount NFS shares using direct and indirect maps, including wildcards.

5.4.1. Mounting NFS Shares with the Automounter

automounter is a service (autofs) that allows for on-demand mounting of NFS shares.

Automounter Benfits

- Standard users can use without needing **mount/umount** privileges
- Available to all users but subject to access control restrictions
- · Configured on client-side, no server-side configuration needed
- · Uses same options as mount command
- · Supports direct and indirect mount-point mapping
- · autofs creates/removes indirect mount mounts automatically
- · autofs is managed like other system services

5.4.1.1. Creating an Automount

1. Install **autofs**

root@servera ~]# yum install autofs

2. Add a master map file to letclauto.master.d called Auto_Mount_Demo.autofs

[root@servera ~]# vim /etc/auto.master.d/Auto_Mount_Demo.autofs

/nfs_shares /etc/autofs.demo ①

- 1 Top-level directory that will allow indirect mapped mounts
- 3. Create mapping Files

[root@servera ~]# vim /etc/autofs.demo

NFS_Work -rw,sync serverb:/ExportDemo ①

① Relative Name/Path for the Directory to be Mounted under the Infs_shares directory

4. Start/Enable the automounter service

```
[root@servera ~]# systemctl enable --now autofs
Created symlink /etc/systemd/system/multi-user.target.wants/autofs.service → /usr/lib/systemd/system/autofs.service.
```

5. Test that it is Working

```
[root@servera ~]# ls /nfs_shares/NFS_Work
testfile
```

5.4.1.2. Direct Maps

Direct maps are used to map NFS shares to an absolute path as a mount point.

See Student Guide

5.4.1.3. Indirect Wildcard Maps

When NFS servers export multiple subdirectories within a directory, automounter can be configured to allow indirect mapping using a single mapping entry containing wildcards.



The mount point (or key) is an asterisk character (*), and the subdirectory on the source location is an ampersand character (&).

6. Managing Containers with the New Runtime

6.1. Deploying Containers with the New Container Runtime

6.1.1. The Podman Container Engine

RHEL8 includes he **container-tools** package module. New engine is **podman** replaces **docker** and **moby**. It also contains new tools **buildah** to build container images and **skopeo** to manage images on registries like **runc**. The new toolset allows building/running containers without daemons.



Figure 14. Docker to RHEL8 Container Runtime

Container Runtime Toolset

- · Docker replaced with new container runtime
- New toolset supports OCI and reuse of third-party images
- Integrates with audit of Docker client-server model
- container-tools module provides new container runtime tools and engine.

41



Figure 15. New Container Runtime

Describing new Container Runtime Tool

- The **podman** engine is daemonless and supporting container execution.
- · podman syntax is similar to the docker command, supporting Dockerfile use
- Buildah builds container images, from scratch or a Dockerfile.
- · Copy and inspect container images in registries with Skopeo
- Skopeo supports Docker and private registries, the Atomic registry, and local directories, including those which use OCI



RHEL8 includes **Pacemaker** containers with **podman** as a tech preview. Pacemaker supports execution of the container across multiple hosts.

Example 6. Using Container Tools

Listing 40. Installation of Container Tools

[root@servera ~]# yum module install container-tools

Listing 41. Creating a Custom Container

[root@servera ~]# buildah from scratch
working-container

Listing 42. Naming and Inspecting a Custom Container

[root@servera ~]# buildah config --label name=My-Container working-container [root@servera ~]# buildah inspect working-container

Listing 43. Installing Packages on Working Container

<pre>[root@servera ~]# buildah mount working-container ①</pre>
[root@servera ~] # yumdownloaderdestdir=/tmp redhat-release-server ②
[root@servera ~] # rpm -ivhroot /var/lib/containers/storage/overlay/a6a136063f0ada2b1ed4b01eff9a04b4d6419ae828bc4b49e742bca594e08560/merged /tmp/redhat-release-8.0- 0.39.el8.x86_64.rpm ③
[root@servera ~] # cp /etc/yum.repos.d/rhel_dvd.repo /var/lib/containers/storage/overlay/a6a136063f0ada2b1ed4b01eff9a04b4d6419ae828bc4b49e742bca594e08560/merged/etc/yum.repos.d/ ④
[root@servera ~]# yum installinstallroot /var/lib/containers/storage/overlay/a6a136063f0ada2b1ed4b01eff9a04b4d6419ae828bc4b49e742bca594e08560/merged httpd ⑤
[root@servera ~]# echo "This is a custom webserver container for me" >> /var/lib/containers/storage/overlay/a6a136063f0ada2b1ed4b01eff9a04b4d6419ae828bc4b49e742bca594e08560/merged/var/www/html/index.html ⑥
[root@servera ~] # yum installinstallroot /var/lib/containers/storage/overlay/a6a136063f0ada2b1ed4b01eff9a04b4d6419ae828bc4b49e742bca594e08560/merged httpd-manual ⑦
[root@servera ~] # buildah configcmd "/usr/sbin/httpd -DFOREGROUND" working-container 🛞
<pre>[root@servera ~]# buildah configport 80/tcp working-container ③</pre>
[root@servera ~] # yum clean allinstallroot /var/lib/containers/storage/overlay/a6a136063f0ada2b1ed4b01eff9a04b4d6419ae828bc4b49e742bca594e08560/merged ⑩
[root@servera ~]# buildah unmount working-container
[root@servera ~]# buildah commit working-container my-container-image 🔞
<pre>[root@servera ~]# buildah images (3)</pre>
 Mount container image filesystem for modification Download Red Hat Release RPM for installation

- 3 Install Red Hat Release RPM
- ④ Create repository for container image so files can be installed
- (5) Install the HTTP package for a webserver
- 6 Create an index.html file for the webserver
- ⑦ Install the Apache manual for reference documentation
- (8) Configure webserver to run
- (9) Configure and open port 80 for the TCP protocol for the container
- 1 Clean up yum data to minimize required disk space
- 1 Unmount the container image filesystem
- 1 Commit the container image
- 1 List container images

Listing 44. Testing the Container Image

[root@servera ~]# podman run -d -p 8080:80 localhost/my-container-image					
[root@servera ~]# curl localhost:8080 This is a custom webserver container for me					
<pre>[root@servera ~]# curl http://localhost:8080/manual/</pre>					
Apache HTTP Server Version 2.4 Documentation - Apache HTTP Server Version 2.4 - Mozilla Firefox	×				
Apache HTTP Server Vers x +					

\leftarrow	\rightarrow C $$	i servera:8080/manual/			···· 🖸 🏠	li\ © ≡
	APACH	ΙE			Modules Directive	s <u>FAQ Glossary</u> <u>Sitemap</u>
KI	HTTP SERVER PRO	OJECT	Ap	ache HTTP Server Version 2	2.4	
A	pache > HTTP Server > Documenta	tion				
	Apache HTTP Server	Version 2.4 Documenta	tion			
				Av	ailable Languages: da de en es fr ja	ko pt-br tr zh-cn
				Google Search		
	Release Notes		Users' Guide		How-To / Tutorials	
	New features with Apache 2.3/2.4		Getting Started		Authentication and Authorization	
	New features with Apache 2.1/2.2		Binding to Addresses and Ports		Access Control	
	New features with Apache 2.0		Configuration Files		CGI: Dynamic Content	
	Upgrading to 2.4 from 2.2		Configuration Sections		.htaccess files	
	Apache License		Content Caching		Server Side Includes (SSI)	
	Deference Menuel		Content Negotiation		Per-user Web Directories (public_html)	
	Reference Manual		Dynamic Shared Objects (DSO)		Reverse proxy setup guide	
	Compiling and Installing		Environment Variables		HTTP/2 guide	
	Starting		Log Files		Blatform Chasifie Notes	
	Stonning or Restarting		and the second second second second		Platform Specific Notes	

Figure 16. Testing Container

Listing 45. Stopping and Cleanup of Image

<pre>[root@servera ~]# podman list ① CONTAINER ID IMAGE 9bd572633953 localhost/my-container-image:latest cranky_stonebraker</pre>	COMMAND /usr/sbin/httpd	CREATED 2 seconds ago	STATUS Up 1 second ago	PORTS 0.0.0.0:8080->80/tcp	NAMES				
<pre>[root@servera ~]# podman stop 9bd572633953 ② 9bd572633953276ac75417db3ac8e70875a0f2713e8cdfd32253fe343d06153d</pre>									
<pre>[root@servera ~]# podman stop -a ③</pre>	[root@servera ~]# podman stop -a ③								
[root@servera ~]# podman rm cranky_stonebraker ④									
[root@servera ~]# podman rm 9bd572633953276ac75417db3ac8e70875a0f2713e8cdfd32253fe343d06153d ⑤									
[root@servera ~]# podman rmi localhost/my-container-image ⑥									
[root@servera ~]# buildah delete working-container ⑦									

1 Listing Running Containers

② Stopping Single Container by ID

③ Stopping All Running Containers

4 Remove Container by Name

- 5 Remove Container by ID
- 6 Removing Container Image from Registry
- ⑦ Delete Working Container from System

7. Implementing Enhanced Networking Features

Objectives

- Explain the new nftables firewall back-end design, advantages and configuration.
- Explain how the NetworkManager service has become an integral and mandatory component in modern network management, able to configure complex, layered network interfaces and components.

7.1. Managing Server Firewalls in RHEL 8

Objective

• Explain **nftables** firewall

7.1.1. Introducing Nftables

nftables with firewalld

- firewalld uses nftables as its back-end.
- The nft command replaces the iptables, ip6tables, arptables, and ebtables commands.
- · firewalld is the recommended way to manage the firewall
- iptables commands are links to the xtables-nft-multi command, which accepts iptables syntax but creates nftables rules instead.



It is recommended to continue using **firewalld** to manage firewalls as the syntax remains the same even for the Rich Rules. It should be noted that **firewalld** can also be managed via **Cockpit**.

47



Figure 17. nftables Chains Process

Example 7. Firewall Management

Listing 46. Listing nftables
<pre>[root@servera ~]# nft list tables</pre>
Listing 47. firewalld and nftables
<pre>[root@servera ~]# nft list ruleset grep http</pre>
<pre>[root@servera ~]# firewall-cmdadd-service=http</pre>
<pre>[root@servera ~]# nft list ruleset grep http</pre>



It is possible to use **iptables** with RHEL8, but you must mask/disable the **firewalld** and **nftables** services as well as install the **iptables-services** package and enable it.

7.2. Configuring Server Networking with NetworkManager

- 8. Adapting to Virtualization Improvements
- 8.1. Configuring Virtual Machines

Appendix A: System Security Policy and Compliance

System security and compliance is a primary concern for people when thinking about the protection of systems and integrity of data. This set of hands-on procedures will focus on obtaining the content and necessary packages to perform a basic scan and remediate the system based on scan results. As part of the lab, you will be customizing your own SCAP content for a scan, view the results, and generate an Ansible playbook based on the failed results.

A.1. Customizing SCAP Content

Red Hat includes the SCAP Workbench application as a GUI application which allows scanning, customizing, and saving SCAP scans and results. The SCAP Workbench can be used to perform scans of remote systems over an SSH connection or it can be utilized to scan the local system. Since the SCAP Workbench is a GUI application, it must run on a system with X-Windows installed.

Servers to Configure

workstation

Packages to Install

scap-workbench



Since we are using an application on a VM that requires a GUI, we can use X11 forwarding with the SSH connection by specifying a **-X** on the command line.

Step 1 - SSH to Workstation

The fist step is to connect to the workstation and forward X11 traffic back to the local system.

Example 8. Connecting to the Workstation VM

Listing 48. Connecting to Workstation Using SSH

ssh root@workstation -X

Step 2 - Install Packages

After connecting to the Workstation VM, you will need to install the SCAP Workbench and its dependencies.

Example 9. Installing SCAP Workbench

Listing 49. Using Yum to Install SCAP Workbench

yum install scap-workbench

Loaded plugins: langpacks, search-disabled-repos Resolving Dependencies --> Running transaction check ---> Package scap-workbench.x86_64 0:1.1.6-1.el7 will be installed --> Processing Dependency: openscap-utils >= 1.2.0 for package: scap-workbench-1.1.6-1.el7.x86_64 --> Processing Dependency: scap-security-guide for package: scap-workbench-1.1.6-1.el7.x86_64 --> Running transaction check ---> Package openscap-utils.x86_64 0:1.2.16-8.el7_5 will be installed --> Processing Dependency: openscap-containers = 1.2.16-8.el7_5 for package: openscap-utils-1.2.16-8.el7_5.x86_64 ---> Package scap-security-guide.noarch 0:0.1.36-9.el7_5 will be installed --> Processing Dependency: openscap-scanner >= 1.2.5 for package: scap-security-guide-0.1.36-9.el7_5.noarch --> Running transaction check ---> Package openscap-containers.noarch 0:1.2.16-8.el7_5 will be installed ---> Package openscap-scanner.x86_64 0:1.2.16-8.el7_5 will be installed --> Finished Dependency Resolution Dependencies Resolved _____ Arch Version Repository Size Package ------Installing: scap-workbench x86_64 1.1.6-1.el7 rhel--server-dvd 1.8 M Installing for dependencies: openscap-containersnoarch1.2.16-8.el7_5rhel_updatesopenscap-scannerx86_641.2.16-8.el7_5rhel_updatesopenscap-utilsx86_641.2.16-8.el7_5rhel_updatesscap-security-guidenoarch0.1.36-9.el7_5rhel_updates **27** k **61** k 27 k **2.**6 M Transaction Summary _____ Install 1 Package (+4 Dependent packages) Total download size: 4.5 M Installed size: 64 M Is this ok [y/d/N]:



Some things might already be installed for you, if SCAP Workbench is already installed, please move on to the next step. Also note the dependencies for SCAP Workbench as they are automatically installed.

Step 3 - Launching SCAP Workbench

In order to run a scan or customize SCAP content, you will need to launch the SCAP Workbench application.



You **must** use SCAP Workbench from a GUI, so it will need to run either locally or through an SSH connection with X11 forwarded.

Example 10. Launching SCAP Workbench

Lis	sting 50. Using SCAP Workbench
‡ scap-workbench	
	SCAP Workbench _
File Help Title Customization Customization Profile Target Image: Customization Rules Customization	Customize Customize Customial Expand All
	Open SCAP Security Guide ×
SCA SECURITY O	 Also, these guides are a good starting point if you'd like to customize a policy or profile for your own needs. Select one of the default guides to load, or select Other SCAP Content option to load your own content. Select content to load: RHEL8 Close SCAP Workbench Load Content
Generate remediation role *	Dry run Fetch remote resources Remediate
Fig	jure 18. SCAP Workbench Startup

Step 4 - Creating Custom Content

Once SCAP Workbench has been launched, select the content to load. For this lab, we will be using the RHEL7 content.

Example 11. Creating Custom Content

- 1. For Select content to load: select "RHEL8", then click "Load Content"
- 2. Select the Profile you want to use to start customization



For this example, we will use the **OSPP - Protection Profile for General Purpose Operating Systems** Baseline

ssg-rhel8-ds.xml - SCAP Workbench						
<u>F</u> ile <u>H</u> elp						
	Cuide to the Course Configuration of Red Upt Entermaine Linux R					
Title	Guide to the Secure Configuration of Red Hat Enterprise Linux 8			_		
Customization	None selected			•		
Profile	OSPP - Protection Profile for General Purpose Operating Systems (187)	Custo	omize	2		
Target	Local Machine O Remote Machine (over SSH)					
Rules		Expa	nd al	II]		
				A		
Ensure Log	gs Sent To Remote Host					
Ensure cro	in Is Logging To Rsyslog					
Ensure au	ditd Collects Information on Kernel Module Loading and Unloading - modprobe					
Ensure au	ditd Collects Information on Kernel Module Loading - init_module					
Ensure au	ditd Collects Information on Kernel Module Unloading - delete_module					
Ensure au	ditd Collects Information on Kernel Module Loading - insmod					
Ensure au	ditd Collects Information on Kernel Module Unloading - rmmod					
Record Un	successul Permission Changes to Files - chmod					
Record Un	Record Unsuccessul Ownership Changes to Files - chown					
Record Un	Record Unsuccessul Permission Changes to Files - removexattr					
Record Un	Record Unauthorized Access Attempts to Files (unsuccessful) - ftruncate					
Record Un	Record Unsuccessul Ownership Changes to Files - Ichown					
Record Un	Record Unsuccessul Permission Changes to Files - fremovexattr					
Record Un	authorized Modification Attempts to Files - open O_TRUNC					
Record Un	successul Ownership Changes to Files - fchown					
Record Un	authorized Modification Attempts to Files - openat O_TRUNC					
Record Un	successul Permission Changes to Files - fsetxattr					
Record Un	successul Ownership Changes to Files - fchownat					
Record Un	authorized Creation Attempts to Files - open_by_handle_at O_CREAT					
Record Un	authorized Access Attempts to Files (unsuccessful) - creat					
Record Un	authorized Creation Attempts to Files - open O_CREAT			-		
4			Þ			
	0% (0 results, 187 rules selected)					
Generate reme	diation role • Dry run Fetch remote resources Remediate	<u>S</u> c	an			

Figure 19. SCAP Workbench OSPP Profile

3. Click "Customize" to create custom SCAP content based on the chosen profile, and give it a name.

		ssg-rhel8-ds.xml - SCAP Workbench _		×	
<u>F</u> ile	<u>H</u> elp				
Title		Guide to the Secure Configuration of Red Hat Enterprise Linux 8			
Customization None selected					
Profi	le	OSPP - Protection Profile for General Purpose Operating Systems (187)	ustomiz	e	
Targ	et	Local Machine O Remote Machine (over SSH)			
Rules	5	Custom Profile Name	xpand a	II	
•	Ensure Log	gs Sent To Remote Host		-	
►	Ensure cro	on Is Logging To Rsyslog			
►	Ensure aud	ditd Collects Information on Kernel Module Loading and Unloading - modprobe			
•	Ensu	Customize Profile ×			
•	Ensu				
	Enst Cho	oose the ID of your profile.			
	Enst Wa	arning: Choose it wisely. It cannot be changed later and may be required if you choose to use			
	Rect Con	Initialid line cools of various integrations of OpenSCAF.			
	Rec. The	e ID has to have a format of "xccdf_{reverse DNS}_profilest of the ID}.			
	For	example "xccdf_org.mycorporation_profile_server".			
•	Rect Nev	w Profile ID xccdf_org.ssgproject.content_profile_ospp_customized			
•	Reco				
•	Reco	OK Cancel			
►	Record Uns	successul Ownership Changes to Files - fchown			
•	Record Una	authorized Modification Attempts to Files - openat O_TRUNC			
•	Record Uns	successul Permission Changes to Files - fsetxattr			
•	Record Uns	successul Ownership Changes to Files - fchownat			
	Record Una	authorized Creation Attempts to Files - open_by_handle_at O_CREAT			
	Record Una	authorized Access Attempts to Files (unsuccessful) - creat			
	Record Una	authorized Creation Attempts to Files - open O_CREAT		•	
		0% (0 results 197 piles celected)	•		
		0% (0 results, 107 rules selected)			
Gene	erate remed	diation role • Dry run 🗌 Fetch remote resources 🗌 Remediate	Scar	1	
				:	

Figure 20. SCAP Custom OSPP Profile Creation



The name for this is: xccdf_org.ssgproject.content_profile_ospp_customized

4. Click "Deselect All" so that you can select the items you wish to include in your custom scan profile. **NOTE:** we are doing this to also limit it to a few checks for the example.



Figure 21. SCAP Custom Profile Selections

a

For this lab, we will be setting the minimum password length and PAM Password quality settings

5. Search for Password to set minimum password length and set the values in login.defs. Check Set Password Minimum Length in login.defs and click on the minimum password legnth and set the value to 18



Figure 22. SCAP Custom Profile Password Settings for Login.Defs

 Set password quality requirements with PAM. Search for the minlen and set it to 18. Also, place a checkbox in Set Password Quality Requirements with pam_quality. Then click "OK"



Figure 23. SCAP Custom Profile PAM Quality Requirements

7. At this point, we have taken the default settings from the OSPP profile with only the tailored pieces that we selected. The next step is to click "File \Rightarrow Save Customization Only" to save the custom content

	ssg-rhel8-ds.xm	- SCAP Workbench	_ 0 X			
<u>F</u> ile <u>H</u> elp						
Title	Guide to the Secure Configura	ation of Red Hat Enterprise Linux	(8			
Customization	(unsaved changes)					
Profile	OSPP - Protection Profile for General Purpos	e Operating Systems [CUSTOMIZED] (11)	Customize			
Tarmat	· · · · · ·					
Target	Local Machine	 Remote Machine (over SSH) 				
Rules			Expand all			
Set Passwo	ord Minimum Length in login.defs					
Set Passwo	ord Strength Minimum Different Characters					
Set Passwo	ord Strength Minimum Uppercase Characters					
Set Passwo	ord Minimum Length					
Set Passwo	ord Retry Prompts Permitted Per-Session					
Set Passwo	ord Strength Minimum Different Categories					
Set Passwo	ord Maximum Consecutive Repeating Charac	ters				
Set Passwo	ord Strength Minimum Special Characters					
Set Passwo	ord Strength Minimum Lowercase Characters					
Set Passwo	ord Strength Minimum Digit Characters					
Set Passwo	 Set Password to Maximum of Consecutive Repeating Characters from Same Character Class 					
	0% (U results,	11 rules selected)				
Generate remed	diation role *	y run 🗌 Fetch remote resources 🗌 Remediat	e <u>S</u> can			

Figure 24. SCAP Custom Profile Selected Settings View

		ssg-rhel8-ds.xml - SCAP Workbench	. • ×
<u>F</u> ile <u>H</u> elp			
Title	Guide to	to the Secure Configuration of Red Hat Enterprise Linux 8	
Customization	(unsaved o	changes)	-
Profile	OSPP - Pro	otection Profile for General Purpose Operating Systems [CUSTOMIZED] (11) 🔹 🔹 🗌	Customize
Target	Local Ma	lachine O Remote Machine (over SSH)	
Rules		Save Customization As ×	Expand all
 Set Pa Loo 	ok in:	💳 /root 🔹 🔾 🗘 🖓 🖽 🔳	
🕨 Set Pa	Computer	Name Size Type Date Modified	
Set Pa	root		
 Set Pa 			
Set Pa			
Set Pa		Name of Tailoring File	
Set Pa			
 Set Pa Set Pa 			
File	e <u>n</u> ame:	ssg-rhel8-ds-tailoring.xml	
File	es of type:	XCCDF Tailoring file (*.xml)	
		0% (0 results. 11 rules selected)	
Generate remed	liation role ·	Dry run Fetch remote resources Remediate	<u>S</u> can

Figure 25. SCAP Custom Profile Creation Saving

	ssg-rhel8-ds.xml - SCAP Workbench	_		×
<u>F</u> ile <u>H</u> elp				
Title	Guide to the Secure Configuration of Red Hat Enterprise Linux 8			
Customization	/root/ssg-rhel8-ds-tailoring.xml			•
Profile	OSPP - Protection Profile for General Purpose Operating Systems [CUSTOMIZED] (11)	Cust	omiz	e
Target	Local Machine O Remote Machine (over SSH)			
Rules		Expa	nd a	
 Set Passwi <	ord Minimum Length in login.defs ord Strength Minimum Different Characters ord Strength Minimum Uppercase Characters ord Minimum Length ord Retry Prompts Permitted Per-Session ord Strength Minimum Different Categories ord Maximum Consecutive Repeating Characters ord Strength Minimum Special Characters ord Strength Minimum Lowercase Characters ord Strength Minimum Digit Characters ord Strength Minimum Digit Characters ord to Maximum of Consecutive Repeating Characters from Same Character Class			
	0% (0 results, 11 rules selected)			
Generate remed	diation role * Dry run Fetch remote resources Remediate	<u>S</u> c	ar	ו

Figure 26. SCAP Custom Profile Final View

8. Copy the custom tailoring file to the server(s) being scanned. In this case, we will want to copy the file to servera

Listing 51. Copy custom content

```
[root@workstation ~]# scp ssg-rhel8-ds-tailoring.xml root@servera:
ssg-rhel8-ds-tailoring.xml 100% 28KB 10.7MB/s 00:00
[root@workstation ~]#
```

A.2. Running a SCAP Scan with Custom Content

Servers to Configure

servera

Packages to Install

- openscap-scanner
- scap-security-guide

Step 1 - SSH to serverc

The fist step is to connect to the server.

Example 12. Connecting to the servera VM

Listing 52. Connecting to servera Using SSH

ssh root@servera

Step 2 - Install packages on servera

The second step is to install software on the server.

Example 13. Install software on servera

Listing 53. Installing Software on servera

```
[root@servera ~]# yum install scap-security-guide
Last metadata expiration check: 0:47:44 ago on Thu 16 Apr 2020 08:26:15 AM EDT.
Dependencies resolved.
-----
Package Arch Version Repository
                                                               Size
-----
Installing:
scap-security-guide
              noarch 0.1.42-11.el8 rhel-8.0-for-x86_64-appstream-rpms 3.4 M
Installing dependencies:
openscap x86_64 1.3.0-7.el8 rhel-8.0-for-x86_64-appstream-rpms 3.3 M
openscap-scanner x86_64 1.3.0-7.el8 rhel-8.0-for-x86_64-appstream-rpms 66 k
xml-common noarch 0.6.3-50.el8 rhel-8.0-for-x86_64-baseos-rpms
                                                                39 k
Transaction Summary
______
Install 4 Packages
Total download size: 6.9 M
Installed size: 132 M
Is this ok [y/N]: y
... output omitted ...
 Verifying : openscap-1.3.0-7.el8.x86_64
Verifying : openscap-scanner-1.3.0-7.el8.x86_64
Verifying : scap-security-guide-0.1.42-11.el8.noarch
Verifying : xml-common-0.6.3-50.el8.noarch
                                                                 1/4
                                                                2/4
                                                               3/4
                                                                 4/4
Installed:
 scap-security-guide-0.1.42-11.el8.noarch openscap-1.3.0-7.el8.x86_64
 openscap-scanner-1.3.0-7.el8.x86_64 xml-common-0.6.3-50.el8.noarch
```

Learning about SCAP Commands

The SSG man page is a very good source of information for usage of the **oscap** tool as well as provides examples of how to use the SCAP SSG Guide profiles itself.

Complete!

# man scap-security-guide									
<pre>scap-security-guide(8) System Manager's Manual scap-security-guide(8)</pre>									
NAME SCAP Security Guide - Delivers security guidance, baselines, and asso- ciated validation mechanisms utilizing the Security Content Automation Protocol (SCAP).									
output omitted									
EXAMPLES To scan your system utilizing the OpenSCAP utility against the ospp- rhel7 profile:									
oscap xccdf evalprofile ospp-rhel7results /tmp/`hostname`-ssg- results.xmlreport /tmp/`hostname`-ssg-results.htmloval-results /usr/share/xml/scap/ssg/content/ssg-rhel7-xccdf.xml									

Listing 54. Looking at SCAP Security Guide (SSG) Man Page



# man	oscap						
OSCAP <mark>(</mark>	8) System Administration Utilities OSCAP(8)						
NAME	oscap - OpenSCAP command line tool						
SYNOPS	IS oscap [general-options] module operation [operation-options-and-argu- ments]						
DESCRIPTION oscap is Security Content Automation Protocol (SCAP) toolkit based on OpenSCAP library. It provides various functions for different SCAP specifications (modules).							
	OpenSCAP tool claims to provide capabilities of Authenticated Configu- ration Scanner and Authenticated Vulnerability Scanner as defined by The National Institute of Standards and Technology.						
ou	tput omitted						
EXAMPLES Evaluate XCCDF content using CPE dictionary and produce html report. In this case we use United States Government Configuration Baseline (USGCB) for Red Hat Enterprise Linux 5 Desktop.							
	<pre>oscap xccdf evalfetch-remote-resourcesoval-results \</pre>						
X	report usgcb-rhel5desktop.report.html \ results usgcb-rhel5desktop-xccdf.xml.result.xml \ cpe usgcb-rhel5desktop-cpe-dictionary.xml \ usgcb-rhel5desktop-xccdf.xml						

Step 3 - Running oscap scan

We will run the **oscap** utility to generate a report and a results file that can be sent back to the **workstation** system so that we can create an Ansible playbook for remediation and view the results of the report.



Be very careful about the name of the profile as this was selected during the creation of the custom profile/tailoring file portion when doing SCAP Workbench customizations.

Example 14. Scanning servera

Listina 56.	Usina os	cap and	the	tailoring	profile	to scan	servera
				··· · J	I		

```
# [root@servera ~]# oscap xccdf eval \
--profile xccdf_org.ssgproject.content_profile_ospp_customized \
--tailoring-file ssg-rhel8-ds-tailoring.xml \
--results custom_scan_results.xml \
/usr/share/xml/scap/ssg/content/ssg-rhel8-ds.xml
Title Set Password Minimum Length in login.defs
       xccdf_org.ssgproject.content_rule_accounts_password_minlen_login_defs
Rule
Ident CCE-80652-1
Result fail
Title Set Password Strength Minimum Different Characters
Rule xccdf_org.ssgproject.content_rule_accounts_password_pam_difok
Ident CCE-80654-7
Result fail
       Set Password Strength Minimum Uppercase Characters
Title
Rule
       xccdf_org.ssgproject.content_rule_accounts_password_pam_ucredit
Ident
       CCE-80665-3
Result fail
Title Set Password Minimum Length
       xccdf_org.ssgproject.content_rule_accounts_password_pam_minlen
Rule
Ident
       CCE-80656-2
Result fail
Title Set Password Retry Prompts Permitted Per-Session
Rule
       xccdf_org.ssgproject.content_rule_accounts_password_pam_retry
Ident CCE-80664-6
Result fail
Title Set Password Strength Minimum Different Categories
Rule
       xccdf_org.ssgproject.content_rule_accounts_password_pam_minclass
Result fail
Title Set Password Maximum Consecutive Repeating Characters
Rule
       xccdf_org.ssgproject.content_rule_accounts_password_pam_maxrepeat
Result fail
Title Set Password Strength Minimum Special Characters
Rule
       xccdf_org.ssgproject.content_rule_accounts_password_pam_ocredit
Ident
       CCE-80663-8
Result fail
Title Set Password Strength Minimum Lowercase Characters
Rule
       xccdf_org.ssgproject.content_rule_accounts_password_pam_lcredit
Ident CCE-80655-4
Result fail
Title Set Password Strength Minimum Digit Characters
Rule
       xccdf_org.ssgproject.content_rule_accounts_password_pam_dcredit
Ident CCE-80653-9
Result fail
       Set Password to Maximum of Consecutive Repeating Characters from Same Character Class
Title
Rule
       xccdf_org.ssgproject.content_rule_accounts_password_pam_maxclassrepeat
Result fail
```

Getting Custom Profile Name from Tailoring File

 \mathbf{O}

If you need to locate the profile used for the custom scanning content from the tailoring file, you can search for it with **grep**.

Step 4 - Creating a Results Report

You can create a results report file from the results file so you have a nice HTML file that is easy to ready with the results from the SCAP scan.

Example 15. Creating a SCAP Report from a Results File

Listing 57. Generating a Report

[root@servera ~]# oscap xccdf generate report \
custom_scan_results.xml > Custom_Scan_Report.html

Combining Steps 3 & 4

It is possible to perform a custom content scan which will generate the results file and the report for transfer back to the workstation for review.

Need to Specify

- · --results
- --report

[root@servera ~]# oscap xccdf eval \ --profile xccdf_org.ssgproject.content_profile_ospp_customized \ --tailoring-file ssg-rhel8-ds-tailoring.xml \ --results custom_scan_results_2.xml \ --report Custom_Scan_Report_2.html \ /usr/share/xml/scap/ssg/content/ssg-rhel8-ds.xml Title Set Password Minimum Length in login.defs Rule xccdf_org.ssgproject.content_rule_accounts_password_minlen_login_defs Ident CCE-80652-1 Result fail Title Set Password Strength Minimum Different Characters Rule xccdf_org.ssgproject.content_rule_accounts_password_pam_difok Ident CCE-80654-7 Result fail Title Set Password Strength Minimum Uppercase Characters Rule xccdf_org.ssgproject.content_rule_accounts_password_pam_ucredit Ident CCE-80665-3 Result fail Title Set Password Minimum Length Rule xccdf_org.ssgproject.content_rule_accounts_password_pam_minlen Ident CCE-80656-2 Result fail Title Set Password Retry Prompts Permitted Per-Session Rule xccdf_org.ssgproject.content_rule_accounts_password_pam_retry Ident CCE-80664-6 Result fail Title Set Password Strength Minimum Different Categories Rule xccdf_org.ssgproject.content_rule_accounts_password_pam_minclass Result fail Title Set Password Maximum Consecutive Repeating Characters Rule xccdf_org.ssgproject.content_rule_accounts_password_pam_maxrepeat Result fail Title Set Password Strength Minimum Special Characters Rule xccdf_org.ssgproject.content_rule_accounts_password_pam_ocredit Ident CCE-80663-8 Result fail Title Set Password Strength Minimum Lowercase Characters Rule xccdf_org.ssgproject.content_rule_accounts_password_pam_lcredit Ident CCE-80655-4 Result fail Title Set Password Strength Minimum Digit Characters Rule xccdf_org.ssgproject.content_rule_accounts_password_pam_dcredit Ident CCE-80653-9 Result fail Set Password to Maximum of Consecutive Repeating Characters from Same Character Class Title Rule xccdf_org.ssgproject.content_rule_accounts_password_pam_maxclassrepeat Result fail

Listing 58. Creating a Results File and Report During Custom Content Scan

Step 5 - Transferring Results File and Report to Workstation
After you have the results files and the report, you should transfer it to your graphical workstation (**workstation**) for further analysis.

Example 16. Transferring Results

	Listing 59. Tran	sferring the	e Results and Report Files
[root@servera ~]# scp *.xml *.html	root@workstation:		
The authenticity of host 'workstatic	on (<no for="" hostip="" pr<="" td=""><td>oxy command></td><td>)' can't be established.</td></no>	oxy command>)' can't be established.
ECDSA key fingerprint is SHA256:pOQ	10JmyF2PFI+jxyFoOSCfi	+1oWNsUruy2D	ZNjg+N0.
Are you sure you want to continue co	onnecting (yes/no)? y	es	
Narning: Permanently added 'workstat	tion' (ECDSA) to the	list of know	n hosts.
root@workstation's password:			
custom_scan_results_2.xml	100% 4086K	3 32.4MB/s	00:00
custom_scan_results.xml	100% 4086K	3 60 .0MB/s	00:00
sg-rhel8-ds-tailoring.xml	100% 28K	3 16. 2MB/s	00:00
Custom Scan Report 2.html	100% 332K	3 44. 3MB/s	00:00
Custom Scan Report.html	100% 332K	3 37.6MB/s	00:00
root@servera ~]#			

Step 6 - Viewing the SCAP scan report

After you have transferred the results file to **workstation** you can open the HTML report in a web browser. In this case we will use *firefox* to open the file.

Example 17. Viewing the SCAP Report

V. 0000 0000	tastros V 💭 Firafay Drivery Nation 🛛 V 上			
rg.open-scap_	testres × Firefox Privacy Notice - 1 × +			has
, C W	The:///roou/Custom_Scan_Report.nt	n	V W	IIIN
Evalua	tion Characteristics			
Evaluation target	servera.lab.example.com	CPE Platforms • cpe:/o:redhat:enterprise_linux:8	Addresses	
Benchmark URL	/usr/share/xml/scap/ssg/content/ssg-rhel8-ds.xml		IPv4 172.25.250.10 IPv6 0:0:0:0:0:0:0:1 IPv6 fe80:0:0:0:e6c5:468e:edb6:9b52	
Benchmark ID	xccdf_org.ssgproject.content_benchmark_RHEL-8		• MAC 00:00:00:00:00 • MAC 52:54:00:00:FA:0A	
Benchmark version	0.1.42			
Profile ID	xccdf_org.ssgproject.content_profile_ospp_customized	ł		
Started at	2020-04-16T09:21:32			
Finished at	2020-04-16T09:21:33			
Performed by	root			
Test system	cpe:/a:redhat:openscap:1.3.0			
Compli The target s Rule rest	ance and Scoring system did not satisfy the conditions of 11 rules! Ple ults	ease review rule results and consider ap	olying remediation.	
		11 failed		
Severity	of failed rules			
Saara		11 medium		
Score			Percent	
Scoring syst	tem Score	Maximum	1 crocite	



Firefox may not open the file based on SELinux context triggers. In order to get around this you can use the command prompt and do **setenforce 0** to allow you to open the report.

A.3. Creating an Ansible Remediation Playbook Based on SCAP Scan Results

The OpenSCAP project and content created by Red Hat can automatically remediate findings from OpenSCAP scans. The findings can be remediated in many ways (**BASH**, **Ansible**, etc.). While things are mostly complete, there are some automated remediations that have not yet been developed.



There are multiple automatic remediation methods developed, but at this time, there isn't a script to fix everything.

Servers to Configure

• servera



We will continue to use **workstation** as our master SCAP system as it should have Ansible and SCAP Workbench installed.

Step 1 - Creating an Ansible Playbook from Results

The first step will be to generate an Ansible playbook from the SCAP scan results for system remediation.

Example 18. Generating Ansible Playbook

Listing 61. Ansible Playbook Generation [root@workstation ~]# oscap xccdf generate fix \ --profile xccdf_org.ssgproject.content_profile_ospp_customized \ --tailoring-file ssg-rhel8-ds-tailoring.xml \ --fix-type ansible \ --result-id "" \ custom_scan_results.xml > Custom_Scan_Fix.yml Viewing Remediation Playbook It is also a good idea to view the created playbook for the system prior to running it. [root@workstation ~]# cat Custom_Scan_Fix.yml # # Ansible remediation role for the results of evaluation of profile xccdf_org.ssgproject.content_profile_ospp_customized # XCCDF Version: unknown # # Evaluation Start Time: 2020-04-16T09:21:32 # Evaluation End Time: 2020-04-16T09:21:33 # This file was generated by OpenSCAP 1.3.0 using: # \$ oscap xccdf generate fix --result-id xccdf_org.openscap_testresult_xccdf_org.ssgproject.content_profile_ospp_customized --template urn:xccdf:fix:script:ansible xccdf-results.xml

This script is generated from the results of a profile evaluation.

```
# It attempts to remediate all issues from the selected rules that failed the test.
#
# How to apply this remediation role:
# $ ansible-playbook -i "localhost," -c local playbook.yml
# $ ansible-playbook -i "192.168.1.155," playbook.yml
# $ ansible-playbook -i inventory.ini playbook.yml
- hosts: all
   vars:
     var_accounts_password_minlen_login_defs: !!str 18
     var_password_pam_difok: !!str 8
     var_password_pam_ucredit: !!str -1
     var_password_pam_minlen: !!str 18
     var_password_pam_retry: !!str 3
     var_password_pam_minclass: !!str 3
     var_password_pam_maxrepeat: !!str 3
     var_password_pam_ocredit: !!str -1
     var_password_pam_lcredit: !!str -1
     var_password_pam_dcredit: !!str -1
     var_password_pam_maxclassrepeat: !!str 4
   tasks:
    - name: "Set Password Minimum Length in login.defs"
     lineinfile:
       dest: /etc/login.defs
       regexp: "^PASS_MIN_LEN *[0-9]*"
       state: present
       line: "PASS_MIN_LEN
                                 {{ var_accounts_password_minlen_login_defs }}"
     tags:
        - accounts_password_minlen_login_defs
       - medium_severity

    restrict_strategy

        - low_complexity
        - low_disruption
        - CCE-80652-1
       - NIST-800-53-IA-5(f)
       - NIST-800-53-IA-5(1)(a)
        - NIST-800-171-3.5.7
        - CJIS-5.6.2.1
... Output Omitted ...
    - name: Ensure PAM variable maxclassrepeat is set accordingly
     lineinfile:
       create: yes
       dest: "/etc/security/pwquality.conf"
       regexp: '^#?\s*maxclassrepeat'
       line: "maxclassrepeat = {{ var_password_pam_maxclassrepeat }}"
      tags:
       - accounts_password_pam_maxclassrepeat
        - medium_severity
        - restrict_strategy
        - low_complexity
        - low_disruption
        - NIST-800-53-IA-5
        - NIST-800-53-IA-5(c)
```



Before we can do the next steps, we will download an Ansible config file and an inventory file so we can properly run the playbook.

Listing 62. Error Output Message
<pre>[root@workstation ~]# ansible-playbook Custom_Scan_Fix.yml [WARNING]: provided hosts list is empty, only localhost is available. Note that the implicit localhost does not match 'all'</pre>
PLAY [all] ***********************************
PLAY RECAP ************************************

Step 2 - Downloading Ansible Config and Ansible Inventory Files

This step is needed so that our Ansible system can be configured with various configuration options and the inventory files so we can run the given playbook.

Example 19. D	ownloading Ansible Files
---------------	--------------------------

	Listing 63. Downloading Ansible Files
[root@workstatior	<pre>n ~]# wget http://people.redhat.com/~tmichett/rh354/inventory</pre>
2020-04-16 09: 3	38:50 http://people.redhat.com/~tmichett/rh354/inventory
Resolving people.	.redhat.com (people.redhat.com) 209.132.183.19
Connecting to peo	<pre>>ple.redhat.com (people.redhat.com) 209.132.183.19 :80 connected.</pre>
HIP request sent	c, awaiting response 200 UK
Saving to: 'inven	itory'
inventory	100%[======>] 24KB/s in 0s
2020-04-16 09:38:	:50 (2.69 MB/s) - 'inventory' saved [24/24]
[root@workstatior	<pre>n ~]# wget http://people.redhat.com/~tmichett/rh354/ansible.cfg</pre>
2020-04-16 09: 3	39:34 http://people.redhat.com/~tmichett/rh354/ansible.cfg
Resolving people.	.redhat.com (people.redhat.com) 209.132.183.19
Connecting to peo	<pre>sple.redhat.com (people.redhat.com)[209.132.183.19]:80 connected.</pre>
HIIP request sent	t, awaiting response 200 UK
Saving to: 'ansib	le.cfg'
ansible.cfg	100%[>] 159KB/s in 0s
2020 04 16 00.20	(13.8 MR/c) = (ancible of a' caved [150/150])

Reviewing Ansible Configurations

The **inventory** file provided only has a single host **servera** in there. On real systems, you must be very cautious of running remediation playbooks against an inventory file as it could apply to unintended systems. Additionally the **ansible.cfg** file provided was created for use in this lab environment. Both of these items should be taken into account when doing going through the process on production systems.



[root@workstation ~]# cat inventory
servera.lab.example.com
[root@workstation ~]# cat ansible.cfg
[defaults]
roles_path = /etc/ansible/roles:/usr/share/ansible/roles
log_path = /tmp/ansible.log
inventory = ./inventory
[privilege_escalation]
become=True
[root@workstation ~]#

Step 3 - Run the Ansible Playbook

This step will utilize the **workstation** system which is configured as your Ansible management node and will run the playbook to remediate the results on the **servera** system.

Example 20. Remediation of serverc with Ansible Playbook

<pre>[root@workstation ~]# ansible-playbook Custom_Scan_Fix.yml PLAY [all] ***********************************</pre>
PLAY [all] ***********************************
TASK [Gathering Facts] ************************************
TASK [Set Password Minimum Length in login.defs] ******************************** changed: [servera.lab.example.com]
TASK [Ensure PAM variable difok is set accordingly] ************************************
TASK [Ensure PAM variable ucredit is set accordingly] ************************************
TASK [Ensure PAM variable minlen is set accordingly] ************************************
TASK [Set Password Retry Prompts Permitted Per-Session - system-auth (change)] *** ok: [servera.lab.example.com]
TASK [Set Password Retry Prompts Permitted Per-Session - system-auth (add)] **** changed: [servera.lab.example.com]
TASK [Ensure PAM variable minclass is set accordingly] ************************************
TASK [Ensure PAM variable maxrepeat is set accordingly] ************************************
TASK [Ensure PAM variable ocredit is set accordingly] ************************************
TASK [Ensure PAM variable lcredit is set accordingly] ************************************
TASK [Ensure PAM variable dcredit is set accordingly] ************************************
TASK [Ensure PAM variable maxclassrepeat is set accordingly] ************************************
PLAY RECAP ************************************



After running the playbook, you can see that there were 10 changes that were made to the system and exactly which parameters were changed. The next thing to do is perform another scan of the system to ensure that it is now fully compliant.

Step 4 - Rescan System and Review Results

Example 21. Scanning System after Fixes and Verifying Results

Listing 65. Performing SCAP Verification Scan

[root@servera ~]# oscap xccdf eval \
--profile xccdf_org.ssgproject.content_profile_ospp_customized \
--tailoring-file ssg-rhel8-ds-tailoring.xml \
--results custom_scan_results_fixed.xml \
--report Custom_Scan_Report_Fixed.html \
/usr/share/xml/scap/ssg/content/ssg-rhel8-ds.xml

Listing 66. Copying Results to Workstation

ed.xml Custon	_Scan_Repor	t_Fixed.html workstation:
100% 4086KB	60. 3MB/s	00:00
100% 282KB	48. 4MB/s	00 :00
	ed.xml Custom 100% 4086KB 100% 282KB	ed.xml Custom_Scan_Repor 100% 4086KB 60.3MB/s 100% 282KB 48.4MB/s

Listing 67. Viewing Results on Workstation

[root@workstation ~]# firefox Custom_Scan_Report_Fixed.html

xccar_org.open-scap_testres x | +

Test

system

		L				
-) C' 🏠		i file:///root/Custom_Scan_Report_Fixed	J.html	… ◙ ☆	
	Performed by	root				

Compliance and Scoring

cpe:/a:redhat:openscap:1.3.0

There were no failed or uncertain rules. It seems that no action is necessary.

Rule results

Severity of fai	led rules						
Score							
Scoring system		Score	Maximum	Pe	rcent		
urn:xccdf:scoring:defa	ult	100.000000	100.000000	1	00%		
Rule Overv	view						
✓ pass	🕑 fail	✓ notchecked	Search throu	gh XCCDF rules			Search
 ✓ fixed ✓ informational 	error v unknown	votapplicable	Group rules by:	Default	~		1
Title					Severity	R	esult
Guide to the Secure	Configuration of Red Hat	Enterprise Linux 8					

Figure 28. Fixed SCAP Scan Results Report in Firefox

Appendix B: IdM Server on RHEL8

B.1. Installing the IdM Server on RHEL8



Installing IdM Server on RHEL8

Red Hat RHEL8 IdM KB Article: https://access.redhat.com/articles/3623971

* Requires login

Installing the IdM Server

1. List yum modules

Listing 68. Source Description

[root@servera ~]	# sudo yum modul	e list gr	ep -i idm					
idm	DL1	common [d]	, adtrust,	client,	dns,	server	r The Red Hat Enterprise Linux Identity Management system module	
idm	client [d]	common [d]					RHEL IdM long term support client module	



RHEL8 Application Streams and Modules

Because RHEL8 uses DNF/YUM4 and has AppStreams and Modules, many packages may be setup as defaults while other modules aren't available. The IdM server package is one such package as the IdM server package is installed from the DL1 application stream.

2. Enable IdM Server packages by enabling the IdM DL1 stream

Listing 69. Enabling a Module Stream

Enabling module streams: 389-ds 1.4 httpd 2.4 ide D11
389-ds 1.4 httpd 2.4 idm D11
httpd 2.4
idm DI1
Tulii DL I
pki-core 10.6
pki-deps 10.6
Transaction Summary

3. Obtain information on the Module Stream

Listing 70. Obtaining Module Stream Packages

4. Installing the **ipa-server** package

[root@servera ~]# yum Last metadata expirat Dependencies resolved	n module ion che I.	e install idm:DL1/server ick: 0:11:29 ago on Tue 14 Apr 2020	02:20:48 AM EDT.	
Package	Arch	Version	Repository	Size
Installing group/modu ipa-server	ile pack x86_64	ages: ↓ 4.7.1-11.module+el8+2842+7481110c	rhel-8.0-for-x86_64-appstream-rpms	502 k
<pre> output omitted python3-sss-murmur- redhat-logos-httpd- sssd-dbus-2.0.0-43. sssd-tools-2.0.0-43</pre>	2.0.0-4 80.7-1. el8.x86 8.el8.x8	3.el8.x86_64 el8.noarch i_64 i6_64		
Complete!				



IPA Server Options

It is important to note that for this demonstration, the DNS server functionality was not installed with the YUM command and it is not installed or configured as part of the **ipa-server-install** command.

5. Run the IdM Installer command ipa-server-install with all options required

Listing 72. Performing IPA Server Installation



```
Excluded by options:
  * Configure the NTP client (chronyd)
To accept the default shown in brackets, press the Enter key.
Do you want to configure integrated DNS (BIND)? [no]:
Enter the fully qualified domain name of the computer
on which you're setting up server software. Using the form
<hostname>.<domainname>
Example: master.example.com.
Server host name [servera.lab.example.com]:
The domain name has been determined based on the host name.
Please confirm the domain name [lab.example.com]:
The IPA Master Server will be configured with:
Hostname:
              servera.lab.example.com
IP address(es): 172.25.250.10
Domain name: lab.example.com
Realm name:
               LAB.EXAMPLE.NET
The CA will be configured with:
Subject DN: CN=Certificate Authority,O=LAB.EXAMPLE.NET
Subject base: O=LAB.EXAMPLE.NET
Chaining:
             self-signed
WARNING: Realm name does not match the domain name.
You will not be able to establish trusts with Active Directory unless
the realm name of the IPA server matches its domain name.
Continue to configure the system with these values? [no]: yes
... output omitted ...
Setup complete
Next steps:
    1. You must make sure these network ports are open:
       TCP Ports:
         * 80, 443: HTTP/HTTPS
         * 389, 636: LDAP/LDAPS
         * 88, 464: kerberos
       UDP Ports:
          * 88, 464: kerberos
    2. You can now obtain a kerberos ticket using the command: 'kinit admin'
       This ticket will allow you to use the IPA tools (e.g., ipa user-add)
       and the web user interface.
Be sure to back up the CA certificates stored in /root/cacert.p12
These files are required to create replicas. The password for these
files is the Directory Manager password
The ipa-server-install command was successful
```

6. Open Firewall

Listing 73. IPA Firewall Service Provisioning

[root@servera ~]# firewall-cmd --add-service=freeipa-ldap --add-service=freeipa-ldaps --add-service=dns --permanent
success
[root@servera ~]# firewall-cmd --reload
success

7. Login to IdM Web Console

🐣 SELinux - servera.lab.exar 🗙	🔇 Identity	Manageme	ent ×	+			•					
← → ♂ ✿ ③	🔒 https://s	ervera.lab.	example.co	m/ipa/ui/#/e/us	er/search		⊌	☆	liiX		ē)	≡
RED HAT IDENTITY MANAGEMENT										🛎 Admin	istrate	or ~
Identity Policy Authe	ntication	Netwo	rk Services	IPA Serve								
Users Hosts Services	Grou	ps ID	Views	Automember	~							
User categories	Δcti		rc									
Active users		ve use	13									_
Stage users	Searci	h	Q		2 Refresh	🖻 Delete	+ Add	— Disabl	e 🖌 Enabl	e Acti	ions	~
Preserved users		User login	First name	Last name	Status	UID	Email addres	Email Teleph address Numb		Jo Ti	b tle	
		admin		Administrator	✓ Enabled	297400000						
	Show	/ing 1 to 1 of	1 entries.									

Figure 29. IdM Web Console

B.2. Installing the IdM Client on RHEL8



IdM Client Installation

The default Application Stream and Module for RHEL8 already supports and provides the **ipa-client** package so there is nothing special to perform in order to get the client installed.

1. Install the **ipa-client** package

Listing 74. YUM Installation of ipa-client

<pre>[root@serverb ~]# yum install ipa-client Red Hat Enterprise Linux 8.0 AppStream (dvd) Red Hat Enterprise Linux 8.0 BaseOS (dvd)</pre>	281 kB/s 3 .2 kB 476 kB/s 2 .7 kB	00:00 00:00
output omitted		
sssd-tools-2.0.0-43.el8.x86_64 xmlrpc-c-1.51.0-5.el8.x86_64 xmlrpc-c-client-1.51.0-5.el8.x86_64		
Complete!		

2. Modify DNS settings of system to point to IdM Server

Listing 75. Using Network Manager to Change DNS Settings

<pre>[root@serverb ~]# nmcli connection modify "Wired connection 1" ipv4.dns 172.25.250.10</pre>
<pre>[root@serverb ~]# nmcli connection up "Wired connection 1" Connection successfully activated (D-Bus active path: /org/freedesktop/NetworkManager/ActiveConnection/1325)</pre>
<pre>[root@serverb ~]# cat /etc/resolv.conf # Generated by NetworkManager</pre>
search lab.example.com example.com
nameserver 172.25.250.10

3. Run the client installation command

<pre>[root@serverb ~]# ipa-client-installprincipal adminpassword redhat123mkhomedirunattendedntp-server=172.25.254.254domain lab.example.comserver servera.lab.example.comrealm LAB.EXAMPLE.NET This program will set up IPA client. Version 4.7.1</pre>
Client hostname: serverb.lab.example.com Realm: LAB.EXAMPLE.NET DNS Domain: lab.example.com IPA Server: servera.lab.example.com BaseDN: dc=lab,dc=example,dc=net
Synchronizing time Configuration of chrony was changed by installer. Attempting to sync time with chronyc. Time synchronization was successful. Successfully retrieved CA cert Subject: CN=Certificate Authority,O=LAB.EXAMPLE.NET Issuer: CN=Certificate Authority,O=LAB.EXAMPLE.NET Valid From: 2020-04-14 06:38:52 Valid Until: 2040-04-14 06:38:52
output omitted
SSSD enabled Configured /etc/openldap/ldap.conf Configured /etc/ssh/ssh_config Configured /etc/ssh/sshd_config Configuring lab.example.com as NIS domain. Client configuration complete. The ipa-client-install command was successful

Listing 76. Running the Installer Command for the IdM Client