Memory Tables Answer Key

Chapter 2

Table 2-2 Standard Input, Output, and Error Overview

Name	Default Destination	Use in Redirection	File Descriptor Number
STDIN	Computer keyboard	< (same as 0<)	0
STDOUT	Computer monitor	> (same as 1>)	1
STDERR	Computer monitor	2>	2

Table 2-3 Common Bash Redirectors

Redirector	Explanation
> (same as 1>)	Redirects STDOUT. If redirection is to a file, the current contents of that file are overwritten.
>> (same as 1>>)	Redirects STDOUT. If output is written to a file, the output is appended to that file.
2>	Redirects STDERR.
2>&1	Redirects STDERR to the same destination as STDOUT. Notice that this has to be used in combination with normal output redirection, as in **ls whuhiu > errout 2>&1**.
< (same as 0<)	Redirects STDIN.

Table 2-4 **vim** Essential Commands

vim Command	Explanation
Esc	Switches from input mode to command mode. Press this key before typing any command.
i, **a**	Switches from command mode to input mode at (**i**) or after (**a**) the current cursor position.
o	Opens a new line below the current cursor position and goes to input mode.
:wq	Writes the current file and quits.
:q!	Quits the file without applying any changes. The **!** forces the command to do its work. Add the **!** only if you really know what you are doing.
:w filename	Writes the current file with a new filename.
dd	Deletes the current line.
yy	Copies the current line.
p	Pastes the current selection.
v	Enters visual mode, which allows you to select a block of text using the arrow keys. Use **d** to cut the selection or **y** to copy it.
u	Undoes the last command. Repeat as often as necessary.
Ctrl-r	Redoes the last undo.
gg	Goes to the first line in the document.
G	Goes to the last line in the document.
/text	Searches for *text* from the current cursor position forward.
?text	Searches for *text* from the current cursor position backward.
^	Goes to the first position in the current line.
$	Goes to the last position in the current line.
!ls	Adds the output of **ls** (or any other command) in the current file.
:%s/old/new/g	Replaces all occurrences of *old* with *new*.

Chapter 3

Table 3-2 FHS Overview

Directory	Use
/	Specifies the root directory. This is where the file system tree starts.
/boot	Contains all files and directories that are needed to boot the Linux kernel.

Directory	Use
/dev	Contains device files that are used for accessing physical devices. This directory is essential during boot.
/etc	Contains configuration files that are used by programs and services on your server. This directory is essential during boot.
/home	Used for local user home directories.
/media, /mnt	Contain directories that are used for mounting devices in the file system tree.
/opt	Used for optional packages that may be installed on your server.
/proc	Used by the proc file system. This is a file system structure that gives access to kernel information.
/root	Specifies the home directory of the root user.
/run	Contains process and user-specific information that has been created since the last boot.
/srv	May be used for data by services like NFS, FTP, and HTTP.
/sys	Used as an interface to different hardware devices that are managed by the Linux kernel and associated processes.
/tmp	Contains temporary files that may be deleted without any warning during boot.
/usr	Contains subdirectories with program files, libraries for these program files, and documentation about them.
/var	Contains files that may change in size dynamically, such as log files, mail boxes, and spool files.

Chapter 4

Table 4-2 Essential Tools for Managing Text File Contents

Command	Explanation
less	Opens the text file in a pager, which allows for easy reading
cat	Dumps the contents of the text file on the screen
head	Shows the first ten lines of the text file
tail	Shows the last ten lines of the text file
cut	Used to filter specific columns or characters from a text file
sort	Sorts the contents of a text file
wc	Counts the number of lines, words, and characters in a file

Table 4-3 Most Significant Regular Expressions

Regular Expression	Use
^text	Matches line that starts with specified text.
text$	Matches line that ends with specified text.
.	Wildcard. (Matches any single character.)
[abc]	Matches *a*, *b*, or *c*.
*	Matches zero to an infinite number of the previous character.
\{2\}	Matches exactly two of the previous character.
\{1,3\}	Matches a minimum of one and a maximum of three of the previous character.
colou?r	Matches zero or one of the previous character. This makes the previous character optional, which in this example would match both *color* and *colour*.

Table 4-4 Most Useful **grep** Options

Option	Use
-i	Not case sensitive. Matches upper- and lowercase letters.
-v	Shows only lines that do *not* contain the regular expression.
-r	Searches files in the current directory and all subdirectories.
-e	Searches for lines matching more than one regular expression.
-A <number>	Shows <number> of lines after the matching regular expression.
-B <number>	Shows <number> of lines before the matching regular expression.

Chapter 5

Table 5-2 Common **ssh** Options

Option	Use
-v	Verbose; shows in detail what is happening while establishing the connection
-Y	Enables support for graphical applications
-p <PORT>	Used to connect to an SSH service that is not listening on the default port 22

Table 5-3 Common **rsync** Options

Option	Use
-r	Synchronizes the entire directory tree
-l	Also synchronizes symbolic links
-p	Preserves symbolic links
-n	Performs only a dry run, not actually synchronizing anything
-a	Uses archive mode, thus ensuring that entire subdirectory trees and all file properties will be synchronized
-A	Uses archive mode, and in addition synchronizes ACLs
-X	Synchronizes SELinux context as well

Chapter 6

Table 6-2 Methods to Run Tasks with Elevated Permissions

Method	Description
su	Opens a subshell as a different user, with the advantage that commands are executed as root only in the subshell
sudo	Enables you to set up an environment where specific tasks are executed with administrative privileges
PolicyKit	Enables you to set up graphical utilities to run with administrative privileges

Chapter 7

Table 7-2 Use of Read, Write, and Execute Permissions

Permission	Applied to Files	Applied to Directories
Read	Open a file	List contents of directory
Write	Change contents of a file	Create and delete files
Execute	Run a program file	Change to the directory

Table 7-3 Numeric Representation of Permissions

Permission	Numeric Representation
Read	4
Write	2
Execute	1

Table 7-4 Working with SUID, SGID, and Sticky Bit

Permission	Numeric Value	Relative Value	On Files	On Directories
SUID	4	u+s	User executes file with permissions of file owner.	No meaning.
SGID	2	g+s	User executes file with permissions of group owner.	Files created in directory get the same group owner.
Sticky bit	1	+t	No meaning.	Prevents users from deleting files from other users.

Table 7-5 umask Values and Their Result

Value	Applied to Files	Applied to Directories
0	Read and write	Everything
1	Read and write	Read and write
2	Read	Read and execute
3	Read	Read
4	Write	Write and execute
5	Write	Write
6	Nothing	Execute
7	Nothing	Nothing

Chapter 8

Table 8-2 Binary-Decimal Conversion Overview

Binary Value	Decimal Value
00000000	0
00100000	32
01000000	64
01100000	96
10000000	128
10100000	160
11000000	192
11100000	224

Chapter 9

Table 9-2 Key Options in .repo Files

Option	Explanation
[label]	Contains the label used as an identifier in the repository file.
name=	Specifies the name of the repository.
mirrorlist=	Refers to a URL where information about mirror servers for this server can be obtained. Typically used for big online repositories only.
baseurl=	Refers to the base URL where the RPM packages are found.
gpgcheck=	Set to 1 if a GNU Privacy Guard (GPG) integrity check needs to be performed on the packages. If set to 1, a GPG key is required.
gpgkey=	Specifies the location of the GPG key that is used to check package integrity.

Table 9-3 Common **yum** Tasks

Task	Explanation
search	Search for the exact name of a package.
[**what**]**provides */name**	Perform a deep search in the package to look for specific files within the package.
info	Provide more information about the package.
install	Install the package.
remove	Remove the package.
list [**all** | **installed**]	List all or installed packages.
group list	List package groups.
group install	Install all packages from a group.
update	Update packages specified.
clean all	Remove all stored metadata.

Table 9-4 Yum Module Terminology

Item	Explanation
RPM	The default package format. Contains files, as well as metadata that describes how to install the files. Optionally may contain pre- and post-installation scripts as well.
Module	A delivery mechanism to install RPM packages. In a module different versions and profiles can be provided.
Application stream	A specific version of the module.
Profile	A collection of packages that are installed together for a particular use case.

Table 9-5 Common RPM Query Commands

Command	Use
rpm -qf	Uses a filename as its argument to find the specific RPM package a file belongs to.
rpm -ql	Uses the RPM database to provide a list of files in the RPM package.
rpm -qi	Uses the RPM database to provide package information (equivalent to **yum info**).
rpm -qd	Uses the RPM database to show all documentation that is available in the package.
rpm -qc	Uses the RPM database to show all configuration files that are available in the package.
rpm -q --scripts	Uses the RPM database to show scripts that are used in the package. This is particularly useful if combined with the **-p** option.
rpm -qp <*pkg*>	The **-p** option is used with all the previously listed options to query individual RPM package files instead of the RPM package database. Using this option before installation helps you find out what is actually in the package before it is installed.
rpm -qR	Shows dependencies for a specific package.
rpm -V	Shows which parts of a specific package have been changed since installation.
rpm -Va	Verifies all installed packages and shows which parts of the package have been changed since installation. This is an easy and convenient way to do a package integrity check.
rpm -qa	Lists all packages that are installed on this server.

Chapter 10

Table 10-2 Job Management Overview

Command	Use
& (used at the end of a command line)	Starts the command immediately in the background.
Ctrl-Z	Stops the job temporarily so that it can be managed. For instance, it can be moved to the background.
Ctrl-D	Sends the End Of File (EOF) character to the current job to indicate that it should stop waiting for further input.
Ctrl-C	Can be used to cancel the current interactive job.
bg	Continues the job that has just been frozen using Ctrl-Z in the background.
fg	Brings back to the foreground the last job that was moved to background execution.
jobs	Shows which jobs are currently running from this shell. Displays job numbers that can be used as an argument to the commands **bg** and **fg**.

Table 10-3 Linux Process States Overview

State	Meaning
Running (R)	The process is currently active and using CPU time, or in the queue of runnable processes waiting to get services.
Sleeping (S)	The process is waiting for an event to complete.
Uninterruptible sleep (D)	The process is in a sleep state that cannot be stopped. This usually happens while a process is waiting for I/O.
Stopped (T)	The process has been stopped, which typically has happened to an interactive shell process, using the Ctrl-Z key sequence.
Zombie (Z)	The process has been stopped but could not be removed by its parent, which has put it in an unmanageable state.

Table 10-4 Tuned Profile Overview

Profile	Use
balanced	The best compromise between power usage and performance
desktop	Based on the balanced profile, but tuned for better response to interactive applications
latency-performance	Tuned for maximum throughput
network-latency	Based on latency-performance, but with additional options to reduce network latency
network-throughput	Based on throughput-performance, optimizes older CPUs for streaming content
powersave	Tunes for maximum power saving
throughput-performance	Tunes for maximum throughput
virtual-guest	Optimizes Linux for running as a virtual machine
virtual-host	Optimizes Linux for use as a KVM host

Chapter 11

Table 11-2 Systemd Status Overview

Status	Description
Loaded	The unit file has been processed and the unit is active.
Active(running)	The unit is running with one or more active processes.
Active(exited)	The unit has successfully completed a one-time run.
Active(waiting)	The unit is running and waiting for an event.
Inactive(dead)	The unit is not running.

Status	Description
Enabled	The unit will be started at boot time.
Disabled	The unit will not be started at boot time.
Static	The unit cannot be enabled but may be started by another unit automatically.

Table 11-3 **systemctl** Unit Overview Commands

Command	Description
systemctl --type=service	Shows only service units
systemctl list-units --type=service	Shows all active service units (same result as the previous command)
systemctl list-units --type=service --all	Shows inactive service units as well as active service units
systemctl --failed --type=service	Shows all services that have failed
systemctl status -l your.service	Shows detailed status information about services

Chapter 13

Table 13-2 System Log Files Overview

Log File	Explanation
/var/log/messages	Contains the most commonly used log file; it is the generic log file where most messages are written to.
/var/log/dmesg	Contains kernel log messages.
/var/log/secure	Contains authentication-related messages. Look here to see which authentication errors have occurred on a server.
/var/log/boot.log	Contains messages that are related to system startup.
/var/log/audit/audit.log	Contains audit messages. SELinux writes to this file.
/var/log/maillog	Contains mail-related messages.
/var/log/samba	Provides log files for the Samba service. Notice that Samba by default is not managed through rsyslog, but writes directly to the /var/log directory.
/var/log/sssd	Contains messages that have been written by the sssd service, which plays an important role in the authentication process.
/var/log/cups	Contains log messages that were generated by the print service CUPS.
/var/log/httpd/	Contains log files that are written by the Apache web server. Notice that Apache writes messages to these files directly and not through rsyslog.

Table 13-3 rsyslogd Facilities

Facility	Used by
auth / authpriv	Messages related to authentication.
cron	Messages generated by the **crond** service.
daemon	Generic facility that can be used for nonspecified daemons.
kern	Kernel messages.
lpr	Messages generated through the legacy lpd print system.
mail	Email-related messages.
mark	Special facility that can be used to write a marker periodically.
news	Messages generated by the NNTP news system.
security	Same as auth/authpriv. Should not be used anymore.
syslog	Messages generated by the syslog system.
user	Messages generated in user space.
uucp	Messages generated by the legacy UUCP system.
local0-7	Messages generated by services that are configured by any of the local0 through local7 facilities.

Table 13-4 rsyslogd Priorities

Priority	Description
debug	Debug messages that will give as much information as possible about service operation.
info	Informational messages about normal service operation.
notice	Informational messages about items that might become an issue later.
warning / warn	Something is suboptimal, but there is no real error yet.
err /error	A noncritical error has occurred.
crit	A critical error has occurred.
alert	Message used when the availability of the service is about to be discontinued.
emerg / panic	Message generated when the availability of the service is discontinued.

Chapter 14

Table 14-3 Common Disk Device Types

Device Name	Description
/dev/sda	A hard disk that uses the SCSI driver. Used for SCSI and SATA disk devices. Common on physical servers but also in VMware virtual machines.
/dev/nvme0n1	The first hard disk on an NVM Express (NVMe) interface. NVMe is a server-grade method to address advanced SSD devices. Note at the end of the device name that the first disk in this case is referred to as *n1* instead of *a* (as is common with the other types).
/dev/hda	The (legacy) IDE disk device type. You will seldom see this device type on modern computers.
/dev/vda	A disk in a KVM virtual machine that uses the virtio disk driver. This is the common disk device type for KVM virtual machines.
/dev/xvda	A disk in a Xen virtual machine that uses the Xen virtual disk driver. You see this when installing RHEL as a virtual machine in Xen virtualization. RHEL 8 cannot be used as a Xen hypervisor, but you might see RHEL 8 virtual machines on top of the Xen hypervisor using these disk types.

Table 14-4 File System Overview

File System	Description
XFS	The default file system in RHEL 8.
Ext4	The default file system in previous versions of RHEL; still available and supported in RHEL 8.
Ext3	The previous version of Ext4. On RHEL 8, there is no need to use Ext3 anymore.
Ext2	A very basic file system that was developed in the early 1990s. There is no need to use this file system on RHEL 8 anymore.
BtrFS	A relatively new file system that is not supported in RHEL 8.
NTFS	A Windows-compatible file system that is not supported on RHEL 8.
VFAT	A file system that offers compatibility with Windows and Mac and is the functional equivalent of the FAT32 file system. Useful on USB thumb drives that exchange data with other computers but not on a server's hard disks.

Table 14-5 /etc/fstab Fields

Field	Description
Device	The device that must be mounted. A device name, UUID, or label can be used.
Mount Point	The directory or kernel interface where the device needs to be mounted.
File System	The file system type.
Mount Options	Mount options.
Dump Support	Use 1 to enable support to back up using the **dump** utility. This may be necessary for some backup solutions.
Automatic Check	This field specifies whether the file system should be checked automatically when booting. Use 0 to disable automated check, 1 if this is the root file system and it has to be checked automatically, and 2 for all other file systems that need automatic checking while booting. Network file systems should have this option set to 0.

Table 14-6 Common Mount Options

Option	Use
auto / noauto	Mounts [does not mount] the file system automatically.
acl	Adds support for file system access control lists (see Chapter 7, "Permissions Management").
user_xattr	Adds support for user-extended attributes (see Chapter 7).
ro	Mounts the file system in read-only mode.
atime / noatime	Disables or enables access time modifications.
noexec / exec	Denies or allows execution of program files from the file system.
_netdev	Mounts a network file system. This option tells fstab to wait until the network is available before mounting this file system.

Chapter 15

Table 15-2 LVM Management Essential Commands

Command	Explanation
pvcreate	Creates physical volumes
pvs	Shows a summary of available physical volumes
pvdisplay	Shows a list of physical volumes and their properties

Command	Explanation
pvremove	Removes the physical volume signature from a block device
vgcreate	Creates volume groups
vgs	Shows a summary of available volume groups
vgdisplay	Shows a detailed list of volume groups and their properties
vgremove	Removes a volume group
lvcreate	Creates logical volumes
lvs	Shows a summary of all available logical volumes
lvdisplay	Shows a detailed list of available logical volumes and their properties
lvremove	Removes a logical volume

Chapter 16

Table 16-2 Linux Kernel Module Management Overview

Command	Use
lsmod	Lists currently loaded kernel modules
modinfo	Displays information about kernel modules
modprobe	Loads kernel modules, including all of their dependencies
modprobe -r	Unloads kernel modules, considering kernel module dependencies

Chapter 18

Table 18-2 Boot Phase Configuration and Troubleshooting Overview

Boot Phase	Configuring It	Fixing It
POST	Hardware configuration (F2, Esc, F10, or another key).	Replace hardware.
Selecting the bootable device	BIOS/UEFI configuration or hardware boot menu.	Replace hardware or use rescue system.
Loading the boot loader	**grub2-install** and edits to /etc/defaults/grub.	Use the GRUB boot prompt and edits to /etc/defaults/grub, followed by **grub2-mkconfig**.
Loading the kernel	Edits to the GRUB configuration and /etc/dracut.conf.	Use the GRUB boot prompt and edits to /etc/defaults/grub, followed by **grub2-mkconfig**.

Boot Phase	Configuring It	Fixing It
Starting /sbin/init	Compiled into initramfs.	Use the **init= kernel** boot argument, **rd.break** kernel boot argument.
Processing initrd.target	Compiled into initramfs.	Use the **dracut** command. (You won't often have to troubleshoot this.)
Switch to the root file system	Edits to the /etc/fstab file.	Apply edits to the /etc/fstab file.
Running the default target	Using **systemctl set-default** to create the /etc/systemd/system/default.target symbolic link	Start the rescue.target as a kernel boot argument.

Chapter 20

Table 20-2 Most Useful sshd Configuration Options

Option	Use
Port	Defines the TCP listening port.
PermitRootLogin	Indicates whether to allow or disallow root login.
MaxAuthTries	Specifies the maximum number of authentication tries. After reaching half of this number, failures are logged to syslog.
MaxSessions	Indicates the maximum number of sessions that can be open from one IP address.
AllowUsers	Specifies a space-separated list of users who are allowed to connect to the server.
PasswordAuthentication	Specifies whether to allow password authentication. This option is on by default.
GSSAPIAuthentication	Indicates whether authentication through the GSSAPI needs to be enabled. Used for Kerberos-based authentication.
TCPKeepAlive	Specifies whether or not to clean up inactive TCP connections.
ClientAliveInterval	Specifies the interval, in seconds, that packets are sent to the client to figure out if the client is still alive.
ClientAliveCountMax	Specifies the number of client alive packets that need to be sent.
UseDNS	If on, uses DNS name lookup to match incoming IP addresses to names.
ServerAliveInterval	Specifies the interval, in seconds, that a client sends a packet to a server to keep connections alive.
ServerAliveCountMax	Specifies the maximum number of packets a client sends to a server to keep connections alive.

Chapter 22

Table 22-2 SELinux Core Elements

Element	Use
Policy	A collection of rules that define which source has access to which target.
Source domain	The object that is trying to access a target. Typically a user or a process.
Target domain	The thing that a source domain is trying to access. Typically a file or a port.
Context	A security label that is used to categorize objects in SELinux.
Rule	A specific part of the policy that determines which source domain has which access permissions to which target domain.
Labels	Same as a context label, defined to determine which source domain has access to which target domain.

Chapter 23

Table 23-2 Firewalld Default Zones

Zone Name	Default Settings
block	Incoming network connections are rejected with an "icmp-host-prohibited" message. Only network connections that were initiated on this system are allowed.
dmz	For use on computers in the demilitarized zone. Only selected incoming connections are accepted, and limited access to the internal network is allowed.
drop	Any incoming packets are dropped and there is no reply.
external	For use on external networks with masquerading (Network Address Translation [NAT]) enabled, used especially on routers. Only selected incoming connections are accepted.
home	For use with home networks. Most computers on the same network are trusted, and only selected incoming connections are accepted.
internal	For use in internal networks. Most computers on the same network are trusted, and only selected incoming connections are accepted.
public	For use in public areas. Other computers in the same network are not trusted, and limited connections are accepted. This is the default zone for all newly created network interfaces.
trusted	All network connections are accepted.
work	For use in work areas. Most computers on the same network are trusted, and only selected incoming connections are accepted.

Table 23-3 Common **firewall-cmd** Options

firewall-cmd Options	Explanation
--get-zones	Lists all available zones
--get-default-zone	Shows the zone currently set as the default zone
--set-default-zone=<*ZONE*>	Changes the default zone
--get-services	Shows all available services
--list-services	Shows services currently in use
--add-service=<*service-name*> [--zone=<*ZONE*>]	Adds a service to the current default zone or the zone that is specified
--remove-service=<*service-name*>	Removes a service from the configuration
--list-all-zones	Shows configuration for all zones
--add-port=<port/protocol> [--zone=<*ZONE*>]	Adds a port and protocol
--remove-port=<port/protocol> [--zone=<*ZONE*>]	Removes a port from the configuration
--add-interface=<*INTERFACE*> [--zone=<*ZONE*>]	Adds an interface to the default zone or a specific zone that is specified
--remove-interface=<*INTERFACE*> [--zone=<*ZONE*>]	Removes an interface from a specific zone
--add-source=<*ipaddress/netmask*> [--zone=<*ZONE*>]	Adds a specific IP address
--remove-source=<*ipaddress/netmask*> [--zone=<*ZONE*>]	Removes an IP address from the configuration
--permanent	Writes configuration to disk and not to runtime
--reload	Reloads the on-disk configuration

Chapter 25

Table 25-2 Understanding Linux Time

Concept	Explanation
Hardware clock	The hardware clock that resides on the main card of a computer system
Real-time clock	Same as the hardware clock
System time	The time that is maintained by the operating system
Software clock	Similar to system time

Concept	Explanation
Coordinated Universal Time (UTC)	A worldwide standard time
Daylight saving time	Calculation that is made to change time automatically when DST changes occur
Local time	The time that corresponds to the time in the current time zone

Table 25-3 Commands Related to RHEL 8 Time Management

Command	Short Description
date	Manages local time
hwclock	Manages hardware time
timedatectl	Developed to manage all aspects of time on RHEL 8

Table 25-4 **timedatectl** Command Overview

Command	Explanation
status	Shows the current time settings
set-time TIME	Sets the current time
set-timezone ZONE	Sets the current time zone
list-timezone	Shows a list of all time zones
set-local-rtc [0|1]	Controls whether the RTC (the real-time clock, normally referred to as the hardware clock) is in local time
set-ntp [0|1]	Controls whether NTP is enabled