Lesson Topic: Lustre File System Fundamentals

Instructional Goal: Computational scientists who are novice users of the Lustre File System, will be able to use basic commands that enable parallel I/O in their computing application. They will have access to a high-performance computing system with a Lustre File System and documentation of basic Lustre commands.

Terminal Objective 1.0: Compute the data storage capacity of a Lustre file system

<u>Enabling Objective 1:</u> Given a list of the four main Lustre filesystem components, the student will identify the component on which user data is stored. One correct answer will be selected from the given list.

<u>Enabling Objective 2:</u> Given the number of OSSs and size of the OSTs on a Lustre filesystem, the student will compute the total storage capacity of the filesystem. The correct answer will be selected from a list of possible values.

Terminal Objective 2.0: Recognize how Lustre uses striping to enable parallel I/O

Enabling Objective 3: Given a list of the three Lustre stripe settings, the student will identify the stripe setting that determines the degree of parallelism that can be achieved by an application. One correct answer will be selected from the given list. Enabling Objective 4: Given the number of OSTs and their storage capacity on a Lustre filesystem, the student will compute the maximum file size allowed in the given situation when using striping. The correct value will be selected from a list of possible values. Enabling Objective 5: Given a written description, the student will correctly indicate whether it does or does not describe a benefit of striping. The correct answer will be selected from a list of possible values. <u>Enabling Objective 6:</u> Given a list of possible responses, the student will correctly identify (select) the one response which describes a characteristic of aligned stripes. <u>Enabling Objective 7:</u> Given the name of a stripe setting, the student will correctly identify (select) its description. The correct value will be selected from a list of possible values.

<u>Enabling Objective 8:</u> Given a graphical depiction showing how a specified number of files is distributed across a given number of OSTs, the stripe size value and the stripe index value, correctly identify (select) the stripe count value that describes the pattern depicted. The correct value will be selected from a list of possible values.

<u>Enabling Objective 9:</u> Given descriptions of how two applications have implemented striping, the student will correctly identify (select) the application that will achieve the best performance.

Terminal Objective 3.0 Use basic Lustre utility commands

- 3.1 List OSTs in the file system
- 3.2 Search the directory tree
- 3.3 Check disk space usage
- 3.4 Get striping information
- 3.5 Set striping patterns

<u>Enabling Objective 10:</u> Given a description of an action enabled by one of the Lustre lfs commands, the student will correctly identify the command which enables the described action. The correct value will be selected from a list of the basic Lustre lfs commands: df, find, getstripe, osts, and setstripe.

<u>Enabling Objective 11:</u> Given a file directory name, the student will correctly identify (select) the command that would most efficiently list the files in the directory and its subdirectories. The correct value will be selected from a list of possible values. <u>Enabling Objective 12:</u> Given output from the lfs osts command, the student will correctly construct the lfs command to recursively list all files on an OST listed in the output. The command will be constructed by selecting each component of the command from a list of possible values.

<u>Enabling Objective 13:</u> Given access to the basic Lustre user commands and their syntax, the student will correctly construct the Lustre command that displays the filesystem disk space usage in human readable format. The command will be constructed by entering it in a short answer text format.

<u>Enabling Objective 14:</u> Given the name of the lfs command to use (setstripe), the student will correctly identify (select) the index value which allows the filesystem to automatically select where to place the first stripe. The correct value will be selected from a list of possible values.

<u>Enabling Objective 15:</u> Given a statement indicating that the stripe pattern of a given directory can be directly changed using the lfs setstripe command, the student will correctly discriminate whether it is a correct statement. The answer will be indicated by selecting either true or false.

<u>Enabling Objective 16:</u> Given a filename, directory name, the number of OSTs the file will be striped across, and the size of each stripe, the student will correctly construct the lfs command to create the file. The command will be constructed by selecting each component of the command from a list of possible values.

Lustre Fundamentals Post-test

Performance	Assessment Item
Objective	
Measured	
1	 Which Lustre file system component stores user data? Select one: Metadata Target Object Storage Servers Metadata Server Object Storage Targets
2	What is the capacity of a Lustre file system with 32 OSSs, each with two 16 TB OSTs? Select one: 512 TB 1024 TB 8192 TB 3072 TB
3	 Which one of the following stripe settings determines the degree of parallelism that can be achieved? Select one: start index stripe count stripe size

4	If a file system has 320 OSTs with a capacity of 2 TBs, using striping the maximum file size would be: Select one: • 640 TBs • a multiple of 320 TBs • a multiple of 640 TBs
5	Indicate whether each of the following is or is not a benefit of striping by choosing Yes or No from the dropdown list. Since files are striped across all servers there is less risk of them being lost if a server breaks down. Increases the available I/O bandwidth. Reduces the overhead and risk associated with I/O operations. No Provides space for very large files.
6	 Which one of the following is a characteristic of aligned stripes? Select one: processes write data as contiguous regions across stripe boundaries some of the OSTs over which the file is striped are simultaneously receiving data from more than one process potential for causing contention resulting in degraded performance processes access the file at offsets that correspond to stripe boundaries
7	Match the following stripe settings with its description in the dropdown list. Stripe size Number of bytes written on one OST before cycling to the next \$ Stripe index Starting number of OST to place the first stripe Stripe count Number of OSTs over which a file is distributed



10	Match the Lustre Ifs utility option that enables the action listed.
	Check disk space usage df
	Search the directory tree find
	Get striping information
	List OSTs in the file system osts
	Set striping patterns setstripe
11	Which command would you enter to efficiently list all files in the directory /kraken/lustre/ and its subdirectories?
	Select one: find /kraken/lustre Is -I Ifs find /kraken/lustre
12	Given the following output from the lfs osts command, what command would you enter to recursively list all files in the directory /kraken/lustre/ that have objects on OST number 334? <pre>\$ lfs osts OBDS: 0: scratch-OST0000_UUID ACTIVE 1: scratch-OST0001_UUID ACTIVE</pre>
	<pre>335: scratch-0ST014f_UUID ACTIVE /lustre/scratch stripe_count: 4 stripe_size: 0 stripe_offset: -1</pre>
	Ifs find \$obd \$ scratch-OST014e_UUID \$ /kraken/lustre/

13	You attempt to write to the file system and you get an OST full error message. To further investigate the problem you decide to check the disk usage on the system. What command would you enter to get this information in human readable format? Answer: Ifs df -h
14	Using the Ifs setstripe command, which of the following index values would allow the system to automatically select where to place the first stripe? Select one: 1 0 1 - 1 - 1
15	You can directly change the stripe pattern on files within a given directory using the Ifs setstripe command. Select one: True False
16	What command would you enter to create a file named astronomy1 in the directory /kraken/lustre that will be striped on six OSTs with 256 KB on each stripe? Ifs setstripe + -s + 256k -c + 6 /kraken/lustre/astronomy1 +