

Lustre Test Tools Environment Variables

This page describes the known about environment variables and users used to drive the Lustre tests. They are typically stored in a configuration script in `lustre/tests/cfg/$NAME.sh`, accessed by the `NAME=name` environment variable within the test scripts. The default configuration for a single-node test is `NAME=local`, which accesses the `lustre/tests/cfg/local.sh` configuration file. If creating a custom test configuration, is recommended to specify the non-default configuration values in your config file, then source the default configuration via `". $LUSTRE/tests/cfg/local.sh"` so that any changes to the default configuration will also be included in the custom configuration.

The list is not complete and so if you find additional variables please add them to this list with a description of what they do. There are also variables on the list that do an unknown function so if you know what they do please describe them.

Cluster Configuration

Variable	Purpose	Typical Value
mds_HOST	The network hostname of the primary MDS host. Uses local hostname if unset.	mdsnode-1
ost_HOST	The network hostname of the primary OSS host. Uses local hostname if unset. This seems to be somewhat misnamed.	ossnode-1
mgs_HOST	The network hostname of the MGS host. Uses <code>\$mds_HOST</code> if unset.	mgsnode
CLIENTS	A comma separated list of the clients to be used for testing	client-1,client-2,client-5
RCLIENTS	A command separated list of the remote clients to be used for testing. One client actually executes the test framework, the other clients are remote client.	client-1,client-5 which would suggest from the above that client-2 is actually running the test framework.
NETTYPE	The network infrastructure to be used in lnet notation.	tcp or o2ib
mdsfailover_HOST	If testing High Availability, the hostname of the backup MDS that can access the MDT storage.	mdsnode-2
ostfailover_HOST	If testing High Availability, the hostname of the backup OSS that can access the OST storage.	ossnode-2
ost1_HOST	The OSS hostname for <code>OST0000</code> , if there are multiple OSS nodes.	ossnode-1
ost1failover_HOST	The backup OSS hostname for <code>OST0000</code> , if there are multiple OSS nodes.	ossnode-2
MDSDEVBASE	Base block device name for generating multiple <code>MDSDEVn</code> device names from, used by appending the MDT number to the base name.	/dev/vg_test/lvmdt
MDSDEV1	The block device that <code>MDT0000</code> will use for its storage, either specified explicitly in the config, or generated from <code>\$MDSDEVBASE</code> .	/dev/vg_test/lvmdt1
MDSCOUNT	Number of MDTs to use. Only valid for Lustre 2.4 and later.	1
MDSSIZE	The size of the MDT storage in kilobytes. Can be smaller than the MDT block device, to speed testing. Use the block device size if unspecified or equal to zero.	200000
MDSOPT	Options to be used when mounting the MDT storage device.	- mountoptions=errors=remount-ro,user_xattr,acl
MGSDEV	Block device to use for the MGS. Defaults to be shared with the MDS block device.	/dev/vg_test/lvmgt
MGSNID	LNNet Node ID if it does not map to the primary address of <code>\$mgs_HOST</code> for <code>\$NETTYPE</code>	<code>\$mgs_HOST</code>
MGSSIZE	The size of the MGT storage in kilobytes, if it is separate from the MDT device. Can be smaller than the MGT block device, to speed testing. Use the block device size if unspecified or equal to zero.	16384
REFORMAT	Whether the filesystems will be reformatted between tests.	true
LNNet_ACCEPT_PORT	The TCP port the LNet acceptor will be listening on for remote client/server connections	988
ACCEPTOR_PORT	Same as <code>LNNet_ACCEPT_PORT</code> (should be removed)	988

OSTCOUNT	The number of OSTs that are being provided by the OSS on <code>\$ost_HOST</code>	5
OSTSIZE	Size of the OST storage in kilobytes. Can be smaller than the OST block device, to speed testing. Use the block device size if unspecified or equal to zero.	1000000
OSTDEVn	The block devices to be used for the OST "n" storage (e.g. OSTDEV1, OSTDEV2, etc.)	/dev/vg_test/lvost1, /dev/vg_test/lvost2, etc.
FSTYPE	Filesystem type to use for all OST and MDT devices.	ldiskfs, zfs
OSTFSTYPE	Filesystem type to use for all OST devices. Defaults to FSTYPE if unspecified.	ldiskfs, zfs
ostn_FSTYPE	Filesystem type to use for the "ostn" device (e.g. ost1, ost2, ...). Defaults to OSTFSTYPE if unspecified.	ldiskfs, zfs
MDTFSTYPE	Filesystem type to use for all MDT devices. Defaults to FSTYPE if unspecified.	ldiskfs, zfs
STRIPE_BYTES	Default stripe size for the whole filesystem.	1048576
STRIPES_PER_OBJ	Default number of stripes per file for the whole filesystem. If zero, use the built-in default value of 1.	0
SINGLEMDS		mds1
TIMEOUT	Lustre timeout to use during testing, in seconds. Reduced from the default to speed testing.	20
FSNAME	The Lustre filesystem name to use for testing. Uses <code>lustre</code> by default, though this can be confusing.	testfs
MOUNT	Mountpoint to use for the client test filesystem(s). Defaults to <code>/mnt/\$FSNAME</code> if unspecified.	/mnt/testfs
MOUNT1	First mountpoint to use for dual-mount testing (e.g. <code>sanityn.sh</code>). Defaults to <code>\$MOUNT</code> if unspecified.	/mnt/testfs
MOUNT2	Second mountpoint to use for dual-mount testing. Defaults to <code>\${MOUNT}2</code> if unspecified	/mnt/testfs2
DIR	Directory in which to run the Lustre tests. Must be within the mountpoint specified by <code>\$MOUNT</code> . Defaults to <code>\$MOUNT</code> if unspecified.	<code>\$MOUNT</code>
DIR1	Directory in which to run the dual-mount Lustre tests. Must be within the mountpoint specified by <code>\$MOUNT1</code> . Defaults to <code>\$MOUNT1</code> if unspecified.	<code>\$MOUNT1</code>
DIR2	Directory in which to run the dual-mount Lustre tests. Must be within the mountpoint specified by <code>\$MOUNT2</code> . Defaults to <code>\$MOUNT2</code> if unspecified.	<code>\$MOUNT2</code>
LDISKFS_MKFS_OPTS	Filesystem type-specific formatting options passed to <code>mkfs.lustre</code> via the <code>--mkfsoptions=</code> argument.	-O large_xattr
MDT_FS_MKFS_OPTS	MDT-specific formatting options passed to <code>mkfs.lustre</code> via the <code>--mkfsoptions=</code> argument.	
OST_FS_MKFS_OPTS	OST-specific formatting options passed to <code>mkfs.lustre</code> via the <code>--mkfsoptions=</code> argument.	
MDSISIZE	Size of inodes on the MDT devices. Uses the <code>mkfs.lustre</code> default of 512 bytes if unspecified.	
MDSJOURNALSIZE	Size of the MDT journal in kilobytes. Uses the <code>mkfs.lustre</code> default (based on MDT device size, between 8MB and 400MB) if unspecified.	128
MDS_MOUNT_OPTS	Mount options passed to the MDT device filesystems.	-o user_xattr,acl
OST_MOUNT_OPTS	Mount options passed to the OST device filesystems.	
MOUNTOPT	Mount options passed to the client <code>mount</code> command.	-o user_xattr

PTLDEBUG	Lustre debug level at which to run tests. Use -1 for full debug.	"vfstrace rpctrace dlmtrace neterror ha config super"
SUBSYSTEM	Lustre subsystems which to include in the debug logs.	"all -lnet -lnd -pinger"
RUNAS_ID	The non-root numeric user ID to use for running tests that need to verify permissions or other behaviour that should not be run as root. There does not need to be an account associated with this user ID.	500
RUNAS_GID	The non-root numeric group ID to use for running tests that need to verify permissions or other behaviour that should not be run as root. There does not need to be an account associated with this group ID.	500
RUNAS	Wrapper program used to change the effective user and group ID of a running program to \$RUNAS_ID and \$RUNAS_GID. This should probably not be changed.	runas
PDSH	The parallel shell command to use for running shell operations on one or more remote hosts.	"pdsh -S -Rrsh -w"
RSYNC_RSH		rsh
FAILURE_MODE	The failure mode to use for failover testing. <i>SOFT</i> failure means that node failures will be simulated by blocking IO to the device and stopping the Lustre services. <i>HARD</i> failure means that the nodes will be rebooted using the \$POWER_DOWN and \$POWER_UP commands to more accurately simulate node failure.	SOFT, HARD
POWER_DOWN	Command used to hard power off a node to simulate node failure.	powerman --off
POWER_UP	Command used to power on a node to simulate node failure, after \$POWER_DOWN is used.	powerman --on
SLOW	Whether the "slow" subtests should be run (those taking five minutes or more per subtest).	no, yes
FAIL_ON_ERROR	Whether test scripts should exit immediately after a failure, or if they should continue running the remaining tests.	true, false
SGPDD_YES	Whether the <i>sgpdd</i> test should be run, which will clobber all data on the tested devices.	true, false
MPIRUN	Command to use for launching MPI test programs.	mpirun
MPIRUN_OPTIONS	Options given to \$MPIRUN when running MPI test programs.	"-mca boot ssh -mca btl tcp,self"
MPI_USER		mpiuser
SHARED_DIRECTORY	This is used by a small number of tests to share state between the client running the tests, or in some cases between the servers (e.g. <i>lfsck.sh</i>). It needs to be a non-lustre filesystem that is available on all the nodes.	/tmp
DAEMONFILE	The filename for Lustre debug logs, if expecting to collect a large amount of debugging information.	/tmp/debug.log
DAEMONSIZE	The maximum size of the \$DAEMONFILE debug log file, in megabytes.	500
SINGLECLIENT	The client node on which to run tests that should run on only a single-threaded test.	\$(hostname)
HSM Environment Variables		
AGTCOUNT	Number of archive agents. Default: Number of remote client nodes; \$CLIENTCOUNT - 1	1
AGTDEV1	Root path of the archive file system mount point. Default: \$TMP/arc1	/mnt/archive
AGTDEV{N}	Root path of the archive file system mount point for each additional agent. Default: \$TMP/arc2, \$TMP/arc3, etc.	/mnt/archive2
agt1_HOST	Host name of the first archive agent.	agtnode

agt{N} _HOST	Host name of any additional archive agents.	agtnode2
SINGLEAGENT	Facet of the single agent. Default: agt1	
HSMTOOL	Name of the HSM copytool. Default: lhsmtool_posix	lhsmtool_posix
HSMTOOL_VERBOSE	Input to the \$HSMTOOL command on start up. Verbose level of the HSM copytool. Default: No verbose output	"-v -v"
HSMTOOL_UPDATE_INTERVAL	Input to the \$HSMTOOL command on start up. Update interval of the HSM copytool. Default: No interval specified	
HSMTOOL_EVENT_FIFO	Input to the \$HSMTOOL command on start up. Default: No event specified	

Users

Test ConfigurationUsers