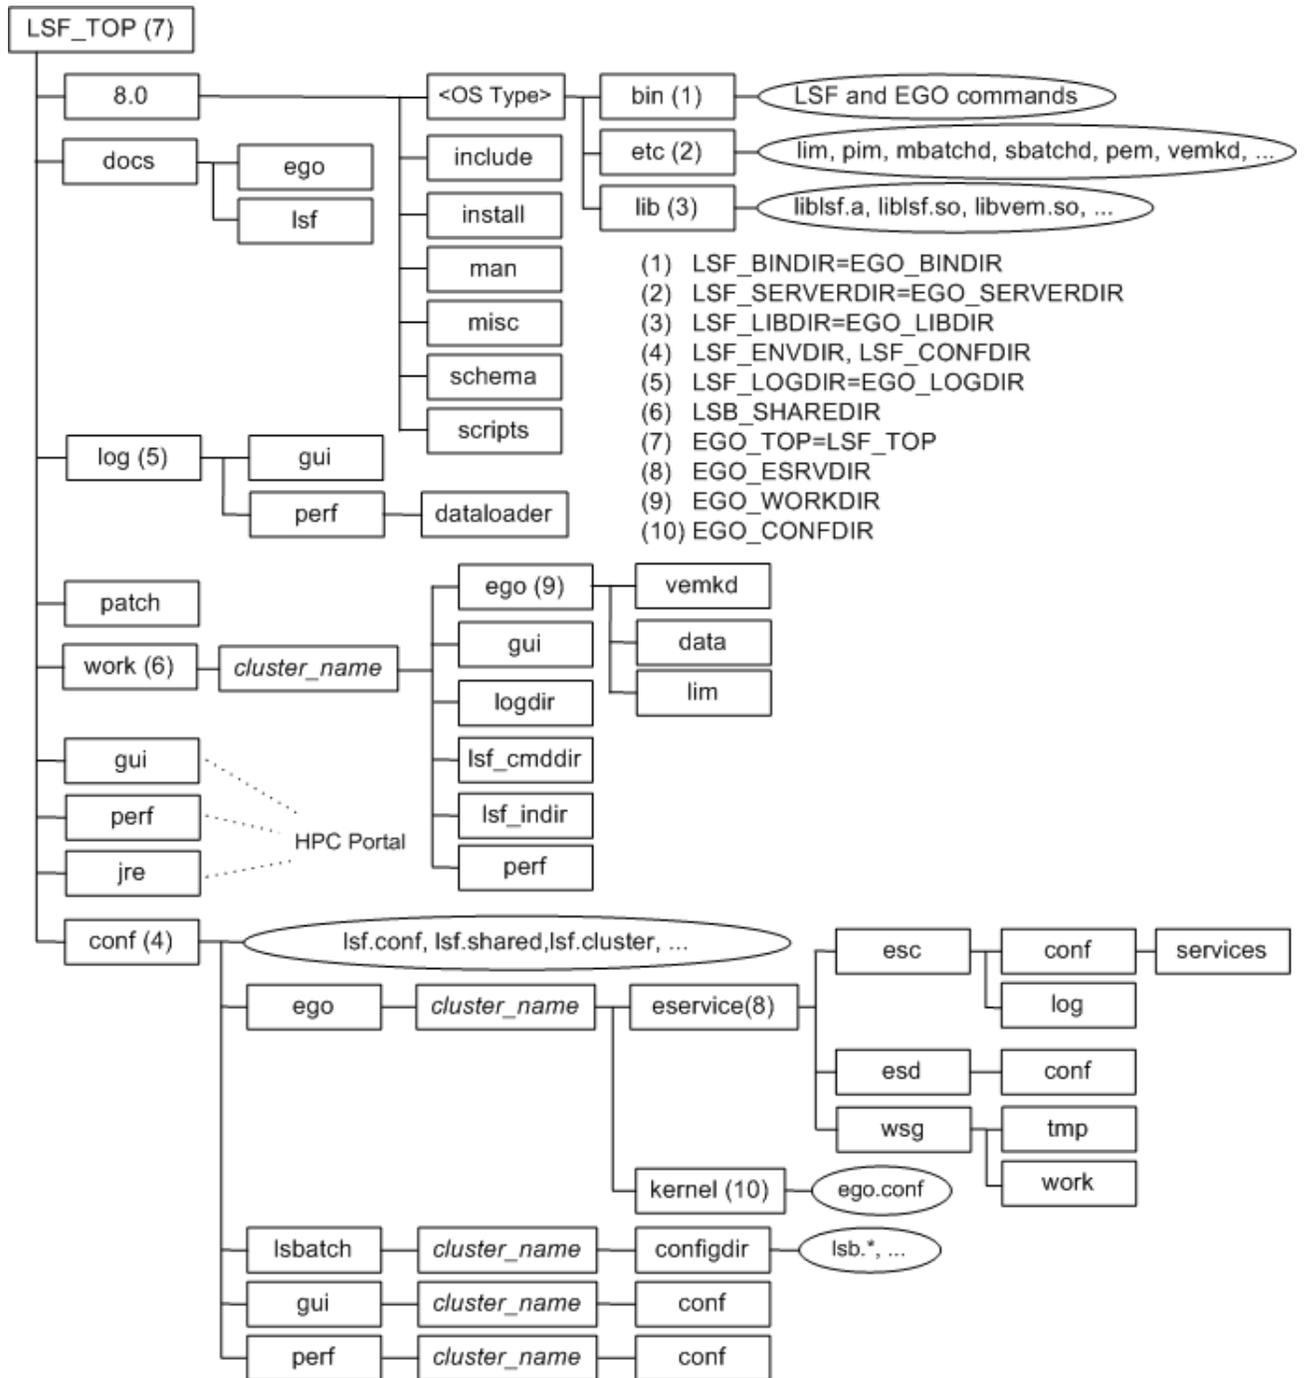

Platform LSF 8.0 Quick Reference

January 2011



Platform LSF 8.0 Quick Reference

Sample UNIX installation directories



Daemon error log files

Daemon error log files are stored in the directory defined by `LSF_LOGDIR` in `lsf.conf`.

LSF base system daemon log files	LSF batch system daemon log files
<code>pi m.log. host_name</code>	<code>mbatchd.log. host_name</code>
<code>res.log. host_name</code>	<code>sbatchd.log. host_name</code>
<code>li m.log. host_name</code>	<code>mbschd.log. host_name</code>

If `EGO_LOGDIR` is defined in `ego.conf`, file `li m.log. host_name` is stored in the directory defined by `EGO_LOGDIR`.

Configuration files

`lsf.conf`, `lsf.shared`, and `lsf.cluster. cluster_name` are located in `LSF_CONFDIR`.

`lsb.params`, `lsb.queues`, `lsb.modules`, and `lsb.resources` are located in `LSB_CONFDIR/cluster_name/confidir/`.

File	Description
<code>install.config</code>	Options for Platform LSF installation and configuration
<code>lsf.conf</code>	Generic environment configuration file describing the configuration and operation of the cluster
<code>lsf.shared</code>	Definition file shared by all clusters. Used to define cluster name, host types, host models and site-defined resources
<code>lsf.cluster. cluster_name</code>	Cluster configuration files used to define hosts, administrators, and locality of site-defined shared resources
<code>lsf.licensescheduler</code>	Configures Platform LSF License Scheduler
<code>lsb.applications</code>	Defines application profiles to define common parameters for the same types of jobs
<code>lsb.params</code>	Configures LSF batch parameters
<code>lsb.queues</code>	Batch queue configuration file
<code>lsb.resources</code>	Configures resource allocation limits, exports, and resource usage limits
<code>lsb.serviceclasses</code>	Defines service-level agreements (SLAs) in an LSF cluster as service classes, which define the properties of the SLA
<code>lsb.users</code>	Configures user groups, hierarchical fairshare for users and user groups, and job slot limits for users and user groups

Cluster configuration parameters (lsf.conf)

Variable	Description	UNIX Default
LSF_BINDIR	Directory containing LSF user commands, shared by all hosts of the same type	LSF_TOP/version/platform/bin
LSF_CONFDIR	Directory for all LSF configuration files	LSF_TOP/conf
LSF_ENVDIR	Directory containing the lsf.conf file, must be owned by root	/etc (if LSF_CONFDIR is not defined)
LSF_INCLUDEDIR	Directory containing LSF API header files lsf.h and lsbatch.h	LSF_TOP/version/include
LSF_LIBDIR	LSF libraries, shared by all hosts of the same type	LSF_TOP/version/platform/lib
LSF_LOGDIR	(Optional) Directory for LSF daemon logs, must be owned by root	/tmp
LSF_LOG_MASK	Specifies the logging level of error messages from LSF commands	LOG_WARNING
LSF_MANDIR	Directory containing LSF man pages	LSF_TOP/version/man
LSF_MISC	Help files for the LSF GUI tools, sample C programs and shell scripts, and a template for an external LIM (elim)	LSF_TOP/version/misc
LSF_SERVERDIR	Directory for all server binaries and shell scripts, and external executables invoked by LSF daemons, must be owned by root, and shared by all hosts of the same type	LSF_TOP/version/platform/etc
LSF_TOP	Top-level installation directory. The path to LSF_TOP must be shared and accessible to all hosts in the cluster. It cannot be the root directory (/).	Not defined Required for installation
LSF_CONFDIR	Directory for LSF Batch configuration directories, containing user and host lists, operation parameters, and batch queues	LSF_CONFDIR/lsbatch
LSF_LIVE_CONFDIR	Directory for LSF live reconfiguration directories written by the bconf command.	LSF_TOP/logdir
LSF_SHAREDIR	Directory for LSF Batch job history and accounting log files for each cluster, must be owned by primary LSF administrator	LSF_TOP/work
LSF_LIM_PORT	TCP service port used for communication with lim	7879
LSF_RES_PORT	TCP service port used for communication with res	6878
LSF_MBD_PORT	TCP service port used for communication with mbatchd	6881
LSF_SBD_PORT	TCP service port used for communication with sbatchd	6882

Administration and accounting commands

Only LSF administrators and root can use these commands.

Command	Description
<code>lsadmin</code>	LSF administrative tool to control the operation of the LIM and RES daemons in an LSF cluster, <code>lsadmin help</code> shows all subcommands
<code>lsinstall</code>	Install LSF using <code>install.config</code> input file
<code>lsfrestart</code>	Restart the LSF daemons on all hosts in the local cluster
<code>lsfshutdown</code>	Shut down the LSF daemons on all hosts in the local cluster
<code>lsfstartup</code>	Start the LSF daemons on all hosts in the local cluster
<code>badmin</code>	LSF administrative tool to control the operation of the LSF Batch system including <code>sbatchd</code> , <code>mbatchd</code> , hosts and queues, <code>badmin help</code> shows all subcommands
<code>bladmin</code>	Reconfigures the Platform LSF License Scheduler daemon (<code>bld</code>)
<code>bconf</code>	Changes LSF configuration in active memory

Daemons

Executable Name	Description
<code>lim</code>	Load Information Manager (LIM) — collects load and resource information about all server hosts in the cluster and provides host selection services to applications through LSLIB. LIM maintains information on static system resources and dynamic load indices
<code>mbatchd</code>	Master Batch Daemon (MBD) — accepts and holds all batch jobs. MBD periodically checks load indices on all server hosts by contacting the Master LIM.
<code>mbschd</code>	Master Batch Scheduler Daemon — performs the scheduling functions of LSF and sends job scheduling decisions to MBD for dispatch. Runs on the LSF master server host
<code>sbatchd</code>	Slave Batch Daemon (SBD) — accepts job execution requests from MBD, and monitors the progress of jobs. Controls job execution, enforces batch policies, reports job status to MBD, and launches MBD.
<code>pim</code>	Process Information Manager (PIM) — monitors resources used by submitted jobs while they are running. PIM is used to enforce resource limits and load thresholds, and for fairshare scheduling
<code>res</code>	Remote Execution Server (RES) — accepts remote execution requests from all load sharing applications and handles I/O on the remote host for load sharing processes.

User commands

Viewing information about your cluster.

Command	Description
bhost s	Displays hosts and their static and dynamic resources
bl i mi t s	Displays information about resource allocation limits of running jobs
bparams	Displays information about tunable batch system parameters
bqueues	Displays information about batch queues
busers	Displays information about users and user groups
lshosts	Displays hosts and their static resource information
l si d	Displays the current LSF version number, cluster name and master host name
l si nfo	Displays load sharing configuration information
l sl oad	Displays dynamic load indices for hosts

Monitoring jobs and tasks.

Command	Description
bacct	Reports accounting statistics on completed LSF jobs
bapp	Displays information about jobs attached to application profiles
bhi st	Displays historical information about jobs
bj obs	Displays information about jobs
bpeek	Displays stdout and stderr of unfinished jobs
bsl a	Displays information about service class configuration for goal-oriented service-level agreement scheduling
bst at us	Reads or sets external job status messages and data files

Submitting and controlling jobs.

Command	Description
bbot	Moves a pending job relative to the last job in the queue
bchkpnt	Checkpoints a checkpointable job
bki ll	Sends a signal to a job
bmi g	Migrates a checkpointable or rerunnable job
bmod	Modifies job submission options
brequeue	Kills and requeues a job

Command	Description
<code>bresize</code>	Releases slots and cancels pending job resize allocation requests
<code>brestart</code>	Restarts a checkpointed job
<code>bresume</code>	Resumes a suspended job
<code>bstop</code>	Suspends a job
<code>bsub</code>	Submits a job
<code>bswitch</code>	Moves unfinished jobs from one queue to another
<code>btop</code>	Moves a pending job relative to the first job in the queue

bsub command

Selected options for `bsub [options] command[arguments]`

Option	Description
<code>-ar</code>	Specifies the job is autoresizable
<code>-H</code>	Holds the job in the PSUSP state at submission
<code>-I -Ip -Is</code>	Submits a batch interactive job. <code>-Ip</code> creates a pseudo-terminal. <code>-Is</code> creates a pseudo-terminal in shell mode.
<code>-K</code>	Submits a job and waits for the job to finish
<code>-r</code>	Makes a job rerunnable
<code>-x</code>	Exclusive execution
<code>-app application_profile_name</code>	Submits the job to the specified application profile
<code>-b begin_time</code>	Dispatches the job on or after the specified date and time in the form <code>[[month:]day:]minute</code>
<code>-C core_limit</code>	Sets a per-process (soft) core file size limit (KB) for all the processes that belong to this job
<code>-c cpu_time[/host_name /host_model]</code>	Limits the total CPU time the job can use. CPU time is in the form <code>[hour:]minutes</code>
<code>-cwd "current_working_directory"</code>	Specifies the current working directory for the job
<code>-D data_limit</code>	Sets the per-process (soft) data segment size limit (KB) for each process that belongs to the job
<code>-E "pre_exec_command [arguments]"</code>	Runs the specified pre-exec command on the execution host before running the job
<code>-Ep "post_exec_command [arguments]"</code>	Runs the specified post-exec command on the execution host after the job finishes
<code>-e error_file</code>	Appends the standard error output to a file

Option	Description
- eo <i>error_file</i>	Overwrites the standard error output of the job to the specified file
- F <i>file_limit</i>	Sets per-process (soft) file size limit (KB) for each process that belongs to the job
- f " <i>local_file op[remote_file]</i> " ...	Copies a file between the local (submission) host and remote (execution) host. <i>op</i> is one of >, <, <<, >>, <>
- i <i>input_file</i> - i s <i>input_file</i>	Gets the the standard input for the job from specified file
- J " <i>job_name[index_list]%job_slot_limit</i> "	Assigns the specified name to the job. Job array <i>index_list</i> has the form <i>start [-end[:step]]</i> , and <i>%job_slot_limit</i> is the maximum number of jobs that can run at any given time.
- k " <i>chkpnt_dir [chkpnt_period] [method=method_name]</i> "	Makes a job checkpointable and specifies the checkpoint directory, period in minutes, and method
- M <i>mem_limit</i>	Sets the per-process (soft) memory limit (KB)
- m " <i>host_name [@cluster_name][[!]] +[pref_level] host_group[[!]] +[pref_level] compute_unit[[!]] +[pref_level]]...</i> "	Runs job on one of the specified hosts. Plus (+) after the names of a host or group indicates a preference. Optionally, a positive integer indicates a preference level with higher numbers indicating a greater preference.
- n <i>min_proc[,max_proc]</i>	Specifies the minimum and maximum numbers of processors required for a parallel job
- o <i>output_file</i>	Appends the standard output to a file
- oo <i>output_file</i>	Overwrites the standard output of the job to the specified file
- p <i>process_limit</i>	Limit the number of processes for the whole job
- q " <i>queue_name ...</i> "	Submits job to one of the specified queues
- R " <i>res_req</i> " [-R " <i>res_req</i> " ...]	Specifies host resource requirements
- S <i>stack_limit</i>	Sets a per-process (soft) stack segment size limit (KB) for each process that belongs to the job
- s a <i>service_class_name</i>	Specifies the service class where the job is to run
- T <i>thread_limit</i>	Sets the limit of the number of concurrent threads for the whole job
- t <i>term_time</i>	Specifies the job termination deadline in the form <i>[[month:]day:]hour:minute</i>
- v <i>swap_limit</i>	Sets the total process virtual memory limit (KB) for the whole job
- W <i>run_time[/host_name /host_mode]</i>	Sets the run time limit of the job in the form <i>[hour:]minute</i>
- h	Prints command usage to stderr and exits
- V	Prints LSF release version to stderr and exits