

Platform LSF
Version 9 Release 1.1

Quick Reference



Platform LSF
Version 9 Release 1.1

Quick Reference



Note

Before using this information and the product it supports, read the information in "Notices" on page 9.

First edition

This edition applies to version 9, release 1 of IBM Platform LSF (product number 5725G82) and to all subsequent releases and modifications until otherwise indicated in new editions.

© Copyright IBM Corporation 1992, 2013.

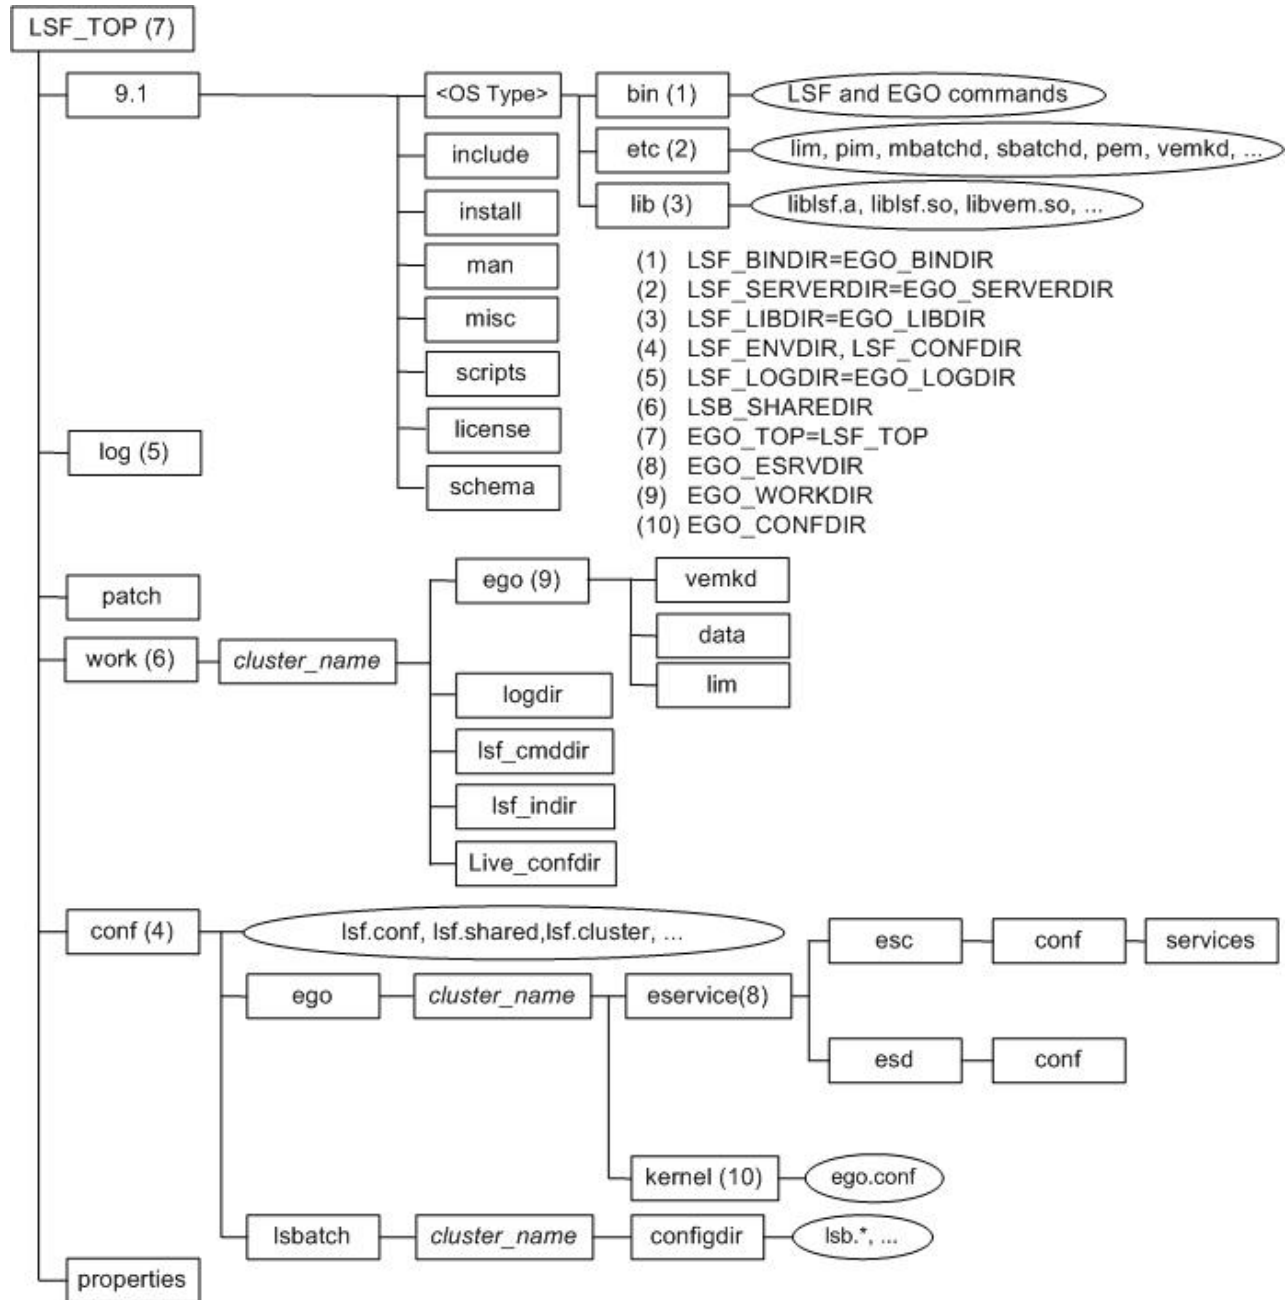
US Government Users Restricted Rights – Use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM Corp.

Contents

Notices	9	Trademarks	11
-------------------	---	----------------------	----

IBM Platform LSF 9.1.1 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>pim.log.host_name</code>	<code>mbatchd.log.host_name</code>
<code>res.log.host_name</code>	<code>sbatchd.log.host_name</code>
<code>lim.log.host_name</code>	<code>mbschd.log.host_name</code>

If **EGO_LOGDIR** is defined in `ego.conf`, file `lim.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/configdir/**.

File	Description
<code>install.config</code>	Options for 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>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
<code>LSF_BINDIR</code>	Directory containing LSF user commands, shared by all hosts of the same type	<code>LSF_TOP/version/platform/bin</code>

Variable	Description	UNIX Default
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
LSB_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.	LSB_SHAREDIR/<cluster_name>/live_confdir
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

Variable	Description	UNIX Default
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 users can use these commands.

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

Daemons

Executable Name	Description
lim	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
mbatchd	Master Batch Daemon (MBD) — accepts and holds all batch jobs. MBD periodically checks load indices on all server hosts by contacting the Master LIM.

Executable Name	Description
mbschd	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
sbatchd	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.
pim	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
res	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
bhosts	Displays hosts and their static and dynamic resources
blimits	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
lsid	Displays the current LSF version number, cluster name and master host name
lsinfo	Displays load sharing configuration information
lsload	Displays dynamic load indices for hosts

Monitoring jobs and tasks.

Command	Description
bacct	Reports accounting statistics on completed LSF jobs

Command	Description
bapp	Displays information about jobs attached to application profiles
bhist	Displays historical information about jobs
bjobs	Displays information about jobs
bpeek	Displays stdout and stderr of unfinished jobs
bsla	Displays information about service class configuration for goal-oriented service-level agreement scheduling
bstatus	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
bkill	Sends a signal to a job
bmig	Migrates a checkpointable or rerunnable job
bmod	Modifies job submission options
brequeue	Kills and requeues a job
bresize	Releases slots and cancels pending job resize allocation requests
brestart	Restarts a checkpointed job
bresume	Resumes a suspended job
bstop	Suspends a job
bsub	Submits a job
bswitch	Moves unfinished jobs from one queue to another
btop	Moves a pending job relative to the first job in the queue

bsub command

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

Option	Description
-ar	Specifies the job is autoresizable
-H	Holds the job in the PSUSP state at submission

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

Option	Description
-m " <i>host_name</i> [<i>@cluster_name</i>][<i>!</i>] <i>+pref_level</i>] <i>host_group</i> [<i>!</i>] <i>+pref_level</i>] <i>compute_unit</i> [<i>!</i>] <i>+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</i> , <i>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
-s1a <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:</i>] <i>day:</i>] <i>hour:minute</i>
-v <i>swap_limit</i>	Sets the total process virtual memory limit (KB) for the whole job
-W <i>run_time</i> [/ <i>host_name</i> / <i>host_model</i>]	Sets the run time limit of the job in the form [<i>hour:</i>] <i>minute</i>
-h	Prints command usage to stderr and exits
-V	Prints LSF release version to stderr and exits

Notices

This information was developed for products and services offered in the U.S.A.

IBM® may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not grant you any license to these patents. You can send license inquiries, in writing, to:

IBM Director of Licensing
IBM Corporation
North Castle Drive
Armonk, NY 10504-1785
U.S.A.

For license inquiries regarding double-byte character set (DBCS) information, contact the IBM Intellectual Property Department in your country or send inquiries, in writing, to:

Intellectual Property Licensing
Legal and Intellectual Property Law
IBM Japan Ltd.
19-21, Nihonbashi-Hakozakicho, Chuo-ku
Tokyo 103-8510, Japan

The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law:

INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some states do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web

sites. The materials at those Web sites are not part of the materials for this IBM product and use of those Web sites is at your own risk.

IBM may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

Licensees of this program who wish to have information about it for the purpose of enabling: (i) the exchange of information between independently created programs and other programs (including this one) and (ii) the mutual use of the information which has been exchanged, should contact:

IBM Corporation
Intellectual Property Law
Mail Station P300
2455 South Road,
Poughkeepsie, NY 12601-5400
USA

Such information may be available, subject to appropriate terms and conditions, including in some cases, payment of a fee.

The licensed program described in this document and all licensed material available for it are provided by IBM under terms of the IBM Customer Agreement, IBM International Program License Agreement or any equivalent agreement between us.

Any performance data contained herein was determined in a controlled environment. Therefore, the results obtained in other operating environments may vary significantly. Some measurements may have been made on development-level systems and there is no guarantee that these measurements will be the same on generally available systems. Furthermore, some measurement may have been estimated through extrapolation. Actual results may vary. Users of this document should verify the applicable data for their specific environment.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

All statements regarding IBM's future direction or intent are subject to change or withdrawal without notice, and represent goals and objectives only.

This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to the names and addresses used by an actual business enterprise is entirely coincidental.

COPYRIGHT LICENSE:

This information contains sample application programs in source language, which illustrates programming techniques on various operating platforms. You may copy, modify, and distribute these sample programs in any form without payment to IBM, for the purposes of developing, using, marketing or distributing application

programs conforming to the application programming interface for the operating platform for which the sample programs are written. These examples have not been thoroughly tested under all conditions. IBM, therefore, cannot guarantee or imply reliability, serviceability, or function of these programs. The sample programs are provided "AS IS", without warranty of any kind. IBM shall not be liable for any damages arising out of your use of the sample programs.

Each copy or any portion of these sample programs or any derivative work, must include a copyright notice as follows:

© (your company name) (year). Portions of this code are derived from IBM Corp. Sample Programs. © Copyright IBM Corp. _enter the year or years_.

If you are viewing this information softcopy, the photographs and color illustrations may not appear.

Trademarks

IBM, the IBM logo, and [ibm.com](http://www.ibm.com)[®] are trademarks of International Business Machines Corp., registered in many jurisdictions worldwide. Other product and service names might be trademarks of IBM or other companies. A current list of IBM trademarks is available on the Web at "Copyright and trademark information" at <http://www.ibm.com/legal/copytrade.shtml>.

Intel, Intel logo, Intel Inside, Intel Inside logo, Intel Centrino, Intel Centrino logo, Celeron, Intel Xeon, Intel SpeedStep, Itanium, and Pentium are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.

Java[™] and all Java-based trademarks and logos are trademarks or registered trademarks of Oracle and/or its affiliates.

Linux is a trademark of Linus Torvalds in the United States, other countries, or both.

LSF[®], Platform, and Platform Computing are trademarks or registered trademarks of International Business Machines Corp., registered in many jurisdictions worldwide.

Microsoft, Windows, Windows NT, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.

Other company, product, or service names may be trademarks or service marks of others.



Printed in USA

GC27-5309-01

