

---

# Release Notes for Platform™ LSF™ Version 7 Update 2

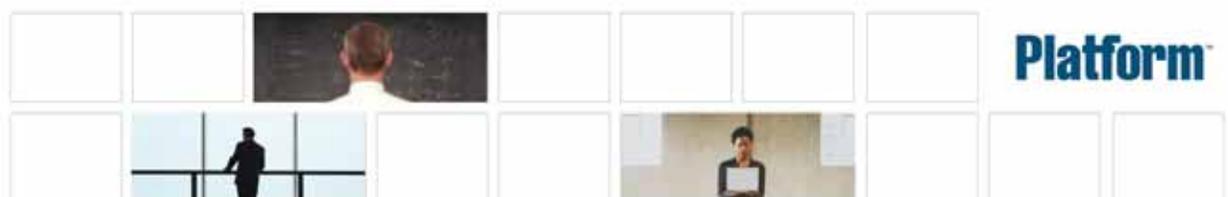
Release date: November 2007  
Last modified: February 20, 2008  
Comments to: [doc@platform.com](mailto:doc@platform.com)  
Support: [support@platform.com](mailto:support@platform.com)

## Contents

- ◆ [Upgrade and Compatibility Notes](#) on page 2
- ◆ [What's Changed in Platform LSF Version 7 Update 2](#) on page 4
- ◆ [Known Issues](#) on page 33
- ◆ [Download the Platform LSF Version 7 Distribution Packages](#) on page 36
- ◆ [Install Platform LSF Version 7](#) on page 39
- ◆ [Learn About Platform LSF Version 7](#) on page 41
- ◆ [Get Technical Support](#) on page 42

## What's new in Platform LSF Version 7 Update 2

For detailed information about what's new in Platform LSF Version 7 Update 2, visit the Platform Computing Web site to see *Features, Benefits & What's New*.



## Upgrade and Compatibility Notes

- ◆ [Server host compatibility](#)
- ◆ [Upgrade from an earlier version of LSF on UNIX and Linux](#)
- ◆ [Migrate your existing LSF Version 7 cluster to Version 7 Update 2 on UNIX and Linux \(Update 2 only\)](#)
- ◆ [Migrate LSF on Windows \(Update 2 only\)](#)
- ◆ [Maintenance pack and update availability](#)
- ◆ [System requirements](#)
- ◆ [API compatibility](#)

### Server host compatibility

---

**IMPORTANT:** To use new features introduced in Platform LSF Version 7, you *must* upgrade all hosts in your cluster to LSF 7.

---

LSF 6.x and 5.x servers are compatible with Platform LSF Version 7 master hosts. All LSF 6.x and 5.x features are supported by LSF 7 master hosts.

### Upgrade from an earlier version of LSF on UNIX and Linux

Run `lsfinstall` to *upgrade* to LSF Version 7 from an earlier version of LSF on UNIX and Linux. Follow the steps in *Upgrading Platform LSF on UNIX and Linux*.

---

**IMPORTANT:** *Do not* use the UNIX and Linux upgrade steps to update an existing LSF Version 7 cluster to LSF Version 7 Update 2. Follow the manual steps in the document *Migrating LSF Version 7 to Update 2 on UNIX and Linux* to migrate your existing LSF Version 7 cluster to LSF Version 7 Update 2.

---

### Migrate your existing LSF Version 7 cluster to Version 7 Update 2 on UNIX and Linux (Update 2 only)

For LSF 7 Update 2 only, you *cannot* use `lsfinstall` to *migrate* your existing LSF Version 7 cluster to LSF 7 Update 2. Follow the manual steps in the document *Migrating LSF Version 7 to Update 2 on UNIX and Linux* (`/distrib/7.0/lsf_migrate_unix_linux_to_update2.pdf`) to migrate your existing LSF Version 7 cluster to LSF Version 7 Update 2. This is required due to the change in LSF directory structure for LSF 7 Update 2. In later LSF updates, `lsfinstall` will automatically handle the update.

### Migrate LSF on Windows (Update 2 only)

To migrate your existing LSF Version 7 cluster on Windows to LSF 7 Update 2, you must follow the manual steps in the document *Migrating LSF Version 7 to Update 2 on Windows* (`/distrib/7.0/lsf_migrate_windows_to_update2.pdf`). This is required due to the change in LSF directory structure for LSF 7 Update 2 only.

## Maintenance pack and update availability

At release, Platform LSF Version 7 Update 2 includes all bug fixes and solutions up to and including bug fixes and solutions before November, 2007. Fixes after November 2007 will be included in the next LSF update.

Fixes in the November 2006 maintenance pack were included in the March 2007 update.

As of February 2007, monthly maintenance packs are no longer distributed for LSF Version 7.

## System requirements

See the Platform Computing Web site for information about supported operating systems and system requirements for the Platform LSF family of products:

- ◆ Platform LSF
- ◆ Platform LSF License Scheduler

## API compatibility

Full backward compatibility: your applications will run under LSF Version 7 without changing any code.

The Platform LSF Version 7 API is fully compatible with the LSF Version 6.x. and 5.x APIs. An application linked with the LSF Version 6.x or 5.x libraries will run under LSF Version 7 without relinking.

To take full advantage of new Platform LSF Version 7 features, you should recompile your existing LSF applications with LSF Version 7.

## New and changed LSF APIs

See the *LSF API Reference* for more information.

No new APIs have been added for LSF Version 7 Update 2.

The following APIs have changed for LSF Version 7 Update 2:

- ◆ `lsb_modify()`, `lsb_submit()`, `lsb_readjobinfo()`, `lsb_readjobinfo_cond()`, `lsb_jsdl2submit()`, and `lsb_submitframe()`—add the following fields to the submit data structure:
  - ❖ `cwd`
  - ❖ `postExecCmd`
  - ❖ `runtimeEstimation`
- ◆ `lsb_geteventrec()` and `lsb_puteventrec()`—add the following fields to the `jobFinishLog`, `jobModLog`, and `jobNewLog` data structures:
  - ❖ `postExecCmd`
  - ❖ `runtimeEstimation`

## What's Changed in Platform LSF Version 7 Update 2

- ◆ LSF directory structure change
- ◆ Enabling and disabling Platform EGO in LSF
- ◆ Changed behavior
- ◆ New and changed configuration parameters and environment variables
- ◆ New and changed commands, options, and output
- ◆ New and changed files
- ◆ New and changed accounting and job event fields
- ◆ Bugs fixed since June 2007

### LSF directory structure change

The installed directory structure has changed for LSF Version 7 Update 2. The LSF directory structure differs from the directory structure of previous releases of LSF Version 7. The LSF directories have been consolidated to resemble the LSF 6.x directory structure, where configuration, work, and logging directories for Platform EGO, Web GUI, and reporting components are located under the standard LSF directories `LSF_TOP/conf`, `LSF_TOP/work`, and `LSF_TOP/log` for easier administration and upgrade.

These directories can be located on their own file systems and have the following access permissions:

- ◆ `LSF_TOP/conf` is writable by the LSF administrator, master host, and master candidate hosts
- ◆ `LSF_TOP/log` is writable by all hosts in the cluster
- ◆ `LSF_TOP/work` is writable by the master host and master candidate hosts, and is accessible to slave hosts

See the chapters “Working with Your Cluster” and “Managing LSF on Platform EGO” in *Administering Platform LSF* for detailed information.

### Enabling and disabling Platform EGO in LSF

You now can choose to enable or disable EGO functionality in LSF Version 7 Update 2. By default, the following LSF features that depend on EGO is enabled:

- ◆ LSF daemon control by EGO Service Controller
- ◆ EGO-enabled SLA scheduling
- ◆ Platform Management Console (PMC)
- ◆ LSF reporting

---

**IMPORTANT:** EGO must be enabled to use EGO commands and the PMC and LSF reporting features.

---

If you do not want to enable EGO functionality at installation, define `ENABLE_EGO="N"` in `install.config`. After installation, use `LSF_ENABLE_EGO` in `lsf.conf` to enable or disable EGO functionality. See the “Managing LSF on Platform EGO” chapter in *Administering Platform LSF* for detailed steps.

---

**IMPORTANT:** After changing the value of `LSF_ENABLE_EGO` in `lsf.conf`, you must shut down and restart the cluster.

---

## lsf.conf parameters corresponding to ego.conf parameters

When Platform EGO is enabled in LSF Version 7, you can configure some LSF parameters in `lsf.conf` that have corresponding EGO parameter names in `EGO_CONFDIR/ego.conf` (`LSF_CONFDIR/lsf.conf` is a separate file from `EGO_CONFDIR/ego.conf`). If both the LSF and the EGO parameters are set in their respective files, the definition in `ego.conf` is used. You must continue to set LSF parameters only in `lsf.conf`.

When EGO is enabled in the LSF cluster (`LSF_ENABLE_EGO=Y` in `lsf.conf`), you also can set the following EGO parameters related to LIM, PIM, and ELIM in either `lsf.conf` or `ego.conf`:

- ◆ `EGO_DAEMONS_CPUS`
- ◆ `EGO_DEFINE_NCPUS`
- ◆ `EGO_DISABLE_UNRESOLVABLE_HOSTS` (applies only to dynamically added hosts)
- ◆ `EGO_SLAVE_CTRL_REMOTE_HOST`
- ◆ `EGO_WORKDIR`
- ◆ `EGO_PIM_SWAP_REPORT`

If EGO is not enabled, you can set these parameters only in `lsf.conf`.

See *Administering Platform LSF* for more information about configuring LSF for EGO. See the *Platform EGO Reference* for information about `ego.conf` parameters.

## Changed behavior

### Multithreaded mbatchd

If you have enabled multithreaded `mbatchd` support, the `bjobs` command may not display up-to-date information if two consecutive query commands are issued before a child `mbatchd` expires because child `mbatchd` job information is not updated. Use `NEWJOB_REFRESH=Y` in `lsb.params` to enable a child `mbatchd` to get up to date new job information from the parent `mbatchd`.

When `NEWJOB_REFRESH=Y` the parent `mbatchd` pushes new job information to a child `mbatchd`. Job queries with `bjobs` display new jobs submitted after the child `mbatchd` was created. `LSB_QUERY_PORT` must be defined in `lsf.conf`.

### Temporary directory for tasks launched by blaunch

By default, LSF creates a temporary directory for a job only on the first execution host. If `LSF_TMPDIR` is set in `lsf.conf` file, the path of the job temporary directory on the first execution host is set to `LSF_TMPDIR/job_ID.tmpdir`.

If `LSB_SET_TMPDIR= Y`, the environment variable `TMPDIR` will be set equal to the path specified by `LSF_TMPDIR`. This value for `TMPDIR` overrides any value that might be set in the submission environment.

Tasks launched through the `b1aunch` distributed application framework make use of the LSF temporary directory specified by `LSF_TMPDIR`:

- ◆ When the environment variable `TMPDIR` is set on the first execution host, the `b1aunch` framework propagates this environment variable to all execution hosts when launching remote tasks  
The job `RES` or the task `RES` creates the directory specified by `TMPDIR` if it does not already exist before starting the job
- ◆ The directory created by the job `RES` or task `RES` has permission `0700` and is owned by the execution user
- ◆ If the `TMPDIR` directory was created by the task `RES`, LSF deletes the temporary directory and its contents when the task is complete
- ◆ If the `TMPDIR` directory was created by the job `RES`, LSF will delete the temporary directory and its contents when the job is done
- ◆ If the `TMPDIR` directory is on a shared file system, it is assumed to be shared by all the hosts allocated to the `b1aunch` job, so LSF does not remove `TMPDIR` directories created by the job `RES` or task `RES`

### Preemption of backfill jobs

With preemption of backfill jobs enabled (`PREEMPT_JOBTYPE=BACKFILL` in `lsb.params`), LSF maintains the priority of jobs with resource or slot reservations by preventing lower-priority jobs that preempt backfill jobs from "stealing" resources from jobs with reservations. Only jobs from queues with a higher priority than queues that define resource or slot reservations can preempt backfill jobs.

To guarantee a minimum run time for interruptible backfill jobs, LSF suspends them upon preemption. To change this behavior so that LSF terminates interruptible backfill jobs upon preemption, you must define the parameter `TERMINATE_WHEN=PREEMPT` in `lsb.queues`.

### `bsub -J` job name length

Options related to command names and job names can contain up to 4094 characters for UNIX and Linux, or up to 255 characters for Windows.

### Job slot limits

Job slot limits can correspond to the maximum number of jobs that can run at any point in time. For example, a queue cannot start jobs if it has no job slots available, and jobs cannot run on hosts that have no available job slots.

Limits such as `QJOB_LIMIT` (`lsb.queues`), `HJOB_LIMIT` (`lsb.queues`), `UJOB_LIMIT` (`lsb.queues`), `MXJ` (`lsb.hosts`), `JL/U` (`lsb.hosts`), `MAX_JOBS` (`lsb.users`), and `MAX_PEND_JOBS` (`lsb.users`) limit the number of job slots. When the workload is sequential, job slots are usually equivalent to jobs. For parallel or distributed applications, these are true job slot limits and not job limits.

### Job limits

Job limits, specified by `JOBS` in a Limit section in `lsb.resources`, correspond to the maximum number of running and suspended jobs that can run at any point in time. If both job limits and job slot limits are configured, the most restrictive limit is applied.

## Job group limits

Job group limits apply to the job group hierarchy. The job group limit is a positive number greater than or equal to zero (0), specifying the maximum the number of running and suspended jobs under the job group (including child groups). If limit is zero (0), no jobs under the job group can run.

By default, a job group has no limit. Limits persist across `mbatchd` restart and reconfiguration.

You cannot specify a limit for the root job group. The root job group has no job limit. Job groups added with no limits specified inherit any limits of existing parent job groups. The `-L` option only limits the lowest level job group created.

The maximum number of running and suspended jobs (including USSUP and SSUP) in a job group cannot exceed the limit defined on the job group and its parent job group.

The job group limit is based on the number of running and suspended jobs in the job group. If you specify a job group limit as 2, at most 2 jobs can run under the group at any time, regardless of how many jobs or job slots are used. If the currently available job slots is zero (0), even if the job group job limit is not exceeded, LSF cannot dispatch a job to the job group.

If a parallel job requests 2 CPUs (`bsub -n 2`), the job group limit is per job, not per slots used by the job.

A job array may also be under a job group, so job arrays also support job group limits.

Job group limits are not supported at job submission for job groups created automatically with `bsub -g`. Use `bgadd -L` before job submission.

## External job submission and execution controls

LSF 7 Update 2 adds the following environment variables in the `esub` execution environment:

- ◆ `LSB_SUB_RERUNNABLE`—"Y" specifies a rerunnable job. "N" specifies a nonrerunnable job (specified with `bsub -rn`). The job is not rerunnable even it was submitted to a rerunable queue or application profile. For `bmod -rn`, the value is `SUB_RESET`.
- ◆ `LSB_SUB2_JOB_PRIORITY`—Job priority (`bsub -sp` and `bmod -sp`). For `bmod -spn`, the value is `SUB_RESET`.
- ◆ `LSB_SUB3_CWD`—Current working directory specified on on the command line with `bsub -cwd`.
- ◆ `LSB_SUB3_POST_EXEC`—Run the specified post-execution command on the execution host after the job finishes. Specified by `bsub -Ep`.
- ◆ `LSB_SUB3_RUNTIME_ESTIMATION`—Runtime estimate specified by `bsub -We`.
- ◆ `LSB_SUB3_USER_SHELL_LIMITS`—Pass user shell limits to execution host. Specified by `bsub -u1`.

## Application profiles

Application profile configuration has been enhanced:

### Pre-execution

Queue-level pre-execution commands run *before* application-level pre-execution commands. Job level pre-execution commands (`bsub -E`) override application-level pre-execution commands.

### Post-execution

When a job finishes, application-level post-execution commands run, followed by queue-level post-execution commands if any.

If both application-level and job-level post-execution commands (`bsub -Ep`) are specified, job level post-execution overrides application-level post-execution commands. Queue-level post-execution commands run after application-level post-execution and job-level post-execution commands

### Estimated runtime and runtime limits

Instead of specifying an explicit runtime limit for jobs, you can specify an *estimated* run time for jobs. LSF uses the estimated value for job scheduling purposes only, and does not kill jobs that exceed this value unless the jobs also exceed a defined runtime limit. The format of runtime estimate is same as run limit set by the `bsub -W` option or the `RUNLIMIT` parameter in `lsb.queues` and `lsb.applications`.

Use `JOB_RUNLIMIT_RATIO` in `lsb.params` to limit the runtime estimate users can set. If `JOB_RUNLIMIT_RATIO` is set to 0 no restriction is applied to the runtime estimate. The ratio does not apply to the `RUNTIME` parameter in `lsb.applications`.

The job-level runtime estimate setting overrides the `RUNTIME` setting in an application profile in `lsb.applications`.

The following LSF features use the estimated runtime value to schedule jobs:

- ◆ Job chunking
- ◆ Advanced reservation
- ◆ SLA
- ◆ Slot reservation
- ◆ Backfill

## Default job group

You can specify a default job group for jobs submitted without explicitly specifying a job group. LSF associates the job with the job group specified with `DEFAULT_JOBGROUP` in `lsb.params`. The `LSB_DEFAULT_JOBGROUP` environment variable overrides the setting of `DEFAULT_JOBGROUP`. The `bsub -g job_group_name` option overrides both `LSB_DEFAULT_JOBGROUP` and `DEFAULT_JOBGROUP`.

Default job group specification supports macro substitution for project name (`%p`) and user name (`%u`). When you specify `bsub -P project_name`, the value of `%p` is the specified project name. If you do not specify a project name at job submission, `%p` is the project name defined by setting the environment variable `LSB_DEFAULTPROJECT`, or the project name specified by `DEFAULT_PROJECT` in `lsb.params`. the default project name is `default`.

## Processor binding for LSF job processes

Rapid progress of modern processor manufacture technologies has enabled the low cost deployment of LSF on hosts with multicore and multithread processors. The default soft affinity policy enforced by the operating system scheduler may not give optimal job performance. For example, the operating system scheduler may place all job processes on the same processor or core leading to poor performance. Frequently switching processes as the operating system schedules and reschedules work between cores can cause cache invalidations and cache miss rates to grow large.

Processor binding for LSF job processes takes advantage of the power of multiple processors and multiple cores to provide hard processor binding functionality for sequential LSF jobs.

---

**RESTRICTION:** Processor binding is supported on hosts running Linux with kernel version 2.6 or higher. For parallel jobs, you should use an MPI that natively supports processor bindin: HP MPI, Scali MPI, OpenMPI, etc.

---

LSF does not bind parallel jobs to any processor. If you require this functionality you should use an MPI that directly supports this.

When processor binding for LSF job processes is enabled on supported hosts, job processes of an LSF job are bound to a processor according to the binding policy of the host. When an LSF job is completed (exited or done successfully) or suspended, the corresponding processes are unbound from the processor.

When a suspended LSF job is resumed, the corresponding processes are bound again to a processor. The process is not guaranteed to be bound to the same processor it was bound to before the job was suspended.

The processor binding affects the whole job process group. All job processes forked from the root job process (the job RES) are bound to the same processor.

Processor binding for LSF job processes does not bind daemon processes.

If processor binding is enabled, but the execution hosts do not support processor affinity, the configuration has no effect on the running processes. Processor binding has no effect on a single-processor host.

## Processor, core, and thread CPU load balancing

By default, the number of CPUs on a host represents the number of physical processors a machine has. For LSF hosts with multiple cores, threads, and processors, `n_cpus` can be defined by the cluster administrator to consider one of the following:

- ◆ Processors
- ◆ Processors and cores
- ◆ Processors, cores, and threads

Globally, this definition is controlled by the parameter `EGO_DEFINE_NCPUS` in `ego.conf` or `lsf.conf`. The default behavior for `n_cpus` is to consider only the number of physical processors (equivalent to `EGO_DEFINE_NCPUS=procs`).

Binding job processes randomly to multiple processors, cores, or threads, may affect job performance. Processor binding configured with `LSF_BIND_JOB` in `lsf.conf` or `BIND_JOB` in `lsb.applications`, detects the `EGO_DEFINE_NCPUS` policy to load balance the job processes by processor, core, or thread (PCT).

### Feature Interactions

- ◆ Existing CPU affinity features
  - Processor binding of LSF job processes will not take effect on a master host with the following parameters configured.
    - ❖ `MBD_QUERY_CPUS`
    - ❖ `LSF_DAEMONS_CPUS`
    - ❖ `EGO_DAEMONS_CPUS`
- ◆ Altix cpusets
  - Processor binding cannot be used with Altix cpusets. If an execution host is configured as part of a cpuset, processor binding is disabled on that host.
- ◆ Job requeue, rerun, and migration
  - When a job is requeued, rerun or migrated, a new job process is created. If processor binding is enabled when the job runs, the job processes will be bound to a processor.
- ◆ `badadmin hrestart`
  - `badadmin hrestart` restarts a new `sbatchd`. If a job process has already been bound to a processor, after `sbatchd` is restarted, processor binding for the job processes are restored.
- ◆ `badadmin reconfig`
  - If the `BIND_JOB` parameter is modified in an application profile, `badadmin reconfig` only affects pending jobs. The change does not affect running jobs.
- ◆ MultiCluster job forwarding model
  - In a MultiCluster environment, the behavior is similar to the current application profile behavior. If the application profile name specified in the submission cluster is not defined in the execution cluster, the job is rejected. If the execution cluster has the same application profile name, but does not enable processor binding, the job processes are not bound at the execution cluster.

### Increasing the job ID limit

By default, LSF assigns job IDs up to 6 digits. This means that no more than 999999 jobs can be in the system at once. The job ID limit is the highest job ID that LSF will ever assign, and also the maximum number of jobs in the system.

LSF assigns job IDs in sequence. When the job ID limit is reached, the count rolls over, so the next job submitted gets job ID "1". If the original job 1 remains in the system, LSF skips that number and assigns job ID "2", or the next available job ID. If you have so many jobs in the system that the low job IDs are still in use when the maximum job ID is assigned, jobs with sequential numbers could have different submission times.

### Increase the maximum job ID

You cannot lower the job ID limit, but you can raise it to 10 digits. This allows longer term job accounting and analysis, and means you can have more jobs in the system, and the job ID numbers will roll over less often.

Use `MAX_JOBID` in `lsb.params` to specify any integer from 999999 to 2147483646 (for practical purposes, you can use any 10-digit integer less than this value).

### Increase the job ID display length

By default, `bjobs` and `bhist` display job IDs with a maximum length of 7 characters. Job IDs greater than 9999999 are truncated on the left.

Use `LSB_JOBID_DISP_LENGTH` in `lsf.conf` to increase the width of the `JOBID` column in `bjobs` and `bhist` display. When `LSB_JOBID_DISP_LENGTH=10`, the width of the `JOBID` column in `bjobs` and `bhist` increases to 10 characters.

## LSF command messages

LSF daemon (LIM) not responding ... still trying

During LIM restart, LSF commands will fail and display this error message. User programs linked to the LIM API will also fail for the same reason. This message is displayed when LIM running on the master host list or server host list is restarted after configuration changes, such as adding new resources, binary upgrade, and so on.

Use `LSF_LIM_API_NTRIES` in `lsf.conf` or as an environment variable to define how many times LSF commands will retry to communicate with the LIM API while LIM is not available. `LSF_LIM_API_NTRIES` is ignored by LSF and EGO daemons (`lim`, `pim`, `pem`, and `vemkd`) and all EGO commands.

When `LSB_API_VERBOSE=Y` in `lsf.conf`, LSF batch commands will display the not responding retry error message to `stderr` when LIM is not available.

When `LSB_API_VERBOSE=N` in `lsf.conf`, LSF batch commands will not display the retry error message when LIM is not available.

## Viewing daemon parameter configuration

- 1 Display all configuration settings for running LSF daemons.
  - ❖ Use `lsadmin showconf` to display all configured parameters and their values in `lsf.conf` or `ego.conf` for LIM.
  - ❖ Use `badmin showconf` to display all configured parameters and their values in `lsb.params` and `lsf.conf` for `mbatchd` and `sbatchd`.

In a MultiCluster environment, `badmin showconf` only displays the parameters of daemons on the local cluster.

Running `badmin showconf` and `lsadmin showconf` from a master candidate host will reach all server hosts in the cluster. Running `badmin showconf` and `lsadmin showconf` from a slave-only host may not be able to reach other slave-only hosts.

`badmin showconf` and `lsadmin showconf` only display the values used by LSF. For example, if you define `LSF_MASTER_LIST` in `lsf.conf`, and `EGO_MASTER_LIST` in `ego.conf`, `badmin showconf` and `lsadmin showconf` display the value of `EGO_MASTER_LIST`.

- 2 Display `mbatchd` and root `sbatchd` configuration.

- ❖ Use `badadmin showconf mbd` to display the parameters configured in `lsf.conf` or `ego.conf` that apply to `mbatchd`.
- ❖ Use `badadmin showconf sbd` to display the parameters configured in `lsf.conf` or `ego.conf` that apply to `root sbatchd`.

### 3 Display LIM configuration.

Use `lsadmin showconf lim` to display the parameters configured in `lsf.conf` or `ego.conf` that apply to root LIM.

By default, `lsadmin` displays the local LIM parameters. You can specify the host to display the LIM parameters.

## LSF reporting data loader plug-ins

The LSF reporting feature adds no new data loader plugins for LSF 7 Update 2.

## New and changed configuration parameters and environment variables

The following configuration parameters and environment variables are new or changed for LSF Version 7:

### `cshrc.lsf` and `profile.lsf`

If Platform EGO is enabled in the LSF cluster (`LSF_ENABLE_EGO=Y` and `LSF_EGO_ENVDIR` are defined in `lsf.conf`), `cshrc.lsf` and `profile.lsf` set the following environment variables.

- ◆ `EGO_BINDIR`
- ◆ `EGO_CONFDIR`
- ◆ `EGO_ESRVDIR`
- ◆ `EGO_LIBDIR`
- ◆ `EGO_LOCAL_CONFDIR`
- ◆ `EGO_SERVERDIR`
- ◆ `EGO_TOP`

See the *Platform EGO Reference* for more information about these variables.

You do not need to use `cshrc.ego` or `profile.ego` to set your LSF environment.

### `ego.conf`

`EGO_DISABLE_UNRESOLVABLE_HOSTS`—Applies only to dynamically added hosts.

### `ego.shared` and `ego.cluster` files

After LSF Version 7 Update 2, the `ego.shared` and `ego.cluster.cluster_name` files are no longer exist in the installation. Platform EGO reads its configuration from `ego.conf`, `lsf.conf`, `lsf.shared`, and `lsf.cluster.cluster_name`.

### `hosts`

The LSF `hosts` file is not installed by default. It is usually located in the directory specified by `LSF_CONFDIR`. The format of `LSF_CONFDIR/hosts` is similar to the format of the `/etc/hosts` file on UNIX machines.

### `install.config` and `slave.config`

- ◆ `DYNAMIC_HOST_WAIT_TIME`—the default is now 60 seconds.
- ◆ `ENABLE_EGO`—Enables Platform EGO functionality in the LSF cluster. `ENABLE_EGO="Y"` (the default) causes `lsfinstall` uncomment `LSF_EGO_ENVDIR` and sets `LSF_ENABLE_EGO="Y"` in `lsf.conf`.

ENABLE\_EGO="N" causes `lsfinstall` to comment out `LSF_EGO_ENVDIR` and sets `LSF_ENABLE_EGO="N"` in `lsf.conf`. Set the value to "N" if you do not want to take advantage of the LSF features that depend on EGO:

- ❖ LSF daemon control by EGO Service Controller
- ❖ EGO-enabled SLA scheduling
- ❖ Platform Management Console (PMC)
- ❖ LSF reporting
- ◆ EGO\_TOP—now the same as LSF\_TOP. You no longer need to define EGO\_TOP in `install.config`,

## lsb.applications

- ◆ BIND\_JOB=Y—Enables LSF job process processor binding for all sequential jobs submitted to the application profile. On Linux execution hosts that support this feature, job processes will be hard bound to selected processors. If processor binding feature is not configured with the BIND\_JOB parameter in an application profile in `lsb.applications`, the `lsf.conf` configuration setting takes effect. The application profile configuration for processor binding overrides the `lsf.conf` configuration. LSF does not bind parallel jobs to any processor. If you require this functionality you should use an MPI that directly supports this.
- ◆ POST\_EXEC—Enables post-execution processing at the application level. The `POST_EXEC` command runs on the execution host after the job finishes. Post-execution commands can be configured at the job, application, and queue levels. If both application-level (`POST_EXEC` in `lsb.applications`) and job-level post-execution commands are specified, job level post-execution overrides application-level post-execution commands. Queue-level post-execution commands (`POST_EXEC` in `lsb.queues`) run after application-level post-execution and job-level post-execution commands.
- ◆ PRE\_EXEC—Enables pre-execution processing at the application level. The `PRE_EXEC` command runs on the execution host before the job starts. If the `PRE_EXEC` command exits with a non-zero exit code, LSF requeues the job to the front of the queue. Pre-execution commands can be configured at the application, queue, and job levels and run in the following order:
  - a The queue-level command
  - b The application-level or job-level command. If you specify a command at both the application and job levels, the job-level command overrides the application-level command; the application-level command is ignored.

The `PRE_EXEC` command uses the same environment variable values as the job, and runs under the user account of the user who submits the job. To run pre-execution commands under a different user account (such as `root` for privileged operations), configure the parameter `LSB_PRE_POST_EXEC_USER` in `lsf.sudoers`.

## lsb.params

- ◆ **DEFAULT\_JOBGROUP**—specifies the name of the default job group. When you submit a job to LSF without explicitly specifying a job group, LSF associates the job with the specified job group. The `LSB_DEFAULT_JOBGROUP` environment variable overrides the setting of `DEFAULT_JOBGROUP`. The `bsub -g job_group_name` option overrides both `LSB_DEFAULT_JOBGROUP` and `DEFAULT_JOBGROUP`.
- ◆ **JOB\_INCLUDE\_POSTPROC**—Specifies whether LSF includes the post-execution processing of the job as part of the job. When set to Y:
  - ❖ Prevents a new job from starting on a host until post-execution processing is finished on that host
  - ❖ Includes the CPU and run times of post-execution processing with the job CPU and run times
  - ❖ `sbatchd` sends both job finish status (`DONE` or `EXIT`) and post-execution processing status (`POST_DONE` or `POST_ERR`) to `mbatchd` at the same time
- ◆ **JOB\_POSTPROC\_TIMEOUT**—Specifies a timeout in minutes for job post-execution processing. The specified timeout must be greater than zero. If post-execution processing takes longer than the timeout, `sbatchd` reports that post-execution has failed (`POST_ERR` status), and kills the entire process group of the job's post-execution processes on UNIX and Linux. On Windows, only the parent process of the post-execution command is killed when the timeout expires. The child processes of the post-execution command are not killed.
- ◆ **JOB\_RUNLIMIT\_RATIO**—Specifies a ratio between a job run limit and the runtime estimate specified by `bsub -We` or `bmod -We`. The ratio does not apply to the `RUNTIME` parameter in `lsb.applications`. This ratio can be set to 0 and no restrictions are applied to the runtime estimate. `JOB_RUNLIMIT_RATIO` prevents abuse of the runtime estimate. The value of this parameter is the ratio of run limit divided by the runtime estimate.
- ◆ **MAX\_JOBID**—By default, LSF assigns job IDs up to 6 digits. This means that no more than 999999 jobs can be in the system at once. Specify any integer from 999999 to 2147483646 (for practical purposes, you can use any 10-digit integer less than this value). You cannot lower the job ID limit, but you can raise it to 10 digits. This means you can have more jobs in the system, and the job ID numbers will roll over less often.
- ◆ **MAX\_SBD\_CONNS**—The default maximum number of file descriptors `mbatchd` can have open and connected concurrently to `sbatchd` is 64.
- ◆ **MBD\_QUERY\_CPUS**—To improve scheduling and dispatch performance of all LSF daemons, you should use `MBD_QUERY_CPUS` together with `EGO_DAEMONS_CPUS` (in `ego.conf`), which controls LIM CPU allocation, and `LSF_DAEMONS_CPUS`, which binds `mbatchd` and `mbschd` daemon processes to specific CPUs so that higher priority daemon processes can run more efficiently. To get best performance, CPU allocation for all four daemons should be assigned their own CPUs. For example, on a 4 CPU SMP host, the following configuration will give the best performance:

```
EGO_DAEMONS_CPUS=0 LSF_DAEMONS_CPUS=1:2 MBD_QUERY_CPUS=3
```

- ◆ NEWJOB\_REFRESH—Enables a child `mbatchd` to get up to date information about new jobs from the parent `mbatchd`. When set to Y, job queries with `bjobs` display new jobs submitted after the child `mbatchd` was created. If you have enabled multithreaded `mbatchd` support, the `bjobs` command may not display up-to-date information if two consecutive query commands are issued before a child `mbatchd` expires because child `mbatchd` job information is not updated. Use `NEWJOB_REFRESH=Y` to enable the parent `mbatchd` to push new job information to a child `mbatchd`.
- ◆ NO\_PREEMPT\_FINISH\_TIME—Specifies that jobs due to finish within the specified number of minutes should not be preempted, where *minutes* is wall-clock time, not normalized time. The job must have a run limit or run time specified for this parameter to affect it.
- ◆ NO\_PREEMPT\_RUN\_TIME—Specifies that jobs that have been running for the specified number of minutes or longer should not be preempted, where *minutes* is wall-clock time, not normalized time. `NO_PREEMPT_RUN_TIME` does not require jobs to be submitted with a runlimit either with `bsub -W` or with `RUNLIMIT` in the queue.
- ◆ PREEMPT\_FOR—If preemptive scheduling is enabled, this parameter is used to disregard suspended jobs when determining if a job slot limit is exceeded, to preempt jobs with the shortest running time, and to optimize preemption of parallel jobs. Specify one or more of the following keywords. Separate multiple keywords with a space:
  - ❖ GROUP\_JLP
  - ❖ GROUP\_MAX
  - ❖ HOST\_JLU
  - ❖ LEAST\_RUN\_TIME
  - ❖ MINI\_JOB
  - ❖ USER\_JLP
- ◆ PREEMPT\_JOBTYPE—If preemptive scheduling is enabled, this parameter enables preemption of exclusive and backfill jobs. Specify one or both of the following keywords. Separate keywords with a space:
  - ❖ EXCLUSIVE
  - ❖ BACKFILL

## lsb.resources

- ◆ NAME—Specify a Limit section name for both vertical and horizontal formats. Specify any ASCII string 40 characters or less. You can use letters, digits, underscores (`_`) or dashes (`-`). You cannot use blank spaces. If duplicate limit names are defined, the Limit section is ignored. If value of NAME is not defined in vertical format, or defined as (`-`), `blimit` displays `NONAME nnn`.
- ◆ JOBS—Specifies a maximum number of running or suspended (RUN, SSUSP, USUSP) jobs available to resource consumers. Specify a positive integer greater than or equal 0. Job limits can be defined in both vertical and horizontal limit formats.

## lsf.cluster

The following RESOURCES are now defined in `lsf.cluster.cluster_name` (ego.cluster.cluster\_name is not created):

- ◆ `define_ncpus_procs`
- ◆ `define_ncpus_cores`
- ◆ `define_ncpus_threads`

## lsf.conf

- ◆ If you manually set `EGO_DAEMON_CONTROL=Y` after installation, you *must* configure LSF `res` and `sbatchd` startup to `AUTOMATIC` in the EGO configuration files `res.xml` and `sbatchd.xml` under `EGO_ESRVDIR/esc/conf/services` to enable automatic LSF daemon startup.

---

**IMPORTANT:** After installation, `LSF_EGO_DAEMON_CONTROL` alone *does not* change the start type for the `sbatchd` and `res` EGO services to `AUTOMATIC` in `res.xml` and `sbatchd.xml` under `EGO_ESRVDIR/esc/conf/services`. You must edit these files and set the `<sc:StartType>` parameter to `AUTOMATIC`.

---

- ◆ When `LSB_API_VERBOSE=Y`, LSF batch commands will display a retry error message to `stderr` when LIM is not available:

LSF daemon (LIM) not responding ... still trying

When `LSB_API_VERBOSE=N`, LSF batch commands will not display a retry error message when LIM is not available.

- ◆ `LSB_DISABLE_LIMLOCK_EXCL`—If preemptive scheduling is enabled, and when `PREEMPT_JOBTYPE=EXCLUSIVE` in `lsb.params`, `LSB_DISABLE_LIMLOCK_EXCL=Y` enables preemption of exclusive jobs and preemption by exclusive jobs. By default, LSF locks the LIM on a host running an exclusive job and unlocks the LIM when the exclusive job finishes. When `LSB_DISABLE_LIMLOCK_EXCL=Y`, for a host running an exclusive job:
  - ◆ LIM is not locked on a host running an exclusive job
  - ◆ `lsload` displays the host status `ok`.
  - ◆ `bhosts` displays the host status `closed`.
  - ◆ Users can run tasks on the host using `lsrun` or `lsgrun`. To prevent users from running tasks during execution of an exclusive job, the parameter `LSF_DISABLE_LSRUN=y` must be defined in `lsf.conf`.
- ◆ `LSB_JOBID_DISP_LENGTH`—By default, LSF commands `bjobs` and `bhist` display job IDs with a maximum length of 7 characters. Job IDs greater than 9999999 are truncated on the left. When `LSB_JOBID_DISP_LENGTH=10`, the width of the `JOBID` column in `bjobs` and `bhist` increases to 10 characters. Specify an integer between 7 and 10.
- ◆ `LSB_LOAD_TO_SERVER_HOSTS` is now obsolete. By default, the client `sbatchd` contacts the local LIM for host status and load information. The client `sbatchd` only contacts the master LIM or a LIM on one of the `LSF_SERVER_HOSTS` if `sbatchd` cannot find the information locally.
- ◆ `LSB_QUERY_PORT` must be defined when `NEWJOB_REFRESH=Y` in `lsb.params` to enable a child `mbatchd` to get up to date information about new jobs from the parent `mbatchd`.

- ◆ LSF\_BIND\_JOB=Y—Enables cluster-wide processor binding for sequential job processes. On Linux execution hosts that support this feature, job processes will be hard bound to selected processors. If processor binding feature is not configured with the BIND\_JOB parameter in an application profile in `lsb.applications`, the `lsf.conf` configuration setting takes effect. The application profile configuration for processor binding overrides the `lsf.conf` configuration. LSF does not bind parallel jobs to any processor. If you require this functionality you should use an MPI that directly supports this.
- ◆ LSF\_DYNAMIC\_HOST\_WAIT\_TIME—the default is now 60 seconds.
- ◆ LSF\_EGO\_DAEMON\_CONTROL—If LSF\_ENABLE\_EGO="N", this parameter is ignored and EGO Service Controller is not started.
- ◆ LSF\_EGO\_ENVDIR—If LSF\_ENABLE\_EGO="N", this parameter is ignored and `ego.conf` is not loaded.
- ◆ LSF\_ENABLE\_EGO—Enables Platform EGO functionality in the LSF cluster. If you set LSF\_ENABLE\_EGO="Y" (the default), you must set or uncomment LSF\_EGO\_ENVDIR in `lsf.conf`. If you set LSF\_ENABLE\_EGO="N" you must remove or comment out LSF\_EGO\_ENVDIR in `lsf.conf`. Set the value to "N" if you do not want to take advantage of the following LSF features that depend on EGO:
  - ❖ LSF daemon control by EGO Service Controller
  - ❖ EGO-enabled SLA scheduling
  - ❖ Platform Management Console (PMC)
  - ❖ LSF reporting
- ◆ LSF\_LD\_SECURITY—When set, jobs submitted using `bsub -Is` or `bsub -Ip` cause the environment variables LD\_PRELOAD and LD\_LIBRARY\_PATH to be removed from the job environment during job initialization to ensure enhanced security against users obtaining root privileges. Two new environment variables are created (LSF\_LD\_LIBRARY\_PATH and LSF\_LD\_PRELOAD) to allow LD\_PRELOAD and LD\_LIBRARY\_PATH to be restored before the job runs.
- ◆ LSF\_LIM\_API\_NTRIES—Defines the number of times LSF commands will retry to communicate with the LIM API when LIM is not available. LSF\_LIM\_API\_NTRIES is ignored by LSF and EGO daemons and EGO commands. The LSF\_LIM\_API\_NTRIES environment variable overrides the value of LSF\_LIM\_API\_NTRIES in `lsf.conf`.
- ◆ LSF\_LOG\_MASK—If EGO is enabled in the LSF cluster, and EGO\_LOG\_MASK is not defined, LSF uses the value of LSF\_LOG\_MASK for LIM, PIM, and MELIM. EGO `vemkd` and `pem` components continue to use the EGO default values. If EGO\_LOG\_MASK is defined, and EGO is enabled, then EGO value is taken.
- ◆ LSF\_MASTER\_LIST—If EGO is enabled, LSF\_MASTER\_LIST or EGO\_MASTER\_LIST can be defined in `lsf.conf`. You can define only EGO\_MASTER\_LIST in `ego.conf`. LIM reads EGO\_MASTER\_LIST wherever it is defined. `lsf.conf` is read first. If EGO\_MASTER\_LIST is

defined in both `lsf.conf` and `ego.conf` the value in `ego.conf` is taken. If EGO is disabled, `ego.conf` not loaded, the value of `EGO_MASTER_LIST` defined in `lsf.conf` is taken.

- ◆ **LSF\_TMPDIR**—By default, LSF creates a temporary directory for a job only on the first execution host. If `LSF_TMPDIR` is set in `lsf.conf`, the path of the job temporary directory on the first execution host is set to `LSF_TMPDIR/job_ID.tmpdir`. If `LSB_SET_TMPDIR= Y`, the environment variable `TMPDIR` will be set equal to the path specified by `LSF_TMPDIR`. Tasks launched through the `blaunch` distributed application framework make use of the LSF temporary directory specified by `LSF_TMPDIR`.

## Environment variables

- ◆ **The `LSB_DEFAULT_JOBGROUP`**—Overrides the setting of `DEFAULT_JOBGROUP` in `lsb.params`.
- ◆ **`LSB_EXIT_IF_CWD_NOTEXIST`**—Indicates that the job will exit if the current working directory specified by `bsub -cwd` or `bmod -cwd` is not accessible on the execution host
- ◆ **`LSB_JOB_INCLUDE_POSTPROC`**—Overrides the value of `JOB_INCLUDE_POSTPROC` in `lsb.params` and `lsb.applications`.
- ◆ **`LSB_JOBGROUP`**—Specifies the name of the job group associated with the job. When a job is dispatched, if it belongs to a job group, the runtime variable `LSB_JOBGROUP` is defined as its group. For example, if a dispatched job belongs to job group `/X`, `LSB_JOBGROUP=/X`. Set during job execution based on `bsub` options or the default job group defined in `LSB_DEFAULT_JOBGROUP` in `lsb.params` and the `DEFAULT_JOBGROUP` environment variable.
- ◆ **`LSF_LD_LIBRARY_PATH`**—When `LSF_LD_SECURITY=Y` in `lsf.conf`, `LSF_LD_LIBRARY_PATH` contains the value of the `LD_LIBRARY_PATH` environment variable, which is removed from the job environment during job initialization to ensure enhanced security against users obtaining root privileges. `LSF_LD_LIBRARY_PATH` allows the `LD_LIBRARY_PATH` environment variable to be restored before the job runs.
- ◆ **`LSF_LD_PRELOAD`**—When `LSF_LD_SECURITY=Y` in `lsf.conf`, `LSF_LD_PRELOAD` contains the value of the `LD_PRELOAD` environment variable, which is removed from the job environment during job initialization to ensure enhanced security against users obtaining root privileges. `LSF_LD_PRELOAD` allows the `LD_PRELOAD` environment variable to be put back before the job runs.
- ◆ **`LSF_LIM_API_NTRIES`**—Defines the number of times LSF commands will retry to communicate with the LIM API when LIM is not available. `LSF_LIM_API_NTRIES` is ignored by LSF and EGO daemons and EGO commands. The `LSF_LIM_API_NTRIES` environment variable. overrides the value of `LSF_LIM_API_NTRIES` in `lsf.conf`.
- ◆ **`LSF_TS_LOGON_TIME`**—Specifies the time to create a Windows Terminal Service session. Configure `LSF_TS_LOGON_TIME` according to the load on your network environment. The default, 30000 milliseconds, is suitable for most environments. If you set `LSF_TS_LOGON_TIME` too small, the LSF tries

multiple times before it succeeds in making a TS session with the TS server, which can cause the job wait a long time before it runs. For a congested network, set `LSF_TS_LOGON_TIME=1000000`.

## New and changed commands, options, and output

The following command options and output are new or changed for LSF Version 7 Update 2:

**bacct** `TERM_CWD_NOTEXIST`—When LSF detects that a job is terminated because the current working directory is not accessible or does not exist on the execution host, `bacct -l` displays the `TERM_CWD_NOTEXIST` keyword. The integer value 25 is logged to `JOB_FINISH` in `lsb.acct`.

**badmin** `showconf mbd | [sbd [ host_name ... | all ]]`—Displays all configured parameters and their values set in `lsf.conf` or `ego.conf` that affect `mbatchd` and `sbatchd`. In a MultiCluster environment, `badmin showconf` only displays the parameters of daemons on the local cluster. Running `badmin showconf` from a master candidate host will reach all server hosts in the cluster. Running `badmin showconf` from a slave-only host may not be able to reach other slave-only hosts. `badmin showconf` only displays the values used by LSF. For example, if you define `LSF_MASTER_LIST` in `lsf.conf`, and `EGO_MASTER_LIST` in `ego.conf`, `badmin showconf` displays the value of `EGO_MASTER_LIST`.

**bgadd** `-L limit` specifies the maximum number of concurrent jobs allowed to run under the job group (including child groups) `-L` limits the number of started jobs (RUN, SSUSP, USSUP) under the job group. Specify a positive number between 0 and 2147483647. If the specified limit is zero (0), no jobs under the job group can run. You cannot specify a limit for the root job group. The root job group has no job limit. Job groups added with no limits specified inherit any limits of existing parent job groups. The `-L` option only limits the lowest level job group created. By default, a job group has no job limit. Limits persist across `mbatchd` restart or reconfiguration.

**bgmod (new)** Modifies the job group with the job group name specified by `job_group_name`. Only root, LSF administrators, the job group creator, or the creator of the parent job groups can use `bgmod` to modify a job group limit. You must provide full group path name for the modified job group. The last component of the path is the name of the job group to be modified.

**bhist**

- ◆ `bhist -l` can display changes to pending jobs as a result of the following `bmod` options:
  - ❖ Absolute priority scheduling (`-aps | -apsn`)
  - ❖ Runtime estimate (`-We | -Wen`)
  - ❖ Post-execution command (`-Ep | -Epn`)
  - ❖ User limits (`-ul | -uln`)
  - ❖ Current working directory (`-cwd | -cwnd`)

## bjgroup

- ◆ Default output now includes the following fields in the displayed list of job groups:
  - ❖ JLIMIT—The job group limit set by `bgadd -L` or `bgmod -L`. Job groups that have no configured limits or no limit usage are indicated by a dash (-). Job group limits are displayed in a USED/LIMIT format. For example, if a limit of 5 jobs is configured and 1 job is started, `bjgroup` displays the job limit under JLIMIT as 1/5.
  - ❖ OWNER—The job group owner.
- ◆ `-N` displays NSLOTS, PEND, RUN, SSUSP, USUSP, RSV counted in slots rather than number of jobs. By default, NSLOTS, PEND, RUN, SSUSP, USUSP, RSV are all counted as the number of jobs rather than number of job slots.
- ◆ `-s` sorts job groups by hierarchy.

## bjobs -l

- ◆ Displays the post-execution command specified at the job-level, by `bsub -Ep`.
- ◆ Displays estimated run time for the job, specified by `bsub -We` or `bmod -We`.

## blimits

- ◆ `blimits` displays current usage of resource allocation limits for JOBS configured in Limit sections in `lsb.resources`. JOBS shows the number of currently running and suspended jobs and the maximum number of jobs configured for the limit policy.
- ◆ `-w` displays resource allocation limits information in a wide format. Fields are displayed without truncation.
- ◆ If a vertical format Limit section has no name, `blimits` displays `NONAMEnnn` under the NAME column for these limits, where the unnamed limits are numbered in the order the vertical-format Limit sections appear in the `lsb.resources` file.

## bmod

- ◆ `-cwdn` sets the current working directory for the job to the directory where `bmod` is running.
- ◆ `-Epn` cancels the setting of job-level post-execution commands. The job-level post-execution commands will not run. Application-level post-execution commands will run if it exists.
- ◆ `-rn` resets the rerunnable job setting specified by `bsub -rn` or `bsub -r`. The application profile and queue level rerunnable job setting if any is used. `bmod -rn` does not disable or override job rerun if the job was submitted to a rerunnable queue or application profile with job rerun configured. `bmod -rn` is different from `bsub -rn`, which *does* override the application profile and queue level rerunnable job setting
- ◆ `-u1n` sets the user shell limits for pending jobs to their default values. `-u1n` is not supported on Windows.
- ◆ `-Wen` cancels the estimated job runtime. The runtime estimate will not take effect for the job.

## bparams

- a displays all the configurable parameters set in `lsb.params`.

**brequeue**

When `JOB_INCLUDE_POSTPROC=Y` is set in `lsb.params` or in an application profile in `lsb.applications`, job requeue will happen only after post-execution processing, not when the job finishes.

**bsub**

- ◆ Job names specified with `bsub -J` can contain up to 4094 characters for UNIX and Linux, or up to 255 characters for Windows.
- ◆ `-rn` specifies that the job is not rerunnable. `bsub -rn` disables job rerun if the job was submitted to a rerunnable queue or application profile with job rerun configured. The command level job rerunnable setting overrides the application profile and queue level setting. This is different from `bmod -rn`, which *cannot* override the application profile and queue level rerunnable job setting.
- ◆ `-ul` passes the current operating system user shell limits for the job submission user to the execution host. User limits cannot override queue hard limits. If user limits exceed queue hard limits, the job will be rejected.

---

**RESTRICTION:** UNIX and Linux only. `-ul` is not supported on Windows.

---

- ◆ `-cwd` specifies the current working directory for the job. By default, if the current working directory is not accessible on the execution host, the job will run in `/tmp`. If the environment variable `LSB_EXIT_IF_CWD_NOTEXIST` is set to `Y` and the current working directory is not accessible on the execution host, the job exits with the exit code 2.
- ◆ `-Ep` runs the specified post-execution command on the execution host after the job finishes. If both application-level (`POST_EXEC` in `lsb.applications`) and job-level post-execution commands are specified, job level post-execution overrides application-level post-execution commands. Queue-level post-execution commands (`POST_EXEC` in `lsb.queues`) run after application-level post-execution and job-level post-execution commands.
- ◆ `-We` specifies an estimated run time for the job. LSF uses the estimated value for job scheduling purposes only, and does not kill jobs that exceed this value unless the jobs also exceed a defined runtime limit. The format of runtime estimate is same as run limit set by the `-W` option. Use `JOB_RUNLIMIT_RATIO` in `lsb.params` to limit the runtime estimate users can set. If `JOB_RUNLIMIT_RATIO` is set to 0 no restriction is applied to the runtime estimate. The job-level runtime estimate setting overrides the `RUNTIME` setting in an application profile in `lsb.applications`.

**lsadmin**

- ◆ `lsadmin limlock`—If `LSB_DISABLE_LIMLOCK_EXCL=y` (to enable preemption of exclusive jobs, for example) LSF suspends all jobs, including exclusive jobs.
- ◆ `showconf lim [ host_name ... | all ]`—Displays all configured parameters and their values set in `lsf.conf` or `ego.conf` that affect `lim`. Use `lsadmin showconf lim` to display the parameters configured in `lsf.conf` and `ego.conf` that apply to root LIM. By default, `lsadmin` displays the local LIM parameters. You can optionally specify the host to display the LIM parameters. In a MultiCluster environment, `lsadmin showconf` only displays the parameters of daemons on the local cluster. Running `lsadmin showconf` from a master candidate host will reach all server hosts in the cluster. Running

`lsadmin showconf` from a slave-only host may not be able to reach other slave-only hosts. `lsadmin showconf` only displays the values used by LSF. For example, if you define `LSF_MASTER_LIST` in `lsf.conf`, and `EGO_MASTER_LIST` in `ego.conf`, `lsadmin showconf` displays the value of `EGO_MASTER_LIST`.

## lsgrun

If `LSB_DISABLE_LIMLOCK_EXCL=y` (to enable preemption of exclusive jobs, for example), you can use `lsgrun` to start a task on hosts that are currently running exclusive jobs.

## lshosts

If EGO is enabled in the LSF cluster and `EGO_DEFINE_NCPUS` is specified in `lsf.conf` or `ego.conf`, the appropriate value for `ncpus` is displayed, depending on the value of `EGO_DEFINE_NCPUS`:

- ◆ `EGO_DEFINE_NCPUS=procs`—`ncpus`=number of processors
- ◆ `EGO_DEFINE_NCPUS=cores`—`ncpus`=number of processors x number of cores
- ◆ `EGO_DEFINE_NCPUS=threads`—`ncpus`=number of processors x number of cores x number of threads

## lsrun

If `LSB_DISABLE_LIMLOCK_EXCL=y` (to enable preemption of exclusive jobs, for example), you can use `lsrun` to start a task on a host that is currently running an exclusive job.

## All X-GUI commands are obsolete

The UNIX X-GUI commands (`xbmod`, `xlsbatch`, `xbsub`, `xlsmon`) are obsolete. The command binaries remain available in the distribution packages as a courtesy, but are provided "as is" and are no longer officially supported.

## New and changed files

No new files have been added in Platform LSF Version 7 Update 2.

## New and changed accounting and job event fields

### lsb.acct

The following fields are new or changed in the `JOB_FINISH` record:

- postExecCmd (%s)** Post-execution command to run on the execution host after the job finishes
- runtimeEstimation (%d)** Estimated run time for the job
- jobGroupName (%s)** Job group name

### lsb.events

The following records are new in `lsb.events`:

- ◆ `GRP_ADD`—Created when a job group is added
- ◆ `GRP_MOD`—Created when a job group is modified

The following fields are new or changed in the `JOB_NEW` and the `JOB_MODIFY2` records:

- postExecCmd (%s)** Post-execution command to run on the execution host after the job finishes
- runtimeEstimation (%d)** Estimated run time for the job

## Bugs fixed since June 2007

The following bugs have been fixed in the November 2007 update (LSF 7 Update 2) since the June 2007 update (LSF 7 Update 1):

97792	Date	2007-11-21
	Description	Master lim gets high volume of UDP requests
	Component	sbatchd
	Platform	All
	Impact	Master lim slows down
78745	Date	2007-11-16
	Description	Job names with wildcards do not match all matching jobs if JOB_DEP_LAST_SUB is set.
	Component	bparams mbatchd
	Platform	All
	Impact	Some scripts depending on the correct behavior may not work as expected.
93768	Date	2007-11-09
	Description	Adapter windows are not cleaned due to signal sent to clean up program by LSF and conflicts of strtok() library calls
	Component	ntbl_api lsntl_api poe_w poejob lsntl_api
	Platform	AIX 5-32 and AIX 5-64
	Impact	Jobs fail or not be dispatched due to lack of adapter windows
96822	Date	2007-10-26
	Description	Data collection script does not have access to job ID job PID information
	Component	sbatchd
	Platform	All
	Impact	Hard to collect data of a specific job
96563	Date	2007-10-25
	Description	NTBL_JOB_KEY in poejob could be out of bound
	Component	poejob
	Platform	AIX 5-32 and AIX 5-64
	Impact	Job can fail if NTBL_JOB_KEY is not generated correctly
95284	Date	2007-10-19
	Description	Files on AIX are owned by root after file transfer
	Component	res
	Platform	AIX
	Impact	High Security risk - file is changed to root ownership.

## What's Changed in Platform LSF Version 7 Update 2

94721	Date	2007-10-17
	Description	Too many xdr messages in mbatchd log files
	Component	mbatchd
	Platform	All
	Impact	The rusage of the job executing in the remote cluster cannot be sent to the local cluster.
95881	Date	2007-10-11
	Description	Array jobs with POE integration fail as jobs overwrite each other's data files
	Component	poejob
	Platform	AIX5 32/64
	Impact	Array jobs fail
94113	Date	2007-10-08
	Description	If LSB_MOD_ALL_JOBS=y is not set in lsf.conf, "bmod -W" on a running job still responds that the modification was successful. bhist also displays successful modification, although bjobs shows that the job was not modified.
	Component	mbatchd
	Platform	All
	Impact	bmod -W does not modify the job.
95000	Date	2007-10-03
	Description	mbschd stops dispatching any job in a queue using slot pool if one array job was killed with some of its job running and some pending. Must restart mbschd to recover.
	Component	mbschd
	Platform	All
	Impact	No jobs can be dispatched even though resources are available.
95059	Date	2007-09-30
	Description	bsub -M cannot accept a large memlimit (greater than 12 digits)
	Component	bmod liblsf.a bsub
	Platform	hpuxia64, linux2.4-x86-64, linux2.4-ia64, linux2.6-x86-64, linux2.6-ia64, linux2.6-ppc64, sparc-sol7-64, sparc-sol10-64, x86-64-sol10
	Impact	Cannot submit a job with bsub -M and a large memlimit.
94525	Date	2007-09-30
	Description	Equal priority queues using cross queue fairshare not handled correctly
	Component	mbschd

94525	Date	2007-09-30
	Platform	All
	Impact	The problem is how scheduler handles the equal priority queues when using cross queue fairshare. The problem only occurs if one of the slave queues has the same priority as the master queue. It does not make a difference if any of the slave queues have the same priority.
94433	Date	2007-09-30
	Description	All jobs in low priority queue are pending with "The queue has reached its job slot limit"
	Component	schmod_default.so
	Platform	All
	Impact	All jobs in low priority queue are pending.
95121	Date	2007-09-27
	Description	If network ID is configured to be a large number with IPV6 enabled, on AIX POE over InfiniBand system, it becomes 0 when loading nrt windows.
	Component	lsnrt_windows
	Platform	AIX 5.3 32/64
	Impact	nrt window cannot be loaded correctly.
95120	Date	2007-09-27
	Description	poejob should export MP_MSG_API
	Component	poejob
	Platform	AIX 5.3 32/64
	Impact	Job runs over IP.
93025	Date	2007-09-26
	Description	Exported shared resource cannot be used by the consumer cluster. All jobs requiring the resource stay pending forever.
	Component	mbatchd
	Platform	All
	Impact	MultiCluster leasing model does not work correctly.
94860	Date	2007-09-25
	Description	LSF allocation cannot satisfy Task Geometry
	Component	res
	Platform	UNIX/Linux
	Impact	User's sequential job might fail if LSB_PJL_TASK_GEOMETRY set in the submission environment. Users must resubmit their jobs either with LSB_PJL_TASK_GEOMETRY not set or requesting more slots for sequential jobs in job script.

## What's Changed in Platform LSF Version 7 Update 2

93579	Date	2007-09-21
	Description	When multiple HP MPI libraries are defined with LSF_VPLUGIN="/usr/lib/libxmpi.so:/usr/lib32/libxmpi.so:/opt/m pi/lib/pa1.1/libmpirm.sl", only the first library is selected.
	Component	pam
	Platform	HP-UX
	Impact	Cannot use multiple HP MPI lib path in LSF_VPLUGIN in lsf.conf.
87229	Date	2007-09-20
	Description	lshosts reports 1 CPU even though lim -t detects number of CPUs properly
	Component	lim
	Platform	Windows Host Type: NTX86, Host Architecture: i1586_3391 dual core
	Impact	Number of CPUs on the host is not reported correctly, leading to licensing and load balancing problems.
92809	Date	2007-09-19
	Description	bhist hangs when for job array job without specifying end of job array index
	Component	bhist libbat.a
	Platform	All
	Impact	The user cannot read history file.
94160	Date	2007-09-14
	Description	After "badadmin reconfig", PRIORITY form "bqueues -l" output become 100, inconsistent with published dynamic priority calculation formula.
	Component	mbatchd
	Platform	All
	Impact	Fairshare information is not accurate.
92460	Date	2007-09-13
	Description	When a long fairshare name is defined, bhpart fails: "Failed in an LSF library call: Failed in sending/receiving a message: Connection reset by peer"
	Component	mbatchd
	Platform	All
	Impact	Cannot use bhpart when a long fairshare name is defined.
93947	Date	2007-09-12
	Description	LSF 7.0.1 does not clean the <i>time_stamp.jobID.hostfile</i> .
	Component	res

93947	Date	2007-09-12
	Platform	All
	Impact	LSF creates a <i>time_stamp.jobID.hostfile</i> along with <i>time_stamp.jobID.out</i> and <i>time_stamp.jobID</i> in <i>.lsbatch</i> directory in users home directory once the job is started. Once the job is completed these temporary files are suppose to be cleaned. However, <i>time_stamp.jobID.hostfile</i> file is left behind and never cleaned up once the job is done. Files keep on accumulating in <i>.lsbatch</i> or in the location where the JOB_SPOOL_DIR is defined.
93614	Date	2007-09-12
	Description	The customer relies on the exit code of bsub -K to trigger some actions. If the job is terminated by bkill -s, however, the bsub -K exit code is always 0, which is not correct.
	Component	nios
	Platform	UNIX/Linux
	Impact	Cannot take advantage of bsub -K feature.
93595	Date	2007-09-10
	Description	When HIST_HOURS=0 and ENABLE_HIST_RUN_TIME=Y are set in a fairshare queue,for finished jobs after mbatchd reconfig or restart, the RUN_TIME of "bqueues -l" output will be cleared.History run time is not accumulated after mbatchd reconfig.
	Component	mbatchd
	Platform	All
	Impact	fairshare becomes not fair when badmin reconfig was done.
92062	Date	2007-09-10
	Description	bjobs usage pattern results in several hundred megabits of traffic between the client and master resulting in a single client taking over 20% of a gigabit network interface resulting also in very high CPU utilization and therefore limiting the capacity of the grid.
	Component	LSF batch library, bjobs, mbatchd
	Platform	All
	Impact	Decreased network performance
91361	Date	2007-08-28
	Description	GRP_CTRL cannot be switched out from lsb.events
	Component	mbatchd
	Platform	All
	Impact	lsb.events file size grows large and affects mbatchd performance.

## What's Changed in Platform LSF Version 7 Update 2

90456	Date	2007-08-28
	Description	With a complex lsb.users file with lots of user groups, "busers all" command fail with "Failed in an LSF library call: Failed in sending/receiving a message: Connection reset by peer"
	Component	busers libbat.so all libbat.a mbatchd
	Platform	All
	Impact	Cannot use complex usre group configuration reliably.
91308	Date	2007-08-27
	Description	brsvs shows expired Advance Reservation with removed host
	Component	mbatchd
	Platform	All
	Impact	Expired Advance Reservation is still listed, and users can submit jobs into the expired Advance Reservation. Jobs pend forever in the expired Advance Reservation.
90811	Date	2007-08-27
	Description	Jobs using Advance Reservation cannot be dispatched until the next occurrence of the reservation after mbatchd restart. If mbatchd is not restarted, those jobs can go expected.
	Component	libbat.so libbat.a mbatchd
	Platform	All
	Impact	Running badmin restart during Advance Reservation active window, causes jobs to be pending until next occurrence.
90760	Date	2007-08-21
	Description	Incorrect calculation of the time period for the next occurrence of an Advance Reservation.
	Component	mbatchd
	Platform	All
	Impact	Recurrent Advance Reservations not working properly.
92134	Date	2007-08-20
	Description	On HP-UX 11-64 slave, sometimes cannot get remote cluster information from LSF commands.
	Component	bqueues bhosts lslod bjobs lslclusters lim blimits lsrin lshosts
	Platform	HP-UX
	Impact	On HP-UX 11-64 slave, sometimes cannot get remote cluster information with LSF commands.
90002	Date	2007-08-15
	Description	If run limits and CPU limits normalized by CPU factor are greater than 0x7FFFFFFF(2147483647)/60 (=35791394), limits are ignored silently.
	Component	mbatchd

90002	Date	2007-08-15
	Platform	All
	Impact	Run limit and CPU limits cannot be set.
91329	Date	2007-08-03
	Description	When any user submits a job, if the job file could NOT be created under \$HOME/.lsbatch and/or specified JOB_SPOOL_DIR directory on NFS by error 'Permission denied', a job file is created under /tmp.
	Component	sbatchd
	Platform	All
	Impact	Confusing behavior. Incorrect "Permission denied" message is logged to sbatchd.log.
91050	Date	2007-08-03
	Description	tssub and bsub -R "msts" crashes after submitting terminal jobs from Windows Vista.
	Component	bsub tssub
	Platform	All
	Impact	tssub crashes
90897	Date	2007-08-02
	Description	Cannot use resource values containing capital letters with mpich2/intelmpi integration
	Component	esub.poe esub.mpich2 esub.intelmpi esub.tvpoe
	Platform	All
	Impact	Jobs not submitted.
61999	Date	2007-08-02
	Description	bhpart cannot work normally if the hosts list is too long
	Component	bhpart mbatchd
	Platform	All
	Impact	Command fails.
90955	Date	2007-07-26
	Description	When "RERUNNABLE = yes" is defined in a queue, cannot use bmod command to modify interactive job in the queue.
	Component	mbatchd
	Platform	All
	Impact	Cannot use bmod command to modify interactive jobs.
90702	Date	2007-07-24
	Description	Task starter cannot register task start event to pam due to host name resolution issue.
	Component	taskstarter pam

## What's Changed in Platform LSF Version 7 Update 2

90702	Date	2007-07-24
	Platform	All
	Impact	If task registration fails, pam will start killing the job.
90414	Date	2007-07-19
	Description	Some options for job submission cannot be set via environment variables.
	Component	libbatch.so libbat.a bsub
	Platform	All
	Impact	Some options for job submission cannot be set via environment variables.
90302	Date	2007-07-18
	Description	If poe authentication is enabled, no job can run, pmd_d log showed b_connect failure
	Component	pmd_wrapper
	Platform	AIX 5-32 aix5-64
	Impact	No jobs can run.
90643	Date	2007-07-16
	Description	Interactive jobs become non-responsive after blaunch is used.
	Component	res
	Platform	All
	Impact	Interactive jobs become non-responsive.
90510	Date	2007-07-16
	Description	badadmin reconfig and lsadmin reconfig error, mbatchd and lim output is incorrect.
	Component	mbatchd, lim
	Platform	Windows Vista x86
	Impact	lsadmin/badadmin reconfig/ckconfig fail.
90281	Date	2007-07-13
	Description	Submission host environment variables appear in execution host
	Component	presvestarter.exe
	Platform	Windows
	Impact	Cannot submit job only based on the execution host environment.
89856	Date	2007-07-13
	Description	Application fails, but pam itself exits with 0
	Component	pam

89856	Date	2007-07-13
	Platform	All platforms that support multi-threading: Linux, AIX, Solaris
	Impact	Wrong pam exit code may affect user job script logic.
83998	Date	2007-07-06
	Description	Unauthorized user can start sbatchd/res services.
	Component	badmin lsadmin
	Platform	All
	Impact	Security risk.
89231	Date	2007-07-03
	Description	blusers -J does not display space between hostname and project field
	Component	blusers
	Platform	All
	Impact	Customers that rely on blusers -J output cannot parse data properly with scripts.
86438	Date	2007-07-02
	Description	Inconsistent fairshare behavior by restarting mbatchd and reconfig mbatchd after user group is removed
	Component	mbatchd
	Platform	All
	Impact	Fairshare does not work as expected after mbatchd restart
82129	Date	2007-06-21
	Description	LSF licenses available but many jobs pending due to no licenses
	Component	mbatchd
	Platform	All
	Impact	License usage is low.
89429	Date	2007-06-19
	Description	res and sbatchd log file logs unclear message about signal delivery
	Component	sbatchd res
	Platform	All
	Impact	Debugging difficult
87624	Date	2007-06-18
	Description	bhist takes long time to return when mbatchd is accessing lsb.events file
	Component	bhist
	Platform	Windows
	Impact	bhist becomes unusable.

## What's Changed in Platform LSF Version 7 Update 2

89164	Date	2007-06-15
	Description	lsmake built on GNU Make 3.8.1 fails to make jobs if there is NFS file attribute caching.
	Component	lsmake
	Platform	All
	Impact	lsmake fails to make jobs.
87548	Date	2007-05-25
	Description	Authentication fails using -cmdfile option of poe command
	Component	poejob
	Platform	AIX 5-32 and AIX 5-64
	Impact	If authentication feature is enabled, jobs submitted using poe -cmdfile option always fail
87384	Date	2007-05-17
	Description	Spaces get stripped from command output at line wrap with -l option
	Component	bacct bjobs bhist, etc.
	Platform	All
	Impact	Unexpected results from scripts that use output of -l option
87279	Date	2007-05-14
	Description	Parameter LSF_LICENSE_FILE is not added into lsfc.conf file after a new installation
	Component	lsfinstall
	Platform	UNIX and Linux
	Impact	Cluster is unlicensed until LSF_LICENSE_FILE parameter is manually added to lsfc.conf
82614	Date	2007-05-11
	Description	Queue-based fairshare does not work as expected when these queues' priorities are equal to the priority of another non-fairshare queue.
	Component	mbschd
	Platform	All
	Impact	Queue-based fairshare does not work as expected.
87100	Date	2007-05-04
	Description	Parallel job using exec rusage pends forever
	Component	schmod_reserve.so schmod_parallel.so schmod_default.so
	Platform	All
	Impact	Job never dispatched

## Known Issues

- ◆ Platform LSF Version 7 Update 2
- ◆ Platform LSF on Windows Vista

### Platform LSF Version 7 Update 2

- |  |  |
|--|--|
| <b>egosh ego<br/>execpasswd and<br/>lspasswd</b>                                 | The password specified on <code>egosh ego execpasswd</code> and <code>lspasswd</code> must be 31 characters or less.   |
| <b>Non-shared slave host<br/>cannot get master list</b>                          | <p>When Platform EGO is enabled, a non-shared slave <code>lim</code> cannot join the cluster if it has different master host list from the one defined on the master host. For example, if the master host list changes while a non-shared slave host is offline, when the slave host restarts, it cannot make contact with the master host because the slave host has a local master host list with no current master candidate hosts in it.</p> <p>Even if the slave host has correct <code>LSF_SERVER_HOSTS</code> set in <code>lsf.conf</code>, it still cannot join the cluster because its local <code>EGO_MASTER_LIST</code> value is wrong. Though the slave host can get the contents of the master host <code>lsf.conf</code> and <code>ego.conf</code>, it ignores the parameters from the master host if it conflicts with a local definition.</p> <p>To correct the problem, you must manually correct the value of <code>EGO_MASTER_LIST</code> in the local copy of <code>ego.conf</code> on each slave host so that it matches the <code>EGO_MASTER_LIST</code> in the shared <code>ego.conf</code>.</p>   |
| <b>LSF 6.1 passwd.lsfuser<br/>password file is not<br/>compatible with LSF 7</b> | <p>In LSF 6.x, if a domain name is defined with <code>LSF_USER_DOMAIN</code> in <code>lsf.conf</code>, LSF only saves the user name to the password entry in the <code>passwd.lsfuser</code> password file.</p> <p>In LSF 7, the user name part of the password entry in the <code>passwd.lsfuser</code> file is a fully qualified user name (<code>domain_name\user_name</code>), even if <code>LSF_USER_DOMAIN</code> is defined in <code>lsf.conf</code>.</p> <p><i>Workaround:</i> If your cluster defines <code>LSF_USER_DOMAIN</code> in <code>lsf.conf</code>, you must upgrade the entire 6.x cluster to LSF 7, and have all users run <code>lspasswd</code> to reenter their password.</p> <p>Without this workaround, LSF 7 daemons cannot find the 6.x password entry and 6.x daemons cannot see the password saved on LSF 7 servers.</p> <p>If you must keep a mixed LSF 7 and LSF 6.x environment:</p> <ul style="list-style-type: none"> <li>◆ You cannot define <code>LSF_USER_DOMAIN</code> in <code>lsf.conf</code></li> <li>◆ Users must run <code>lspasswd</code> on both the 6.x and LSF 7 server hosts</li> </ul> <p>This problem affects all LSF versions before Version 7, LSF 6.0, 6.1, and 6.2.</p> |

## Platform LSF on Windows Vista

**Upgrade the bundled Java Runtime Environment** The version of the Java SE Runtime Environment (JRE) installed with Windows Vista must be upgraded to version 1.5.0\_12-b04 or above. With the older JRE, the LSF reporting database will display an incorrect time stamp in all tables.

Complete the following steps to upgrade the JRE:

- 1 Run the `java -version` command to check the installed version of the JRE. If it is 1.5.0\_12-b04 or above, you do not need to upgrade your JRE.
- 2 Download the Windows JRE 5.0 Update 14 from Sun Microsystems Web site: [http://java.sun.com/javase/downloads/index\\_jdk5.jsp](http://java.sun.com/javase/downloads/index_jdk5.jsp)
  - ❖ For 64-bit Windows Vista on x86 hosts, download `jre-1_5_0_14-windows-amd64.exe`
  - ❖ For 32-bit Windows Vista on x86 hosts, download `jre-1_5_0_14-windows-i586-p.exe`
- 3 Shut down all LSF reporting and PMC daemons.
- 4 Back up the existing `LSF_TOP\jre` directory to `LSF_TOP\jre.old`. `LSF_TOP` is the top-level LSF 7 Update 2 installation directory; for example, `c:\LSF_702`.
- 5 Run `jre-1_5_0_14-windows-amd64.exe` or `jre-1_5_0_14-windows-i586-p.exe` to install the new JRE.

You can install the JRE in the default location and copy it over `LSF_TOP\jre`, or you can specify `LSF_TOP\jre` as the installation directory. Do one of the following:

- ❖ Select custom installation, and set the installation top directory to `LSF_TOP\jre`.
- ❖ Install to the default location, and copy the JRE directory under `LSF_TOP`. For example, if you installed under `c:\jdk1.5`, after installation, you must create the directory `LSF_TOP\jre`, and copy all the files under `C:\jdk1.5\jre` to `LSF_TOP\jre`.

**lsfmon -install fails** Running `lsfmon -install` fails with the following error:

```
OpenSCManager failed - Access is denied. (0x5)
```

Windows Vista User Account Control (UAC) component causes the Administrators group to run most applications as a standard user.

`lsfmon -install` requires administrator rights to run.

To configure `lsfmon` to run as an administrator, complete the following steps:

- 1 Log onto an LSF on Windows host as a member of the local administrators group.
- 2 Right-click on `LSF_TOP\7.0\bin\lsfmon.exe`.
- 3 Click **Properties**, and then select the **Compatibility** tab.
- 4 Under **Privilege Level**, select **Run this program as an administrator**, and then click **OK**.
- 5 Run `lsfmon -install`.

These steps mark `lsfmon` to prompt users for consent, and if granted, run as an administrative application.

- Cannot open LSF help file** The Windows Help (WinHlp32.exe) program is no longer included in Windows operating systems starting with Windows Vista. To view the LSF on Windows help file in Windows Vista or in a future Windows release, you must download WinHlp32.exe from the Microsoft Download Center, and install it on your system. Visit the following Microsoft Web site to download WinHlp32.exe <http://go.microsoft.com/fwlink/?LinkID=82148>. For more information, see <http://support.microsoft.com/kb/917607>.
- Cannot delete uninstall directory** Windows shows "Access denied" when the local Windows administrator or the cluster administrator tries to delete the LSF uninstall directory. The LSF uninstall directory cannot be deleted because the C:\LSF\_7.0\conf\passwd.lsfuser file is owned by "System". The passwd.lsfuser file must be owned by the cluster administrator.

### Shared directory permissions

When users create a shared directory on Windows Vista, the default owner is the directory creator. For LSF to work properly, the shared directory for LSF must be configured so that cluster administrators have read/write permission and all LSF users must have at least read permission. The shared directory must have the following security settings:

- ◆ Owned or co-owned by all cluster administrators
- ◆ Read for all LSF users

**cmd.exe permissions** For installations on an NTFS file system, users must have "Read" and "Execute" privileges for cmd.exe. The following files:

- ◆ %WINDIR%\system32\cmd.exe
- ◆ %WINDIR%\syswow64\cmd.exe

Require the following access permissions:

- ◆ Administrators: Full Control
- ◆ Users: read+execution

## Platform LSF on Windows Server 2008

**Terminal Services jobs do not run** To run Terminal Services (TS) jobs on 64-bit Windows Server 2008, you must add lstsmgr.exe and tscon.exe to the RemoteApp program list. Complete the following steps:

- 1 Open **Start > Control Panel > Administrative Tools > Terminal Services > TS RemoteApp Manager**.
- 2 In the **Action** menu, choose **Add RemoteApp program**.  
You must enable Allow any command-line argument.

## Download the Platform LSF Version 7 Distribution Packages

Download the LSF distribution packages two ways:

- ◆ Through FTP at `ftp.platform.com`
- ◆ Through the World Wide Web at `my.platform.com`

### Download LSF through FTP

**Prerequisites:** Access to the Platform FTP site is controlled by login name and password. If you cannot access the distribution files for download, send email to `support@platform.com`.

- 1 Log on to the LSF file server.
- 2 Change to the directory where you want to download the LSF distribution files. Make sure that you have write access to the directory. For example:  

```
# cd /usr/share/lsf/tarfiles
```
- 3 FTP to the Platform FTP site:  

```
# ftp ftp.platform.com
```
- 4 Provide the login user ID and password provided by Platform.
- 5 Change to the directory for the LSF Version 7 release:  

```
ftp> cd /distrib/7.0
```
- 6 Set file transfer mode to binary:  

```
ftp> binary
```
- 7 For LSF on UNIX and Linux, get the installation distribution file.  

```
ftp> get platform_lsf_update2/lsf7Update2_lsfinstall.tar.Z
```

---

**TIP:** Before installing LSF on your UNIX and Linux hosts, you must uncompress and extract `lsf7Update2_lsfinstall.tar.Z` to the same directory where you download the LSF product distribution tar files.

- 8 Get the distribution packages for the products you want to install on the supported platforms you need. For example:
  - ◆ For the Solaris 7 64-bit version of LSF Version 7:  

```
ftp> get platform_lsf_update2/lsf7Update2_sparc-sol7-64.tar.Z
```

---

**TIP:** Put the LSF distribution files in the same directory as the installation tar files. *Do not* uncompress and extract the distribution files.

- ◆ For 32-bit LSF Version 7 on Windows:  

```
ftp> get platform_lsf_update2/lsf7Update2_win32.msi
```

- 9 Download the Platform LSF Version 7 documentation from  
`/distrib/7.0/docs/`.

```
ftp> get docs/lsf7Update2_documentation.zip  
ftp> get docs/lsf7Update2_documentation.tar.Z
```

---

**TIP:** After installing LSF, you should extract the Platform LSF Version 7 documentation files to `LSF_TOP/docs/lsf`. Browse `LSF_TOP/docs/lsf/index.html` to access the LSF 7 Knowledge Center. If you install the Platform Management Console, the LSF 7 Knowledge Center is installed automatically to `LSF_TOP/docs/lsf`.

---

- 10 Download the Platform EGO Version 1.2.3 documentation from  
`/distrib/7.0/docs/`.

```
ftp> get docs/ego1.2.3_documentation.zip  
ftp> get docs/ego1.2.3_documentation.tar.Z
```

---

**TIP:** After installing LSF, you should extract the EGO documentation files to `LSF_TOP/docs/ego`. Browse `LSF_TOP/docs/ego/index.html` to access the EGO Knowledge Center. If you install the Platform Management Console, the EGO Knowledge Center is installed automatically to `LSF_TOP/docs/ego`.

---

- 11 Optional. Download the Platform Management Console (PMC) distribution package from `/distrib/7.0/platform_lsf_update2/`.

```
ftp> get platform_lsf_update2/lsf7Update2_pmc_linux-x86.tar.Z
```

OR

```
ftp> get  
platform_lsf_update2/lsf7Update2_pmc_linux-x86_64.tar.Z
```

---

**NOTE:** To take advantage of the Platform LSF reporting feature, you *must* download and install the Platform Management Console. The reporting feature is only supported on the same platforms as the Platform Management Console: 32-bit and 64-bit x86 Windows and Linux operating systems. Platform EGO must be enabled to use the PMC and LSF reporting features.

---

- 12 Exit FTP.

```
ftp> quit
```

---

## Download LSF from my.platform.com

**Prerequisites:** You must provide your Customer Support Number and register a user name and password on `my.platform.com` to download LSF.

If you have not registered at `my.platform.com`, click **New User?** and complete the registration form. If you do not know your Customer Support Number or cannot log in to `my.platform.com`, send email to `support@platform.com`.

- 
- 1 Navigate to `http://my.platform.com`.
  - 2 Choose **Products > Platform LSF Family > LSF 7 Update 2**.
  - 3 Under **Download**, choose **Product Packages**.
  - 4 Select the Updates, Packages, and Documentation you wish to download.
  - 5 Log out of `my.platform.com`.
- 

## Archive location of previous update releases

Directories containing release notes and distribution files for previous LSF Version 7 update releases are located on the Platform FTP site under `/distrib/7.0/archive`. Archive directories are named relative to the current update release:

- ◆ LSF Version 7 Update 1: `/distrib/7.0/archive/update1`

## Install Platform LSF Version 7

Installing Platform LSF involves the following steps:

- 1 Get a DEMO license (`license.dat` file).
- 2 Run the installation programs.

### Get a Platform LSF demo license

Before installing Platform LSF Version 7, you must get a demo license key.

Contact `license@platform.com` to get a demo license.

Put the demo license file `license.dat` in the same directory where you downloaded the Platform LSF product distribution tar files.

### Run the UNIX and Linux installation

Use the `lsfinstall` installation program to install a new LSF Version 7 cluster, or upgrade from an earlier LSF version.

See *Installing Platform LSF on UNIX and Linux* for new cluster installation steps.

See the *Platform LSF Command Reference* for detailed information about `lsfinstall` and its options.

---

**IMPORTANT:** Do not use the UNIX and Linux upgrade steps to update an existing LSF Version 7 cluster to LSF Version 7 Update 2. Follow the manual steps in the document *Migrating LSF Version 7 to Update 2 on UNIX and Linux* to migrate your existing LSF Version 7 cluster to LSF Version 7 Update 2.

---

For LSF 7 Update 2 only, you *cannot* use `lsfinstall` to *migrate* your existing LSF Version 7 cluster to LSF 7 Update 2. Follow the manual steps in the document *Migrating LSF Version 7 to Update 2 on UNIX and Linux* (`/distrib/7.0/lsf_migrate_unix_linux_to_update2.pdf`) to migrate your existing LSF Version 7 cluster to LSF Version 7 Update 2. This is required due to the change in LSF directory structure for LSF 7 Update 2. In later updates, `lsfinstall` will automatically handle the update.

### Run the Windows installation

Platform LSF on Windows 2000, Windows 2003, and Windows XP is distributed in the following packages:

- ◆ `lsf7Update2_win32.msi`
- ◆ `lsf7Update2_win-x64.msi`
- ◆ `lsf7Update2_win-ia64.msi`

See *Installing Platform LSF on Windows* for new cluster installation steps.

To migrate your existing LSF Version 7 cluster on Windows to LSF 7 Update 2, you must follow the manual steps in the document *Migrating LSF Version 7 to Update 2 on Windows* (`/distrib/7.0/lsf_migrate_windows_to_update2.pdf`). This is required due to the change in LSF directory structure for LSF 7 Update 2 only.

Install Platform LSF Version 7

## Install Platform LSF HPC

Use `lsfinstall` to install a new Platform LSF HPC cluster or to upgrade LSF HPC from a previous release.

---

**IMPORTANT:** Make sure `ENABLE_HPC_INST=Y` is specified in `install.config` to enable Platform LSF HPC installation.

---

See *Using Platform LSF HPC* for installation and configuration steps.

## Install Platform LSF Desktop Support

See the *Platform LSF Desktop Support Administrator's Guide* for installation and configuration steps.

## Learn About Platform LSF Version 7

Information about Platform LSF is available from the following sources:

- ◆ [World Wide Web and FTP](#)
- ◆ [Platform LSF documentation](#)
- ◆ [Platform EGO documentation](#)
- ◆ [Platform training](#)

### World Wide Web and FTP

Information about Platform LSF Version 7 is available in the LSF area of the Platform FTP site (<ftp.platform.com/distrib/7.0/>).

The latest information about all supported releases of Platform LSF is available on the Platform Web site at [www.platform.com](http://www.platform.com).

If you have problems accessing the Platform web site or the Platform FTP site, send email to [support@platform.com](mailto:support@platform.com).

### my.platform.com

[my.platform.com](http://my.platform.com)—Your one-stop-shop for information, forums, e-support, documentation and release information. [my.platform.com](http://my.platform.com) provides a single source of information and access to new products and releases from Platform Computing.

On the Platform LSF Family product page of [my.platform.com](http://my.platform.com), you can download software, patches, updates and documentation. See what's new in Platform LSF Version 7, check the system requirements for Platform LSF, or browse and search the latest documentation updates through the Platform LSF Knowledge Center.

### Platform LSF documentation

The Platform LSF Knowledge Center is your entry point for all LSF documentation. You should extract the LSF documentation distribution file to the directory `LSF_TOP/docs/lsf`. After extracting the LSF documentation distribution file, browse `LSF_TOP/docs/lsf/7.0/index.html` to access and search the Platform LSF Knowledge Center.

If you have installed the Platform Management Console, access and search the Platform LSF documentation through the link to the Platform Knowledge Center.

### Platform EGO documentation

The Platform EGO Knowledge Center is your entry point for Platform EGO documentation. It is installed when you install LSF. To access and search the EGO documentation, browse the file `LSF_TOP/docs/ego/1.2.3/index.html`.

If you have installed the Platform Management Console, access the EGO documentation through the link to the Platform Knowledge Center.

## Get Technical Support

### Platform training

Platform's Professional Services training courses can help you gain the skills necessary to effectively install, configure and manage your Platform products. Courses are available for both new and experienced users and administrators at our corporate headquarters and Platform locations worldwide.

Customized on-site course delivery is also available.

Find out more about Platform Training at [www.platform.com/Services/Training/](http://www.platform.com/Services/Training/), or contact [Training@platform.com](mailto:Training@platform.com) for details.

## Get Technical Support

### Contact Platform

Contact Platform Computing or your LSF vendor for technical support. Use one of the following to contact Platform technical support:

**Email** [support@platform.com](mailto:support@platform.com)

**World Wide Web** [www.platform.com](http://www.platform.com)

**Mail**  
Platform Support  
Platform Computing Inc.  
3760 14th Avenue  
Markham, Ontario  
Canada L3R 3T7

When contacting Platform, please include the full name of your company.

See the Platform Web site at [www.platform.com/company/contact-us](http://www.platform.com/company/contact-us) for other contact information.

### Get patch updates and other notifications

To get periodic patch update information, critical bug notification, and general support notification from Platform Support, contact [supportnotice-request@platform.com](mailto:supportnotice-request@platform.com) with the subject line containing the word "subscribe".

To get security related issue notification from Platform Support, contact [securenotice-request@platform.com](mailto:securenotice-request@platform.com) with the subject line containing the word "subscribe".

## We'd like to hear from you

If you find an error in any Platform documentation, or you have a suggestion for improving it, please let us know:

### Email

[doc@platform.com](mailto:doc@platform.com)

### Mail

Information Development  
Platform Computing Inc.  
3760 14th Avenue  
Markham, Ontario  
Canada L3R 3T7

Be sure to tell us:

- ◆ The title of the manual you are commenting on
- ◆ The version of the product you are using
- ◆ The format of the manual (HTML or PDF)

## Copyright

© 1994-2008, Platform Computing Inc.

Although the information in this document has been carefully reviewed, Platform Computing Inc. (“Platform”) does not warrant it to be free of errors or omissions. Platform reserves the right to make corrections, updates, revisions or changes to the information in this document.

UNLESS OTHERWISE EXPRESSLY STATED BY PLATFORM, THE PROGRAM DESCRIBED IN THIS DOCUMENT IS PROVIDED “AS IS” AND WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESSED OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE. IN NO EVENT WILL PLATFORM COMPUTING BE LIABLE TO ANYONE FOR SPECIAL, COLLATERAL, INCIDENTAL, OR CONSEQUENTIAL DAMAGES, INCLUDING WITHOUT LIMITATION ANY LOST PROFITS, DATA, OR SAVINGS, ARISING OUT OF THE USE OF OR INABILITY TO USE THIS PROGRAM.

### Document redistribution policy

This document is protected by copyright and you may not redistribute or translate it into another language, in part or in whole.

### Internal redistribution

You may only redistribute this document internally within your organization (for example, on an intranet) provided that you continue to check the Platform Web site for updates and update your version of the documentation. You may not make it available to your organization over the Internet.

### Trademarks

LSF is a registered trademark of Platform Computing Corporation in the United States and in other jurisdictions.

POWERING HIGH PERFORMANCE, PLATFORM COMPUTING, PLATFORM SYMPHONY, PLATFORM JOBSCHEDULER, and the PLATFORM and PLATFORM LSF logos are trademarks of Platform Computing Corporation in the United States and in other jurisdictions.

UNIX is a registered trademark of The Open Group in the United States and in other jurisdictions.

Linux is the registered trademark of Linus Torvalds in the U.S. and other countries.

Microsoft is either a registered trademark or a trademark of Microsoft Corporation in the United States and/or other countries.

Windows is a registered trademark of Microsoft Corporation in the United States and other countries.

Macrovision, Globetrotter, and FLEXIm are registered trademarks or trademarks of Macrovision Corporation in the United States of America and/or other countries.

Oracle is a registered trademark of Oracle Corporation and/or its affiliates.

Intel, Itanium, and Pentium are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.

Other products or services mentioned in this document are identified by the trademarks or service marks of their respective owners.

## Third Party License Agreements

[www.platform.com/legal-notices/third-party-license-agreements](http://www.platform.com/legal-notices/third-party-license-agreements)

Copyright

