

HP XC How To



Installing LSF-HPC for SLURM into an existing
standard LSF cluster

Version 1.1 Nov, 2005

© 2005 Hewlett-Packard Development Company, L.P.

© 2005 Hewlett-Packard Development Company, L.P. The information contained herein is subject to change without notice. The only warranties for HP products and services are set forth in the express warranty statements accompanying such products and services. Nothing herein should be construed as constituting an additional warranty. HP shall not be liable for technical or editorial errors or omissions contained herein.

Linux is a U.S. registered trademark of Linus Torvalds.

LSF, Platform Computing, and the LSF and Platform Computing logos are trademarks or registered trademarks of Platform.

Contents

Introduction

Requirements	5
--------------------	---

HP XC Preparation

Ensure that LSF on HP XC is shut down	7
Mount the LSF <i>tree</i> to XC	7
Ensure that the XC resource_management nodes have an external connection	8
Lower the firewall on the XC external network	8
Node-to-node communication	9
Setup the expected LSF environment	9
Set up the XC LSF startup script	10
Completing the Preparation	10

Installation of XC LSF

Perform Post-Installation Tasks	12
Configuring the XC LSF alias	14
Starting LSF on XC	14
Sample running jobs	15
Troubleshooting	16
Sample installation dialog	17

Revision history

Revision tables

Table 1 Revisions

Date	Edition	Revision
Jun 2005	V1.0	First Edition
Nov 2005	V1.1	Update and reformat for inclusion in the HP XC Manuals

Introduction

This *HP XC How To* describes how you can install LSF-HPC for SLURM into an existing, standard LSF cluster.

An understanding of standard LSF installation and configuration procedures is required to perform this procedure. You should be familiar with the LSF installation documentation and the README file provided in the LSF installation tar file. You should also be familiar with the normal procedures in adding a node to an existing LSF cluster, such as establishing default communications (`.rhosts` or ssh keys), setting up shared directories, and adding common users.

Before you begin, read the documentation for LSF-HPC with SLURM for XC, located in the *HP XC System Software Administration Guide*.

An existing LSF cluster typically has a single common `LSF_TOP` location, also known as the **LSF root** or **LSF tree**. In this location, multiple versions of LSF are installed and centrally maintained.

NOTE: LSF does not create an actual `LSF_TOP` environment variable; the term is used only to identify the LSF root in the documentation.

The procedure described in this document assumes a single `LSF_TOP` location, shared by both the existing LSF cluster and the XC cluster. The `LSF_TOP` location resides on the existing LSF cluster and it is exported to the XC cluster. The XC cluster mounts the NFS `LSF_TOP` location from the existing LSF cluster. Any changes that you make to the configuration files in `LSF_TOP/conf` are visible to both clusters.

You will install LSF-HPC for XC into the `LSF_TOP` location on the existing LSF cluster, maintaining the existing LSF cluster, while providing the SLURM support needed by the HP XC cluster.

Requirements

The procedure used in this *HP XC How To* has the following requirements:

- You can add LSF-HPC for XC only to an existing, standard LSF cluster running LSF V6.0 or later. Recent versions of standard LSF contain the required `schmod_slurm` interface module.
- HP XC v2.1 contains the necessary support for this procedure. To install on a prior version, obtain the latest `hptc-lsf` RPM package.
- LSF daemons communicate through pre-configured ports in `lsf.conf`. However, LSF commands open random ports for receiving information when they communicate with the LSF daemons. Because an LSF cluster needs this open network environment, maintaining a firewall is beyond the scope of this *XC How To*. You can attempt the procedure, but it is not guaranteed to work with a firewall in place.

The example used in this *HP XC How To* uses a sample cluster with the following characteristics:

- The cluster consists of an existing LSF cluster installed on a single node with a hostname of `plain`.
- An HP XC LSF node with the hostname `xc1sf` is added to the cluster.
- The name `xc1sf` is the LSF alias associated with the HP XC cluster. Using an independent IP and hostname as the LSF alias enables LSF-HPC failover capability. In the event of a failure, another node automatically hosts the alias. (See the `control1sf` man page for more details on setting up LSF failover.)
- The head node of the example HP XC cluster is `xc-head`.
- The node `plain` also serves the LSF tree to the HP XC cluster as an NFS file system.

A summary of the IP addresses and host names used in this example follows:

plain	16.32.1.24	Non-XC node in pre-existing LSF cluster
xc-head	16.32.2.128	Head node of XC cluster
xc	16.32.2.130	XC cluster alias (not used by LSF)
xclsf	16.32.2.140	IP and hostname to be used as the external XC LSF alias

The XC cluster alias is mentioned here to prevent confusion between the XC cluster alias and the XC LSF alias. It is possible to have two separate aliases on HP XC, and depending on the cluster configuration, these aliases might even reside on the same node. However, their purposes are distinct from each other.

There might also be additional IPs and hostnames used by an HP XC cluster. For example, each node with an external connection (either for login or I/O purposes) typically has its own IP and hostname.

For the purpose of the example, the existing cluster characteristics are as follows:

- The value of `LSF_TOP` is set to `/shared/lsf`.
- The account `lsfadmin` is configured as the Primary LSF Administrator.
- The name of the existing LSF cluster is `corplsf`.

HP XC preparation

You must perform the following steps to prepare to install LSF-HPC for SLURM into an existing standard LSF cluster. Read through all of these steps first to ensure that you understand what is to be done. All steps are performed via a login to the HP XC head node, and most of the steps involve propagating changes to the rest of the cluster.

Ensure that LSF on HP XC is shut down

Use the following procedure to shut down and remove LSF:

1. If LSF was installed locally on HP XC and is running, shut it down as follows:

```
[root@xc128 conf]# controllsf stop
```

2. Remove this LSF installation by using the following configuration command. In this step, the LSF installation is removed from the current `LSF_TOP`, which is `/opt/hptc/lsf/top`:

```
[root@xc128 conf]# /opt/hptc/etc/gconfig.d/C55lsf gunconfigure
removing /opt/hptc/lsf/top/conf...
removing /opt/hptc/lsf/top/6.0...
removing /opt/hptc/lsf/top/work...
removing /opt/hptc/lsf/top/log...
removing /hptc_cluster/lsf/conf...
removing /hptc_cluster/lsf/work...
removing /var/lsf...
```

3. Log out and log back in to clear your environment of any old LSF environment settings.

In the following procedure, a new `LSF_TOP` is mounted from the non-XC system `plain`.

Mount the LSF *tree* to XC

This procedure describes how to NFS mount the `LSF_TOP` location, (`/shared/lsf` in this example), from the non-XC system.

1. On the non-XC system, (`plain` in this example), export the directory specified by '`LSF_TOP`' to the XC cluster. A typical existing LSF cluster already has this location exported to the other nodes; you need only add the XC cluster. Refer to the following documentation, if necessary:
 - a. For UNIX and Linux systems, see the `exports(5)` reference (man) page for instructions on exporting directories.
 - b. See the *HP XC System Software Administration Guide* and the `/hptc_cluster/etc/fstab.proto` file for information on mounting external filesystems cluster-wide.
2. Create the mount point cluster-wide.

3. Add the appropriate `fstab` file entry to the `/hptc_cluster/etc/fstab.proto` file in the section titled ALL.
4. Restart the `cluster_fstab` service cluster-wide. For this example:

```
# pdsh -a mkdir -p /shared/lsf
# pdsh -a service cluster_fstab restart
```

Ensure that the XC resource_management nodes have an external connection

The node that is configured as the LSF node (which will be one of the resource_management nodes) must be able to communicate with, and receive communication from, the existing LSF cluster on the external network.

If there are nodes with the resource_management role that do not have external network access, you must reconsider the cluster configuration to ensure that LSF on XC runs only on nodes that have external network access. Options include:

- Adding additional network hardware to current resource_management nodes
- Reassigning the resource_management role.

See the *XC Installation Guide* for more information on configuring and reconfiguring roles in XC.

Use the `shownodes` command to ensure that each node that was configured as a resource_management node during `cluster_config` also has access to the external network:

```
[root@xc128 root]# shownode roles --role resource_management external
resource_management: xc[127-128]
external: xc[125-128]
```

If this command is not available, check the role assignments by running the `cluster_config` command and viewing the node configurations.

NOTE: Use the `quit` command to close `cluster_config` when you have determined the configuration. Do not reconfigure the cluster with any changes at this point. There is another opportunity to reconfigure the cluster later in this procedure.

In the next four steps, you make modifications to the head node configuration and propagate the changes to the rest of the cluster. The recommended method is to use the `updateimage` and the `updateclient` commands as documented in the *HP XC System Software Administration Guide*.

Lower the firewall on the XC external network

For this procedure, open up the unprivileged ports (1024-65535) plus one privileged port (1023) on the external network by adding the following lines to `/etc/sysconfig/iptables.proto` on the head node:

```
-A RH-Firewall-1-INPUT -i External -p tcp -m tcp --dport 1023:65535 -j ACCEPT
-A RH-Firewall-1-INPUT -i External -p udp -m udp --dport 1023:65535 -j ACCEPT
```

This file establishes the initial firewall rules for all nodes. The revised rules open up all the unprivileged ports externally and one privileged port (1023). Opening up the privileged port allows LSF commands run as root on XC to communicate with non-XC LSF daemons. LSF commands executed by the root account always use privileged ports. If necessary, you can avoid opening the privileged port.

You must apply the new rules to every node that might be selected to run the LSF-HPC daemons. A later step in this procedure describes how to generate a new `/etc/sysconfig/iptables` file on each node, using your modified `iptables.proto` file.

Node-to-node communication

LSF can use either `rsh` or `ssh` to control all the LSF daemons in the cluster. The daemons expect the selected mechanism to enable access to all nodes without a password.

Because `rsh` transmits passwords unencrypted through the network, HP does not recommend its use in HP XC clusters. (Unencrypted passwords might be revealed by a standard network-sniffing program.) Consequently, `rsh` and its related packages are not installed on HP XC clusters by default and `ssh` is recommended instead.

If you prefer to use `rsh` within the LSF cluster, install its RPM packages on the head node before proceeding to set up LSF. The `rsh` packages are provided on the HP XC software DVD.

Setup the expected LSF environment

A typical LSF installation provides two environment setup files, named `profile.lsf` and `cshrc.lsf` by default. When sourced by the user, the setup files enable access to the LSF binaries, man pages, and libraries.

When LSF is installed locally, two custom system files named `/etc/profile.d/lsf.sh` and `/etc/profile.d/lsf.csh` are created. The system files automatically source the LSF environment setup files at log in, giving all authorized users access to LSF as soon as they log in to the cluster.

The current contents of the two system files are shown below. The current `LSF_TOP`, (`/opt/hptc/lsf/top`), is replaced by a new `LSF_TOP`, (`/shared/lsf`), which is a location that is shared between the two clusters.

```
[root@xc128 profile.d]# cat lsf.sh

case $PATH in
    *-slurm/etc:*) ;;
    */opt/hptc/lsf/top*) ;;
    *)
        if [ -f /opt/hptc/lsf/top/conf/profile.lsf ]; then
            . /opt/hptc/lsf/top/conf/profile.lsf
        fi
    esac
```

```
[root@xc128 profile.d]# cat lsf.csh

if ( "${path}" !~ *-slurm/etc* ) then
    if ( -f /opt/hptc/lsf/top/conf/cshrc.lsf ) then
        source /opt/hptc/lsf/top/conf/cshrc.lsf
    endif
endif
```

The objective of these custom files is to source (once only) the appropriate LSF environment file, which is:

- `$LSF_ENVDIR/cshrc.lsf`, for `csh` users
- `$LSF_ENVDIR/profile.lsf`, for users of `sh`, `bash`, and other `sh`-based shells.

Create `/etc/profile.d/lsf.sh` and `/etc/profile.d/lsf.csh` on the cluster to setup the LSF environment on. Using `/shared/lsf` for the value of `LSF_TOP` as an example, the new files look similar to the following:

```
[root@xc128 profile.d]# cat lsf.sh
```

```

case $PATH in
  *-slurm/etc:*) ;;
  */shared/lsf/*) ;;
  *)
  if [ -f /shared/lsf/conf/profile.lsf.xc ]; then
    . /shared/lsf/conf/profile.lsf.xc
  fi
esac

```

```
[root@xc128 profile.d]# cat lsf.csh
```

```

if ( "${path}" !~ *-slurm/etc* ) then
  if ( -f /shared/lsf/conf/cshrc.lsf.xc ) then
    source /shared/lsf/conf/cshrc.lsf.xc
  endif
endif

```

To summarize the current state of the installation:

- The scripts `profile.lsf.xc` and `cshrc.lsf.xc` scripts do not exist yet. You will establish them after installing LSF-HPC for XC later in this procedure.
- You have seen that the names of the setup files are appended with the `xc` string. After installing LSF-HPC for XC, you must rename the setup files to maintain the existing setup files.
- In a later step, you will propagate the new custom files `/etc/profile.d/lsf.sh` and `/etc/profile.d/lsf.csh` to every node in the XC cluster.

Set up the XC LSF startup script

The XC `controllsf` command can also double as the Red Hat `/etc/init.d/` service script for starting and stopping XC LSF during boot and shutdown of the cluster. When starting LSF, this command establishes the LSF alias and starts the LSF daemons. The command can start LSF on one of the `resource_management` nodes only. It is safe, and also necessary for failover, to configure `controllsf` on every node in the cluster.

To configure the `controllsf` command as the service script, create the appropriate soft-link:

```

[root@xc128 root]# which controllsf
/opt/hptc/sbin/controllsf
[root@xc128 root]# ln -s /opt/hptc/sbin/controllsf /etc/init.d/lsf

```

This configuration must be present on every node in the XC cluster. At this point, you have completed the changes to the head node.

Completing the preparation

To complete the preparation, update the golden image with the changes that you made in the preceding sections, and propagate the changes to the cluster nodes. See *HP XC System Software Administration Guide* for the recommended procedures to update the golden image and propagate the changes. Complete the preparation as follows:

1. Use the `cluster_config` command to update the golden image if you must change XC roles to ensure that all `resource_management` nodes have direct access to the external network.

Enter `n` (no) when prompted to install LSF locally.

2. If additional software (such as `rsh`) was not installed, use the following procedure to update the golden image and propagate these changes with minimal impact to a running cluster. (This action propagates changes only to those nodes that are up and running):
 - a. Update the golden image with the following command:

```
# updateimage --gc `nodename nh` --image base_image --no-netboot
```

- b. Use the `pdcpl` command to propagate the specific file changes to all nodes:

```
# pdcpl -a /etc/sysconfig/iptables.proto /etc/sysconfig/iptables.proto
# pdcpl -a /etc/profile.d/lsf.sh /etc/profile.d/lsf.sh
# pdcpl -a /etc/profile.d/lsf.csh /etc/profile.d/lsf.csh
# pdcpl -a /etc/init.d/lsf /etc/init.d/lsf
```

- c. Refresh the firewall on all nodes with the following command. This will update the `/etc/sysconfig/iptables` file on all nodes with the changes made in a previous step:

```
# pdsh -a /opt/hptc/etc/nconfig.d/C04iptables nconfigure
```

Installation of XC LSF

Now that you have prepared the cluster, install XC LSF into the LSF tree as described in the following procedure.

1. Preserve the existing environment setup files.

Change directory to the existing `LSF_TOP/conf` directory and rename the setup files by appending a unique identifier. For example:

```
# cd /shared/lsf/conf
# mv profile.lsf profile.lsf.orig
# mv cshrc.lsf cshrc.lsf.orig
```

The installation of XC LSF provides its own `profile.lsf` and `cshrc.lsf` files, and you must rename those files to a unique name, then restore the originals. The environment setup files are needed only on XC, and are automatically sourced by the `/etc/profile.d/lsf.sh` and `/etc/profile.d/lsf.csh` files when users log on to XC.

2. Copy the XC LSF tar files to a temporary location on the node that hosts

`LSF_TOP`.

In the example the temporary location is the node named `plain`. When the copy is complete, unpack the installation scripts as follows:

```
[root@plain lsf]# mkdir hpctmp
[root@plain lsf]# scp root@xc-head:/opt/hptc/lsf/files/hpc* hpctmp/
root@xc-head's password:
hpc6.0_hpcinstall.ta 100% |*****| 237 KB 00:00
hpc6.0_linux2.4-glib 100% |*****| 37039 KB 00:08
[root@plain lsf]# cd hpctmp/
[root@plain hpctmp]# tar zxf hpc6.0_hpcinstall.tar.Z
[root@plain hpctmp]# ls
hpc6.0_hpcinstall hpc6.0_hpcinstall.tar.Z hpc6.0_linux2.4-glibc2.3-amd64-slurm.tar.Z
```

3. Prepare the LSF installation process.

This step consists of manual edits to the `install.config` file. Set the appropriate values as documented in this file. For the example, you set the following:

```
LSF_TOP="/shared/lsf"           - existing cluster's LSF_TOP
LSF_ADMINIS="lsfadmin"         - existing cluster's admin
LSF_CLUSTER_NAME="corplsf"     - name of existing cluster
LSF_LICENSE="<none>"           - XC-specific, for HP OEM licensing
LSF_ADD_SERVERS="xclsf"        - the XC LSF "node" to be added
```

4. Start the LSF installation process.

Use the following command to start the installation process:

```
[root@plain hpc6.0_hpcinstall]# ./hpcinstall -f install.config
```

Perform post-installation tasks

The LSF documentation and instructions mentioned at the end of the `hpc_install` script are generic and do not apply to an HP XC cluster. The following manual procedures describe every task that you must perform on an HP XC cluster:

1. Restore the original environment setup files.

Change directory back to the existing `LSF_TOP/conf` directory and rename the environment setup files to distinguish the XC files and restore the original files. Using the example:

```
# cd /shared/lsf/conf
# mv profile.lsf profile.lsf.xc
# mv cshrc.lsf cshrc.lsf.xc
# mv profile.lsf.orig profile.lsf
# mv cshrc.lsf.orig cshrc.lsf
```

The XC environment setup files now match the files configured in `/etc/profile.d/lsf.sh` and `/etc/profile.d/lsf.csh` from the earlier step.

2. Obtain the XC internal network base.

Examine the `/opt/hptc/config/base_addr.ini` file on the XC cluster and identify the value of `nodeBase`. By default, this value is `172.20`, representing the internal XC network. Using this value, edit the `LSF_TOP/conf/lsf.cluster.<cluster name>` file as follows:

- a. In the **Host** section, find the XC node, and add `slurm` in the `RESOURCES` column. For the example:

```
Begin Host
HOSTNAME model type server rlm mem swp RESOURCES #Keywords
...
xclsf ! ! 1 3.5 () () (slurm)
End Host
```

- b. In the **Parameters** section, set up the floating client address range (`FLOAT_CLIENTS_ADDR_RANGE`) using the `nodeBase` entry from Step 1. For the example:

```
Begin Parameters
PRODUCTS=LSF_Base ... Platform_HPC
FLOAT_CLIENTS_ADDR_RANGE=172.20
End Parameters
```

Save and exit the file

3. Edit the `LSF_TOP/conf/lsf.conf` file.

Make the following changes to the `LSF_TOP/conf/lsf.conf` file:

- a. Create or modify the `LSF_SERVER_HOSTS` variable to add the XC LSF node, along with the other LSF Execution Hosts in the cluster:
`LSF_SERVER_HOSTS="plain xclsf"`
- b. Enable HP OEM licensing by adding the following variable:
`XC_LIBLIC=/opt/hptc/lib/libsyslic.so`
- c. Disable `LSF_NON_PRIVILEGED_PORTS`, by setting its value to `N`, or remove it from the file. In standard LSF v6.1, this parameter is not supported, and generates bad port messages from the `sbatchd` and `mbatchd` daemons on a non-XC node.
- d. If you intend to use `ssh` for node-to-node communication, set the following variable in `lsf.conf` (assuming you configured the `ssh` keys to allow access without a password):

```
LSF_RSH=ssh
```

Save and exit the file.

4. Optional: Configure any special XC-specific queues.

For HP XC V2.1, HP recommends that you use a `JOB_STARTER` script, configured for all queues on a XC system. The default installation of LSF on XC provides queue configurations in the `/opt/hptc/lsf/etc/configdir/lsb.queues` file. The `JOB_STARTER` script and its helper scripts are located in the `/opt/hptc/lsf/bin/` file. Depending on how they are configured for use, you might need to move the scripts to a shared location. If you move the scripts, edit the `job_starter.sh` script, updating its definition of `loc`, which provides the location of the helper scripts.

The `JOB_STARTER` script minimizes the impact when a is launched on a non-XC node and you can configure it on all queues, but use good judgement in employing this script on your cluster.

5. Optional: Adjust the LSF environment files if necessary.

When you install a new or different version of LSF, it renames the old `profile.lsf` and `cshrc.lsf` files before replacing them with the files from the most recent installation. Ensure that these files are consistent with the XC default LSF environment files, as configured previously.

Configuring the XC LSF alias

Having completed the XC LSF installation, log out and log back into XC to establish the proper LSF environment.

The next step is to configure the LSF alias on XC. An alias is used on XC to prevent hardwiring LSF to any one node, so that the LSF node in XC can fail over to another node if the current node becomes compromised (hung or crashed). XC provides infrastructure to monitor the LSF node and fail over the LSF daemons to another node if necessary.

The selected IP and host name must not be in use but must be known on the external network. The example is using 'xclsf' with an IP address of '16.32.2.140', and the head node for the XC cluster is 'xc-head'.

To verify that the selected external host name is not currently in use, ping it as follows:

```
[root@xc128 conf]# ping xclsf
PING xclsf.lab.mycorp.com (16.32.2.140) 56(84) bytes of data.
From xc128.lab.mycorp.com (16.32.2.128) icmp_seq=0 Destination Host
Unreachable
```

Configure `controllsf` (which manages LSF setup on XC) with the new alias, as follows:

```
[root@xc128 root]# controllsf set virtual hostname xclsf
```

Confirm that the alias is set, as follows:

```
[root@xc128 root]# controllsf show
LSF is currently shut down, and assigned to node .
Failover is disabled.
Head node is preferred.
The primary LSF host node is xc128.
SLURM affinity is enabled.
The virtual hostname is "xclsf".
```

Starting LSF on XC

To update LSF with the latest configuration changes use the `lsadmin reconfig` and `badadmin reconfig` commands within the existing LSF cluster (on node `plain` in the example). Subsequent

use of the `lshosts` or `bhosts` commands will display the new XC node, although the node status is `UNKNOWN` and `unavailable`, respectively.

You can now start LSF on XC as follows:

```
[root@xc128 root]# controllsf start
```

This command sets up the virtual LSF alias on the appropriate node and then starts the LSF daemons. It also creates a `$LSF_ENVDIR/hosts` file (in the example `$LSF_ENVDIR = /shared/lsf/conf`). This hosts file is used by LSF to map the LSF alias to the actual host name of the node in XC running LSF. See the Platform LSF documentation for information on `hosts` files.

When the LSF daemons have started up and synchronized their data with the rest of the LSF cluster, use the `lshosts` and `bhosts` commands to display all the nodes with their appropriate values and to verify that they are ready for use as follows:

```
[root@xc128 root]# lshosts
HOST_NAME      type      model    cpuf  ncpus  maxmem  maxswp  server  RESOURCES
plain          LINUX86   PC1133   23.1    2    248M   1026M   Yes  ()
xclsf         SLINUX6   Intel_EM 60.0    256   3456M    -      Yes  (slurm)
[root@xc128 root]# bhosts
HOST_NAME      STATUS    JL/U     MAX   NJOBS    RUN   SSUSP  USUSP   RSV
plain          ok        -        2     0        0     0      0      0
xclsf         ok        -       256    0        0     0      0      0
```

Sample running jobs

The following scenarios provides examples of how you might run jobs on the newly installed system.

- As a user on the non-XC node, launching to a Linux ia32 resource:

```
[test@plain test]$ bsub -I -n1 -R type=LINUX86 hostname
Job <411> is submitted to default queue <normal>.
<<Waiting for dispatch ...>>
<<Starting on plain>>
plain
```

- Launching to an XC resource:

```
[test@plain test]$ bsub -I -n1 -R type=SLINUX64 hostname
Job <412> is submitted to default queue <normal>.
<<Waiting for dispatch ...>>
<<Starting on xclsf>>
xc127
```

- Launching a parallel job:

```
[test@aplain test]$ bsub -I -n6 -R type=SLINUX64 srun hostname
Job <413> is submitted to default queue <normal>.
<<Waiting for dispatch ...>>
<<Starting on xclsf>>
/home/test/.lsbatch/1113947197.413: line 8: srun: command not found
[test@plain test]$ bsub -I -n6 -R type=SLINUX64 /opt/hptc/bin/srun hostname
Job <414> is submitted to default queue <normal>.
<<Waiting for dispatch ...>>
<<Starting on xclsf>>
xc3
xc3
xc2
xc2
```

```
xc1
xc1
```

In this scenario, the `srun` command was not found because the user's `$PATH` did not include `/opt/hptc/bin`, which is specific to XC. There are several standard ways to address this if necessary. For example, you can add `/opt/hptc/bin` to the default `$PATH` on the non-XC node; or create a softlink to the `srun` command from `/usr/bin` on all the nodes in XC.

- User on the XC node, launching to a Linux ia32 resource:

```
[test@xc128 test]$ bsub -I -n1 -R type=LINUX86 hostname
Job <415> is submitted to default queue <normal>.
<<Waiting for dispatch ...>>
<<Starting on plain>>
plain
```

- Launching to an XC resource:

```
[test@xc128 test]$ bsub -I -n6 -R type=SLINUX64 srun hostname
Job <416> is submitted to default queue <normal>.
<<Waiting for dispatch ...>>
<<Starting on xc1sf>>
xc3
xc3
xc2
xc2
xc1
xc1
```

Troubleshooting

If your cluster does not perform or behave as expected after you have applied this *XC HowTo*, use the following procedure to verify the configuration:

- Use the following commands to check your configuration changes:
 - Confirm the firewall settings (use other command options if necessary):
`iptables -L`
 - Confirm startup script:
`pdsh -a 'ls -l /etc/init.d/lsf'`
 - Confirm that the LSF tree was properly mounted (using the example):
`pdsh -a 'ls -ld /shared/lsf/'`
 - Confirm the LSF environment scripts:
`pdsh -a 'ls -l /etc/profile.d/lsf.sh'`
- Monitor the LSF log files when problems arise for communication complaints, unresolved host name issues, or configuration problems.
- Ensure that the XC firewall is disabled. The firewall prevents the non-XC LSF nodes from communicating with the XC LSF node.
- Ensure that `controllsf` has been properly configured with alias. Run `controllsf show` to confirm its settings.
- Check the `ifconfig` output on the XC LSF node to ensure that the LSF alias was properly established. If `eth0` is the external network device, the LSF alias entry is `eth0:lsf`.
- Use the appropriate LSF commands to restart daemons when network communications have been adjusted:
 - `lsadmin reconfig` - to restart the Load Information manager (LIM)
 - `badmin mbdrestart` - to restart the Master Batch Daemon (mbatchd)
 - `badmin reconfig` - to reset the batch settings for the Slave Batch Daemon (sbatchd).

See the Platform LSF documentation for more information on these commands.

Sample installation dialog

A sample installation script is provided below.

```
Logging installation sequence in /shared/lsf/hpctmp/hpc6.0_hpcinstall/Install.log
LSF pre-installation check ...
```

```
Checking the LSF TOP directory /shared/lsf ...
... Done checking the LSF TOP directory /shared/lsf ...
LSF license is defined in "/shared/lsf/conf/lsf.conf", LSF_LICENSE is ignored ...
```

```
Checking LSF Administrators ...
LSF administrator(s):      "lsfadmin smith"
Primary LSF administrator: "lsfadmin"
```

```
Searching hpc distribution tar files in /shared/lsf/hpctmp Please wait ...
```

```
1) linux2.4-glibc2.3-amd64-slurm
```

```
Press 1 or Enter to install this host type: 1
```

```
You have chosen the following tar file(s):
hpc6.0_linux2.4-glibc2.3-amd64-slurm
```

```
Checking selected tar file(s) ...
... Done checking selected tar file(s).
```

```
Pre-installation check report saved as text file:
/shared/lsf/hpctmp/hpc6.0_hpcinstall/prechk.rpt.
```

```
... Done LSF pre-installation check.
```

```
Installing hpc binary files "hpc6.0_linux2.4-glibc2.3-amd64-slurm"...
```

```
Copying hpcinstall files to /shared/lsf/6.0/install
... Done copying hpcinstall files to /shared/lsf/6.0/install
```

```
Installing linux2.4-glibc2.3-amd64-slurm ...
```

```
Please wait, extracting hpc6.0_linux2.4-glibc2.3-amd64-slurm may take up to a few minutes
...
```

```
... Done extracting /shared/lsf/hpctmp/hpc6.0_linux2.4-glibc2.3-amd64-slurm.tar.Z.
```

```
Creating links to LSF commands ...
... Done creating links to LSF commands ...
```

```
Modifying owner, access mode, setuid flag of LSF binary files ...
... Done modifying owner, access mode, setuid flag of LSF binary files ...
```

```
Creating the script file lsf_daemons ...
... Done creating the script file lsf_daemons ...
```

```
... linux2.4-glibc2.3-amd64-slurm installed successfully under /shared/lsf/6.0.
```

```
... Done installing hpc binary files "linux2.4-glibc2.3-amd64-slurm".
```

LSF is already installed ...

Old version of LSF configuration files exist in /shared/lsf.
LSF configuration files under /shared/lsf/conf will be upgraded.

corplsf is an existing cluster ...

Updating PRODUCTS line in /shared/lsf/conf/lsf.cluster.corplsf ...

1. Backup /shared/lsf/conf/lsf.cluster.corplsf to
/shared/lsf/conf/lsf.cluster.corplsf.old.31585

2. Enable Platform_HPC

3. Remove LSF_Data and LSF_Parallel

Setting common HPC external resources to /shared/lsf/conf/lsf.shared

Enabling LSB_SHORT_HOSTLIST in /shared/lsf/conf/lsf.conf ...

Enabling schmod_slurm in /shared/lsf/conf/lsbatch/corplsf/configdir/lsb.modules ...

Setting JOB_ACCEPT_INTERVAL = 0 to /shared/lsf/conf/lsbatch/corplsf/configdir/lsb.params ...

Setting MXJ to ! in /shared/lsf/conf/lsbatch/corplsf/configdir/lsb.hosts ...

Adding server hosts ...

Host(s) "xclsf" has (have) been added to the cluster "corplsf".

... LSF configuration is done.

Creating hpc_getting_started.html ...

... Done creating hpc_getting_started.html

Creating hpc_quick_admin.html ...

... Done creating hpc_quick_admin.html

hpcinstall is done.

To complete your hpc installation and get your
cluster "corplsf" up and running, follow the steps in
"/shared/lsf/hpctmp/hpc6.0_hpcinstall/hpc_getting_started.html".

After setting up your LSF server hosts and verifying
your cluster "corplsf" is running correctly,
see "/shared/lsf/6.0/hpc_quick_admin.html"
to learn more about your new LSF cluster