IBM Spectrum LSF

LSF Getting Started

John Welch

June 24th, 2019
Updated March 15th, 2021 for 10.1 FP11

Agenda

Sourcing LSF

Directory Overview

Is my cluster up?

Start, stop and restart

Add a node to the cluster

Configuration changes

Top commands

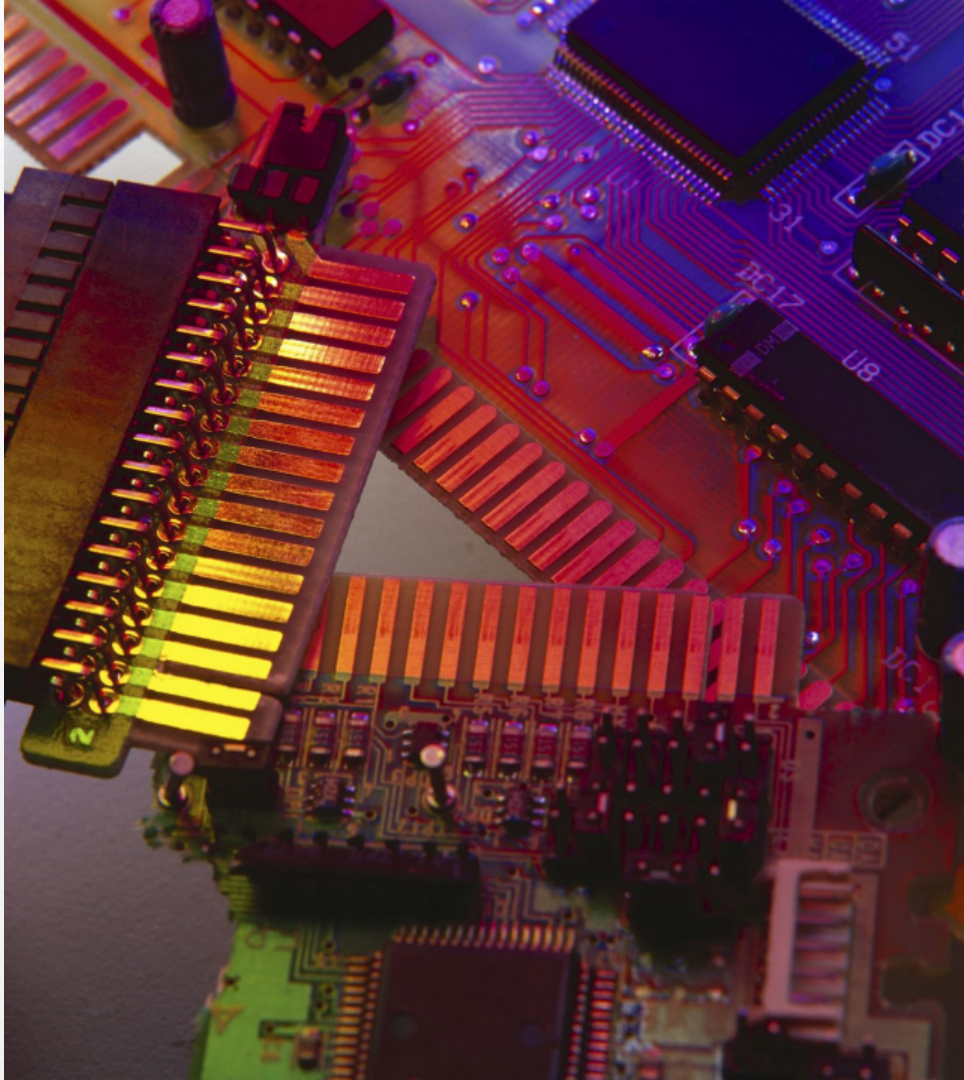

Sourcing LSF

On Bourne or bash shells

. LSF_TOP/conf/profile.lsf

On Cshell

source LSF_TOP/conf/cshrc.lsf

Change LSF_TOP to your site specific directory such as /tools/lsf.

Environmental variables set

```
$ env | grep LSF
LSF_SERVERDIR=/tools/lsf/10.1/linux3.10-glibc2.17-x86_64/etc
LSF_LIBDIR=/tools/lsf/10.1/linux3.10-glibc2.17-x86_64/lib
LSF_BINDIR=/tools/lsf/10.1/linux3.10-glibc2.17-x86_64/bin
LSF_ENVDIR=/tools/lsf/conf
$
```

Notes

- The user's path will be updated to include $LSF_BINDIR
- LSF supports multiple platforms
  - linux2.6-glibc2.3-x86_64 is the platform directory for RHEL 6.X on x86_64.
  - linux3.10-glibc2.17-x86_64 is the platform directory for RHEL 7.X on x86_64.

LSF Directory Overview

Directory	Description	Location
LSF_SERVERDIR	Daemons (lim, res, sbatchd, mbatchd, ...) and other plugins such as elim*, esub*, ...	LSF_TOP/10.1/<platform>/etc
LSF_LIBDIR	Library files	LSF_TOP/10.1/<platform>/lib
LSF_BINDIR	Binary/commands	LSF_TOP/10.1/<platform>/bin
LSF_ENVDIR	Configuration	LSF_TOP/conf
BATCH	Batch configuration (lsb.queues, lsb.hosts, etc)	LSF_ENVDIR/lsbatch/<cluster_name>/configdir
LOG	Log files (check /tmp if not here)	LSF_TOP/log
PATCH	Patch information and backups for roll back of patches	LSF_TOP/patch
INSTALL	Installer, patch installer, hostsetup (not in PATH), etc	LSF_TOP/10.1/install

LSF Base (ls*) and Batch (b*)

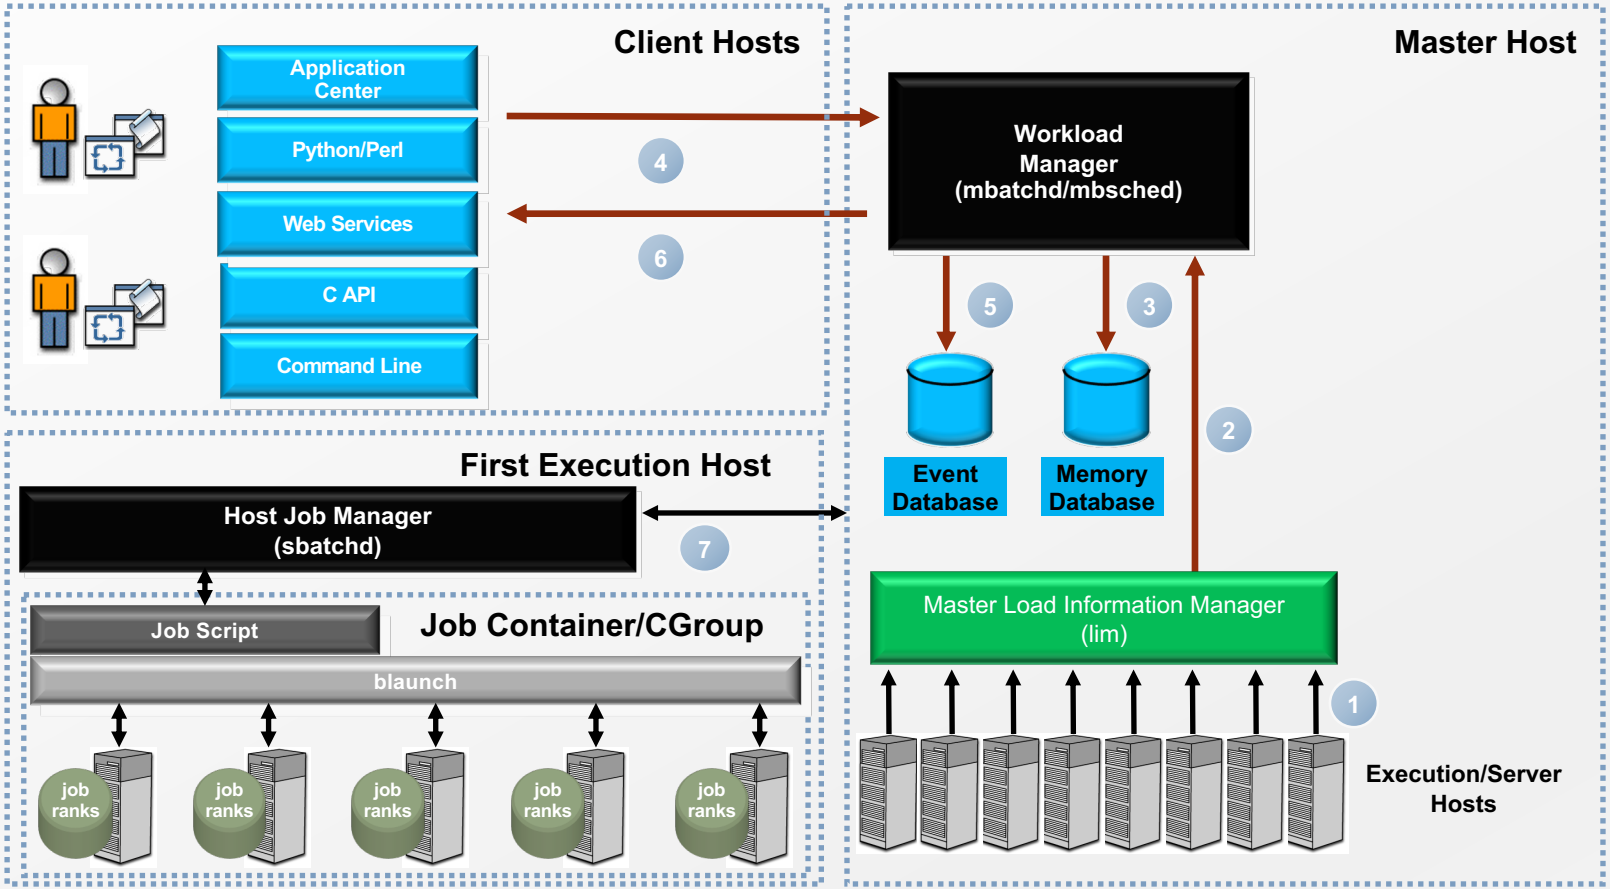

Is your cluster up?

Step 1 – Check the LSF base?

$ lsid
IBM Spectrum LSF Standard 10.1.0.6, Sep 09 2018
Copyright International Business Machines Corp. 1992, 2016.
US Government Users Restricted Rights - Use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM Corp.

My cluster name is YourClusterName
My master name is lsfmaster1

Step 2 – Check the LSF batch?

```
$ bhosts
HOST_NAME          STATUS       JL/U    MAX  NJOBS    RUN  SSUSP  USUSP    RSV
lsfmaster1         closed          -      0      0      0      0      0      0
compute1           ok              -     16      0      0      0      0      0
$
```

lim (base) and sbatchd (batch) runs on every node in the cluster except LSF Clients

$ ps –ef | grep lim
$ ps –ef | grep sbatchd

LSF Master Batch daemon only runs on the current LSF Master
$ ps –ef | grep mbatchd

How to start LSF

Start LSF on current node

Method 1
bctrld start lim
bctrld start res
bctrld start sbd

Method 2

- **On RHEL 6.X**

  /etc/init.d/lsf start

- **On RHEL 7.X**

  systemctl start lsfd

Start LSF on all nodes in the cluster (except LSF Clients). Parallel option: -pdsh [-delay seconds] [-num_hosts number]
lsfstartup

Start on a different node or all nodes
bctrld start lim [hostname | all]
bctrld start res [hostname | all]
bctrld start sbd [hostname | all]

Notes
Run above as root unless lsf.sudoers is configured on cluster nodes
bctrld use LSF_RSH method (defined in $LSF_ENVDIR/lsf.conf) to access remote systems on startup
Typically, passwordless ssh is setup from LSF Master to LSF Servers.

How to stop LSF

Stop LSF on current node

Method 1 (order matters)
bctrld stop sbd
bctrld stop res
bctrld stop lim

Method 2

- **On RHEL 6.X**

  /etc/init.d/lsf stop

- **On RHEL 7.X**

  systemctl stop lsfd

Stop LSF on all nodes in the cluster. Parallel option: -pdsh
lsfshutdown

Extra step on the LSF Master node to shutdown the Master batch daemon (mbatchd)
bctrld stop sbd
badmin mbdrestart –s
bctrld stop res
bctrld stop lim

Other options
bctrld stop sbd [hostname | all]
bctrld stop res [hostname | all]
bctrld stop lim [hostname | all]

How to restart the whole cluster

On the LSF Master node

bctrld restart lim all
bctrld restart res all
bctrld restart sbd all

badmin mbdrestart

Restart LSF on all nodes in the cluster (except LSF Clients). Parallel option: -pdsh

lsfrestart

Adding a Node Part 1

Step 1 – add the new node to the cluster file

a) Login into the LSF Master as root
b) Change directory to $LSF_ENVDIR
c) Edit the lsf.cluster.<cluster_name> file
d) Add the blue line for the new host. For example,

```
NewNode    !    !    1    ()
End      Host
```

Step 2 – restart lim on the LSF Master node and any LSF Master Candidate nodes and the mbatchd

a) Login into the LSF Master as root

bctrld restart lim *lsfmaster1 lsfmaster2*
badmin mbdrestart

Step 3 – setup passwordless ssh from LSF Master to new LSF Server

Adding a Node Part 2

Step 3 – Run the host setup

a) Login into the new node as root
b) Make sure the LSF_TOP directory is accessible on the new node

```
$ cd LSF_TOP/10.1/install
$ ./hostsetup --top="LSF_TOP" --boot="y" --start="y" --profile="y"
```

Replace LSF_TOP with your site specific LSF installation location

Step 4 – Verify host is in cluster and okay with lshosts and bhosts commands

Documentation link on adding a node to your cluster

https://www.ibm.com/support/knowledgecenter/SSWRJV_10.1.0/lsf_admin_foundations/add_hosts.html

Note, documentation specifies a using reconfig, but this can be replaced with restarts as noted in step 2

Recommend setting up lsf.sudoers on each node in the LSF cluster to allow lsfadmin and other specified users to start and stop LSF services

https://www.ibm.com/support/knowledgecenter/en/SSWRJV_10.1.0/lsf_config_ref/lsf.sudoers.5.html

Recommended Configuration changes

Additional lsf.conf parameters

Add these parameters to the bottom of $LSF_ENVDIR/lsf.conf

LSF_RSH="ssh -o 'PasswordAuthentication no' -o 'StrictHostKeyChecking no'"
LSF_STRIP_DOMAIN=.your.domain.com
LSB_STDOUT_DIRECT=Y
LSB_HJOB_PER_SESSION=4
LSB_SIGSTOP=SIGTSTP

To prevent job dispatch on the LSF Master node, set the job slot limit to zero in lsb.hosts

Add the blue line to $LSF_ENVDIR/lsbatch/<your_cluster_name>/lsb.hosts and replace lsfmaster1 with your LSF Master node name

```
default    !     ()       ()     ()      ()      ()             (Y)    # Example
lsfmaster1 0     ()       ()     ()      ()      ()             (Y)    # Example
End Host
```

LSF Documentation link on configuration changes and the required restart or reconfiguration

https://www.ibm.com/support/knowledgecenter/en/SSWRJV_10.1.0/lsf_admin/cluster_reconfig.html

Top LSF Commands

Base level host information

```
$ lshosts
HOST_NAME      type    model  cpuf ncpus maxmem maxswp server RESOURCES
lsfmaster1   X86_64 Intel_EM  60.0    16  63.6G  27.8G    Yes (mg)
compute1     X86_64 Intel_EM  60.0    16  63.6G  27.8G    Yes (docker)
client1      UNKNOWN UNKNOWN   1.0     -      -      -     No ()
$
Use lshosts –w to see wide format when fields like HOST_NAME are truncated
```

Batch level host information

```
$ bhosts
HOST_NAME          STATUS       JL/U    MAX  NJOBS    RUN  SSUSP  USUSP    RSV
lsfmaster1         closed          -      0      0      0      0      0      0
compute1              ok           -     16      0      0      0      0      0
$
```

Current load reported by each hosts

```
$ lsload
HOST_NAME       status  r15s   r1m  r15m   ut    pg  ls    it   tmp   swp   mem
lsfmaster1          ok   0.3   0.3   0.2   1%   0.0   1     0  212G 27.8G 52.7G
compute1            ok   0.3   0.3   0.2   1%   0.0   1     0  212G 27.8G 52.7G
$
```

Top LSF Commands

Job Submission command: bsub

Batch job example
bsub sleep 10

Interactive job example
bsub –I ls

Job status command: bjobs

Shows jobs still in mbatchd memory with various filtering and displaying options

See your active jobs (-r for running jobs, -p for pending jobs, -a all jobs)
```
bjobs
```

Long listing of a job
```
bjobs –l <jobid>
```

Long listing of pending jobs
```
bjobs –l -p
```

See all users jobs
```
bjobs –u all
```

Top LSF Commands

Job history: bhist

Shows historical job information from lsb.events* files

```
bhist -l <jobid>
```

Job kill: bkill

```
bkill <jobid>
```

Job suspend/resume: bstop/bresume

Queue information: bqueues

Thanks You

John Welch
jswelch@us.ibm.com