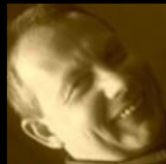


IBM Spectrum LSF

What's New in LSF Service Pack 12



Bill.McMillan@uk.ibm.com
Principal Product Manager
Spectrum Computing



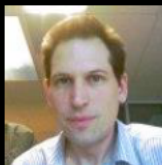
Yun Gao
Technical Support Professional
Spectrum Computing



Mingliang Zu
Technical Account Manager
Spectrum Computing



Yi Sun
Technical Account Manager
Spectrum Computing



Michael Spriggs
STSM Spectrum Computing
Spectrum Computing

Upcoming Seminars

Details and Registration on www.ibm.biz/LSFCommunitySeminars

September 2, 2021

What's New in LSF Service Pack 12

Speaker: LSF Support Team

September 16, 2021

Best Practices for Upgrading your LSF Clusters

Speaker: Larry Adams, Expert Labs

September 30, 2021

What if....? Using the LSF Simulator to answer those hard questions

Speakers: John Welch & Renita Leung, LSF SME

October. Date to be confirmed.

The LSF.Next Beta Program

Speaker: Bill McMillan, Product Management

October 14, 2021

Simplifying HPC – Just push the button! Tips & Tricks

Speaker: Gabor Samu, Product Management

October 25-28, 2021

IBM TechU: <https://www.ibm.com/training/events>

- Applied AI and HPC: Case Studies from IBM
- Deploying LSF with OCP: Tips and Tricks
- HPC Cloud Bursting on IBM Cloud: Tips and Tricks

November 11, 2021

Expediting PMR's with the LSF Support Tools

Speaker: LSF Support Team

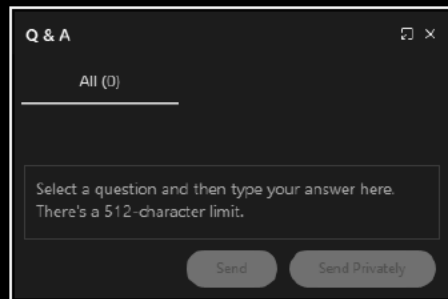
November 29, 2021

High Performance Computing - Health Check Services from Lab Services

Speakers(s): Larry Adams, Expert Labs

Agenda

- Security Enforcement in FP12
- Nvidia A100 MIG Support in FP12
- Resource Connector Enhancement for Google Cloud
- Quick Deployment of LSF on IBM Cloud
- Deprecated Features
- Q&A: You can ask questions at any time in the Q&A panel



Security Enforcement in FP12



Security Enforcement in FP12

- A Brief Review of Major Security Features Introduced Prior to FP12
- New Parameters
- New Command Option
- New Enforced Default Value to The Security Related Parameters
- Compatibility

A Brief Review of Major Security Features Introduced Prior to FP12

- eauth.cve added (site-specific default encryption key and time difference check) in FP2
 - User customized eauth encryption key (LSF_EAUTH_KEY)
- eauth.cve upgrading in FP7
 - New parameter LSF_EXT_SERVERDIR in lsf.conf
 - New parameter LSF_ENV_OVERRIDE in lsf.conf
 - New hostsetup script options (--ext-serverdir --eauth-key)
- eauth.cve replaces the default eauth in FP10

A Brief Review of Major Security Features Introduced Prior to FP12

- Root privilege disabled by default in FP10
 - LSF_ROOT_REX removed in lsf.conf
 - New parameter LSF_ROOT_USER in lsf.conf
- Authentication of query commands (LSF_AUTH_QUERY_COMMANDS) in FP11
- Secured authorization requests (LSF_STRICT_CHECKING=enhanced) in FP11

New Parameters

- LSF_EAUTH_OLDKEY (lsf.sudoers)

Specifies the previous key that eauth used to encrypt and decrypt user authentication data after you specify a new eauth key

- LSF_EAUTH_OLDKEY_EXPIRY (lsf.sudoers)

Specifies the expiry date for the previous eauth key (LSF_EAUTH_OLDKEY_EXPIRY parameter), after which the previous key no longer works and only the new LSF_EAUTH_KEY parameter works.

New Command Option

- badmin security view

No	Category	Comment
1	Authentication	Check eauth version, LSF_EAUTH_KEY, LSF_EAUTH_OLDKEY_EXPIRY
2	Message Integrity	Check LSF_STRIC_CHEKCING
3	Secure Startup	Check LSF_STARTUP_USERS, set-user-id of bctrlld, lsadmin, badmin
4	Secure Configuration	Check /etc/lsf.conf
5	Disallow Root User	Check LSF_ROOT_USER
6	Kerberos Impersonation	Check LSF_KRB_IMPERSONATE
7	Interactive Preload Protection	Check LSF_LD_SECURITY
8	Custom Plugin Location	Check LSF_EXT_SERVERDIR
9	Query Authentication	Check LSF_AUTH_QUERY_COMMANDS

New Command Option

- badmin security view

```
yxu@ib20b01-401: badmin security view
Authentication                      Medium
Enhanced Message Integrity          Enabled
Secure Startup                      Enabled for USERLIST
Secure Configuration                Disabled
Disallow Root User                 Enabled
Interactive Preload Protection      Disabled
Custom Plugin Location              Disabled
Query Authentication                Enabled
```

```
yxu@ib20b01-684: badmin security view -v
Authentication                      Medium
You are using an up-to-date version of eauth with an automatically generated
encryption key.
```

```
Enhanced Message Integrity          Enabled
The default value of LSF_STRICT_CHECKING=ENHANCED is set, ensuring that
communications cannot be tampered with.
```

```
Secure Startup                      Disabled
It is strongly recommended to remove the setuid bit from lsadmin, badmin,
bctrld and eauth. Daemon startup/shutdown should be performed with "bctrld"
and setuid should only be set on bctrld and eauth when LSF_STARTUP_USERS is
defined in lsf.sudoers.
```

```
Secure Configuration                Enabled
Secure configuration is enabled and cluster configuration will always be read
from /etc/lsf.conf and $LSF_ENVDIR will be ignored.
```

```
Disallow Root User                 Enabled
The root user is not allowed to perform actions as a valid user from the
LSF command line.
```

```
Interactive Preload Protection      Disabled
Preload libraries are allowed when running interactive workload. This may
weaken security and is not recommended.
```

```
Custom Plugin Location              Disabled
Your e*plugins (eauth/elim etc) are stored in $LSF_SERVERDIR. This means they
could be accidentally overridden during upgrades.
```

```
Query Authentication                Enabled
Authentication is enabled for LSF query commands.
```

New Enforced Default Value to the Security Related Parameters

Parameters	Prior FP12	FP12
LSF_STRICT_CHECKING	Not defined	ENHANCED
LSF_AUTH_QUERY_COMMANDS	N	Y
LSF_ENV_OVERRIDE	Y	N

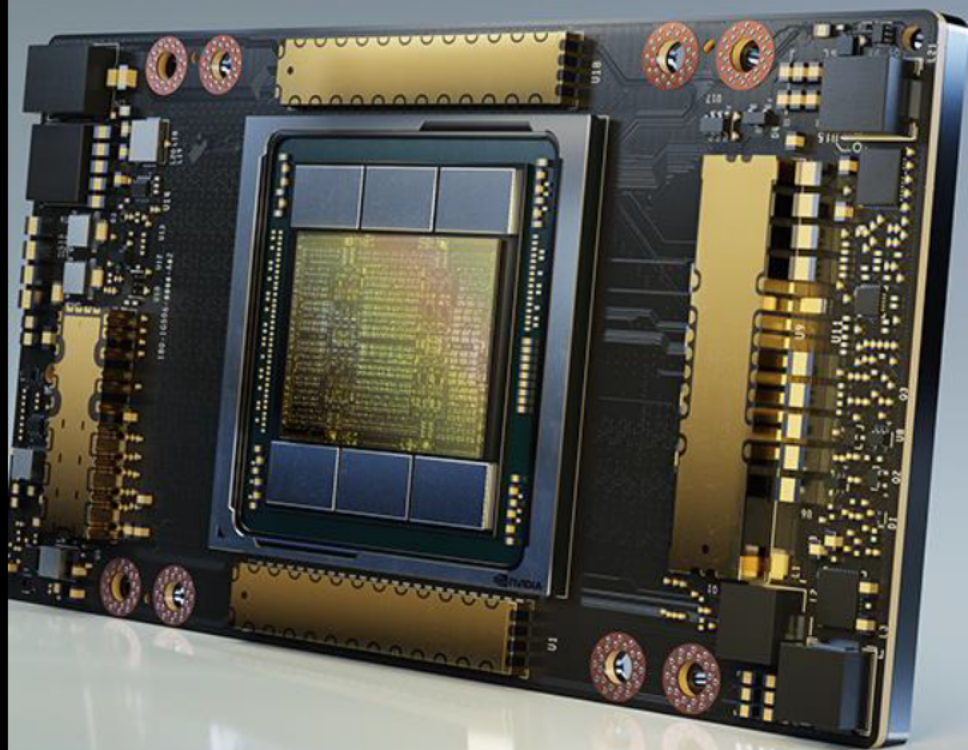
Compatibility

<div>LS, DM, PPM, RTM, MC PAC, Explore</div> <div>LSF</div>	[Prior FP2)	[FP2 ~ FP11]	FP12
[10.1 GA ~ FP2)	✓	✗	✗
[FP2 ~ FP11]	✗	✓	✗
FP12	✗	Condition	✓

- **Condition**

1. If LSF_STRICT_CHECKING is not defined in lsf.conf, "N" should be configured to resolve the compatibility issue.
2. If LSF_AUTH_QUERY_COMMANDS is not defined in lsf.conf, "N" should be configured to resolve the compatibility issue.
3. If above 2 parameters are defined in lsf.conf, everything should be working fine.

Nvidia A100 MIG Support in FP12

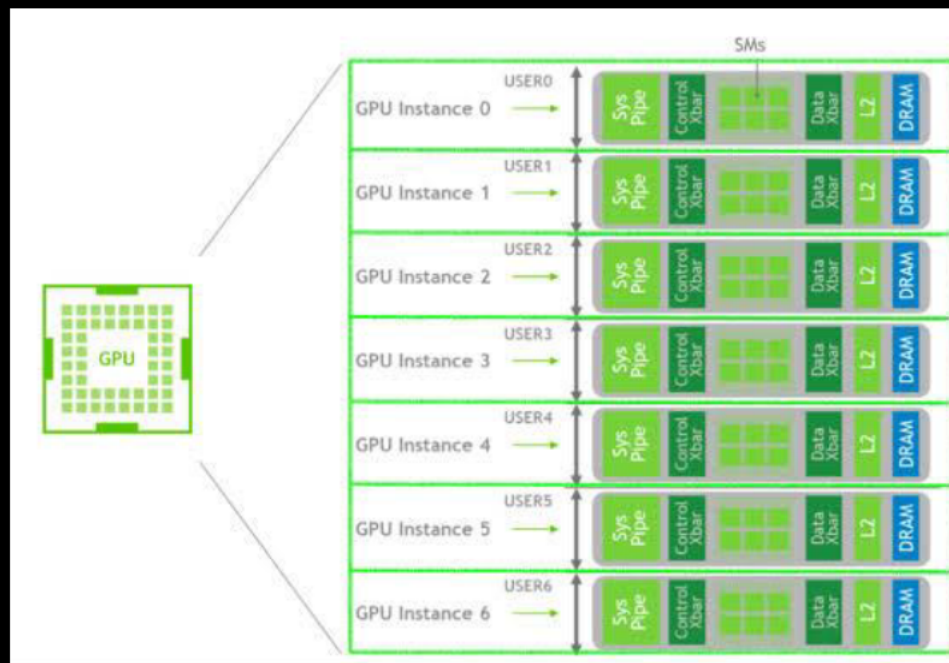


Nvidia A100 MIG Support in FP12

- MIG Introduction
- LSF MIG Support Pre-request
- LSF New Parameters for MIG Support
- New command options for MIG Support
- Use Case of MIG Support

MIG introduction

- NVIDIA releases GPU A100 with Multi-Instance GPU (MIG) feature.



MIG introduction

- GPU Instances (GI)
 - Each GPU can be split to 7 GPU instances(GI) at most.
 - The name of GI indicates the GPU Slices number and GPU memory
“2g.10gb” means the GI have 2 GPU Slices and 10GB GPU memory.

GPU Instance (GI) Profiles			
Profile Name	Fraction of Memory	Fraction of SMs	Number of Instances Available
MIG 1g.5gb	1/8	1/7	7
MIG 2g.10gb	2/8	2/7	3
MIG 3g.20gb	4/8	3/7	2
MIG 4g.20gb	4/8	4/7	1
MIG 7g.40gb	Full	7/7	1

MIG introduction

- Compute Instances (CI)
 - Each GPU instance (GI) can be divided to several GPU Compute Instances(CI)

Memory	20gb			20gb	
GPU Instance	3g			3g	
Compute Instance	1c	1c	1c	2c	1c
MIG Device	1c.3g.20gb	1c.3g.20gb	1c.3g.20gb	2c.3g.20gb	1c.3g.20gb
	GPC	GPC	GPC	GPCGPC	GPC

The above shows that two different “3g.20gb”GPU Instances (GI) can be divided to 3 Compute Instances (CI) named 1c/1c/1c and 2 Compute Instances (CI) named 2c/1c respectively.

MIG introduction

- MIG Device

- A CI and its parent GI can be called a MIG device. And CUDA applications treat the MIG device as a normal GPU.

For example,

- A MIG device called 3g.20gb means the device have 3 GPU slices and 20 GB GPU memory, all compute instance are used.
- A MIG device called 2c.3g.20gb means the device have 3 GPU SM slices and 20 GB GPU memory, and only 2 compute instances(GPU SM slices) of total 3 are used.

LSF MIG Support Pre-request

- Enable Nvidia MIG mode

```
$ sudo nvidia-smi -i 0 -mig 1
Enabled MIG Mode for GPU 00000000:36:00.0
All done.
```

- LSF GPU related parameters setting

```
LSF_GPU_RESOURCE_IGNORE=Y
LSF_GPU_AUTOCONFIG=Y
LSB_GPU_NEW_SYNTAX=extend
LSF_MANAGE_MIG=Y/N
```

LSF New Parameters for MIG Support

- LSF_MANAGE_MIG=N (Default Value)

The GPU instance(GI) and Compute instance(CI) is configured by the administrator on each hosts. LSF dispatches the job to the host based on its GI/CI topology.

For example:

The administrator creates GI/CI with size 1c.3g.20gb, 2c.3g.20gb, and 3g.20gb on host1, and the user submit a job with '-gpu num=1:mig=3/2' request. LSF searches the available MIG resource among the hosts in the cluster, allocate the MIG resource and dispatch the job to host1, run the job on the MIG device 2c.3g.20gb

LSF New Parameters for MIG Support

- `LSF_MANAGE_MIG=Y`

LSF dynamically manages the GI/CI allocation/deallocation for the GPU hosts in the cluster. LSF will create GI/CI according to the job's requirement and destroy GI/CI after job execution finishes.

For example:

The administrator doesn't need to configure any GI/CI on GPU hosts in advance. If the user requests a job 'mig=3/2', LSF allocates a host that MIG 3/2 can be created, dispatch the job to that host, create GI/CI for this job, run job on the MIG device, finally destroy the MIG device.

New command option for MIG Support

- **lshosts -mig**

- Use '*lshosts -gpu -mig*' to check the MIG instances on each GPU hosts.
- '*lshosts -gpu*' indicate which gpu enabled MIG mode on each hosts

```
rladmin@rl-dgxa-c18-u19:~/migtest/lsftop/conf$ lshosts -gpu -mig
```

HOST_NAME	gpu_id	gpu_model	gpu_driver	gpu_factor	numa_id	vendor	devid	gid	cid	inst_name
rl-dgxa-c18	0	TeslaA100_SXM4_	450.51.06	8.0	3	Nvidia	0	11	0	1g.5gb
	0	TeslaA100_SXM4_	450.51.06	8.0	3	Nvidia	1	13	0	1g.5gb
	1	TeslaA100_SXM4_	450.51.06	8.0	3	Nvidia	0	3	0	2g.10gb
	2	TeslaA100_SXM4_	450.51.06	8.0	1	Nvidia	0	2	0	3g.20gb
	3	TeslaA100_SXM4_	450.51.06	8.0	1	Nvidia	0	1	0	4g.20gb
	4	TeslaA100_SXM4_	450.51.06	8.0	7	Nvidia	0	0	0	7g.40gb
	5	TeslaA100_SXM4_	450.51.06	8.0	7	Nvidia	-	-	-	-
	6	TeslaA100_SXM4_	450.51.06	8.0	5	Nvidia	-	-	-	-
	7	TeslaA100_SXM4_	450.51.06	8.0	5	Nvidia	-	-	-	-

```
rladmin@rl-dgxa-c18-u19:~/migtest/lsftop/conf$ lshosts -gpu
```

HOST_NAME	gpu_id	gpu_model	gpu_driver	gpu_factor	numa_id	vendor	mig
rl-dgxa-c18	0	TeslaA100_SXM4_	450.51.06	8.0	3	Nvidia	Y
	1	TeslaA100_SXM4_	450.51.06	8.0	3	Nvidia	Y
	2	TeslaA100_SXM4_	450.51.06	8.0	1	Nvidia	Y
	3	TeslaA100_SXM4_	450.51.06	8.0	1	Nvidia	Y
	4	TeslaA100_SXM4_	450.51.06	8.0	7	Nvidia	Y
	5	TeslaA100_SXM4_	450.51.06	8.0	7	Nvidia	Y
	6	TeslaA100_SXM4_	450.51.06	8.0	5	Nvidia	Y
	7	TeslaA100_SXM4_	450.51.06	8.0	5	Nvidia	Y

Use Case of MIG Support

- Use `'mig=gi_size/ci_size'` in job submission to request a MIG device
- Use `'bjobs -gpu'` shows the GPU_ALLOCATION with MIG info

```
rladmin@rl-dgxa-c18-u19:~/migtest/lsftop/conf$ bsub -gpu num=1:mig=1/1 ../../mig_bin/e06-gpu
Job <416> is submitted to default queue <normal>.
rladmin@rl-dgxa-c18-u19:~/migtest/lsftop/conf$ bjobs
JOBID   USER   STAT  QUEUE      FROM_HOST   EXEC_HOST   JOB_NAME   SUBMIT_TIME
416     rladmin RUN    normal     rl-dgxa-c18 rl-dgxa-c18 *n/e06-gpu Nov 25 06:19
rladmin@rl-dgxa-c18-u19:~/migtest/lsftop/conf$ bjobs -l -gpu

Job <416>, User <rladmin>, Project <default>, Status <RUN>, Queue <normal>, Com
mand <../../mig_bin/e06-gpu>, Share group charged </rladmi
n>
Wed Nov 25 06:19:01: Submitted from host <rl-dgxa-c18-u19-enp226s0>, CWD <$HOME
/migtest/lsftop/conf>, Requested GPU <num=1:mig=1/1>;
Wed Nov 25 06:19:01: Started 1 Task(s) on Host(s) <rl-dgxa-c18-u19-enp226s0>, A
llocated 1 Slot(s) on Host(s) <rl-dgxa-c18-u19-enp226s0>,
Execution Home </home/rladmin>, Execution CWD </home/rladm
in/migtest/lsftop/conf>;

SCHEDULING PARAMETERS:
          r15s  r1m  r15m  ut      pg      io      ls      it      tmp      swp      mem
loadSched  -    -    -    -      -      -      -      -      -      -      -
loadStop   -    -    -    -      -      -      -      -      -      -      -

RESOURCE REQUIREMENT DETAILS:
Combined: select[(ngpus>0) && (type == local)] order[r15s:pg] rusage[ngpus_phy
sical=1.00:mig=1/1]
Effective: select[(ngpus>0)) && (type == local)] order[r15s:pg] rusage[ngpus_
physical=1.00:mig=1/1]

GPU REQUIREMENT DETAILS:
Combined: num=1:mode=shared:mps=no:j_exclusive=yes:gvendor=nvidia:mig=1/1
Effective: num=1:mode=shared:mps=no:j_exclusive=yes:gvendor=nvidia:mig=1/1

GPU_ALLOCATION:
HOST          TASK  GPU_ID  GI_ID/SIZE  CI_ID/SIZE  MODEL          MTOTAL  FACTOR  MRSV  SOCKET  NVLINK/XGMI
rl-dgxa-c18-u19- 0      0      4/1        4/1        TeslaA100_SX  39.5G   8.0     0M     3      -
```

Use Case of MIG Support

- Use `'bhosts -gpu -l'` to shows the MIG info on each hosts.

```
rladmin@rl-dgxa-c18-u19:~/migtest/lsftop/conf$ bhosts -gpu -l
```

HOST: rl-dgxa-c18-u19-enp226s0
 NGPUS 8
 NGPUS_SHARED_AVAIL 8
 NGPUS_EXCLUSIVE_AVAIL 8

STATIC ATTRIBUTES							
GPU_ID	MODEL	MTOTAL	FACTOR	SOCKET	VENDOR	MIG	NVLINK/XGMI
0	TeslaA100_SXM4_40GB	39.5G	8.0	3	Nvidia	Y	-/N/N/N/N/N/N/N
1	TeslaA100_SXM4_40GB	39.5G	8.0	3	Nvidia	Y	N-/N/N/N/N/N/N/N
2	TeslaA100_SXM4_40GB	39.5G	8.0	1	Nvidia	Y	N/N-/N/N/N/N/N/N
3	TeslaA100_SXM4_40GB	39.5G	8.0	1	Nvidia	Y	N/N/N-/N/N/N/N/N
4	TeslaA100_SXM4_40GB	39.5G	8.0	7	Nvidia	Y	N/N/N/N-/N/N/N/N
5	TeslaA100_SXM4_40GB	39.5G	8.0	7	Nvidia	Y	N/N/N/N/N-/N/N/N
6	TeslaA100_SXM4_40GB	39.5G	8.0	5	Nvidia	Y	N/N/N/N/N/N-/N/N
7	TeslaA100_SXM4_40GB	39.5G	8.0	5	Nvidia	Y	N/N/N/N/N/N/N/-

DYNAMIC ATTRIBUTES										
GPU_ID	MODE	MUSED	MRSV	TEMP	ECC	UT	MUT	PSTATE	STATUS	ERROR
0	SHARED	87M	0M	31C	0	0%	0%	0	ok	-
1	SHARED	7M	0M	27C	0	0%	0%	0	ok	-
2	SHARED	11M	0M	28C	0	0%	0%	0	ok	-
3	SHARED	14M	0M	28C	0	0%	0%	0	ok	-
4	SHARED	0M	0M	31C	0	0%	0%	0	ok	-
5	SHARED	0M	0M	30C	0	0%	0%	0	ok	-
6	SHARED	0M	0M	30C	0	0%	0%	0	ok	-
7	SHARED	0M	0M	30C	0	0%	0%	0	ok	-

GPU JOB INFORMATION						
GPU_ID	JEXCL	RUNJOBIDS	SUSPJOBIDS	RSVJOBIDS	GI_ID/SIZE	CI_ID/SIZE
0	Y	416	-	-	4/1	4/1
1	-	-	-	-	-	-
2	-	-	-	-	-	-
3	-	-	-	-	-	-
4	-	-	-	-	-	-
5	-	-	-	-	-	-
6	-	-	-	-	-	-
7	-	-	-	-	-	-

Use Case of MIG Support

- Request MIG device with GPU memory.

- *“bsub -gpu num=1:mig=3/2:gmem=5G gpu_app”*

It will request 1 MIG device with GI size 5 CI size 2 and 5G memory on this GPU Instance(GI).

- Request share/exclusive MIG device.

- *“bsub -gpu num=1:mig=2/1:mode=shared(or mode=exclusive_process or j_exclusive=yes) gpu_app”*

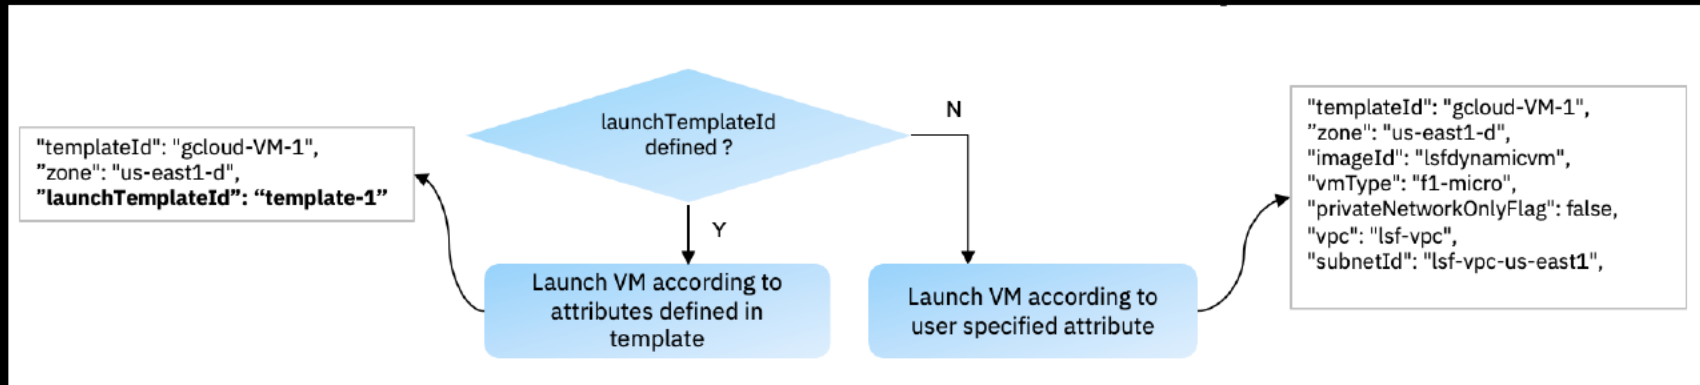
Although in MIG mode, only shared mode are supported, LSF has been enhanced to use MIG device exclusively or shared.

Resource Connector Enhancement for Google Cloud



Launch Instance Template

- Instance template is created on Google Cloud Platform portal
- Use **launchTemplateId** in googleprov_templates.json to enable support of launch instance template
- Properties of the instance can be defined within the instance template
- Properties defined in googleprov_templates.json override instance template



Local SSDs

- Local SSDs are supported through instance template, not in googleprov_templates.json
- In Google Cloud, attach local SSDs to instance template in **Disks _> Add new disk**
- SSDs must be mounted before being used. LSF provides an example of how to mount them in user_data: `<LSF_TOP>/<LSF_VERSION>/resource_connector/google/scripts/example_user_data.sh`

Preemptable VM instances

- Preemptable VM instances are supported through instance template, not in googleprov_templates.json
- In Google Cloud, set the **Preemptibility** field to **On** when creating the launch instance template to enable preemptible VM instances.
- When an instance is preempted, LSF automatically requeue the jobs running on it by default.
 - Using RC_RECLAIM_ACTION=REQUEUE | TERMINATE to change the default behavior

Bulk instance APIs (Bulk API endpoints)

- With bulk API, LSF resource connector creates VM instances and check status in batch
- Specify zone to let LSF call **zonal bulk API**. Instances launch in the specified zone
- Leave zone undefined to let LSF call **regional bulk API**. Zone of the instance is selected by Google Compute Engine based on hardware availability.
- Use **G_CLOUD_REGION** in googleprov_config.json to define the region. Can be override by region in googleprov_template.json

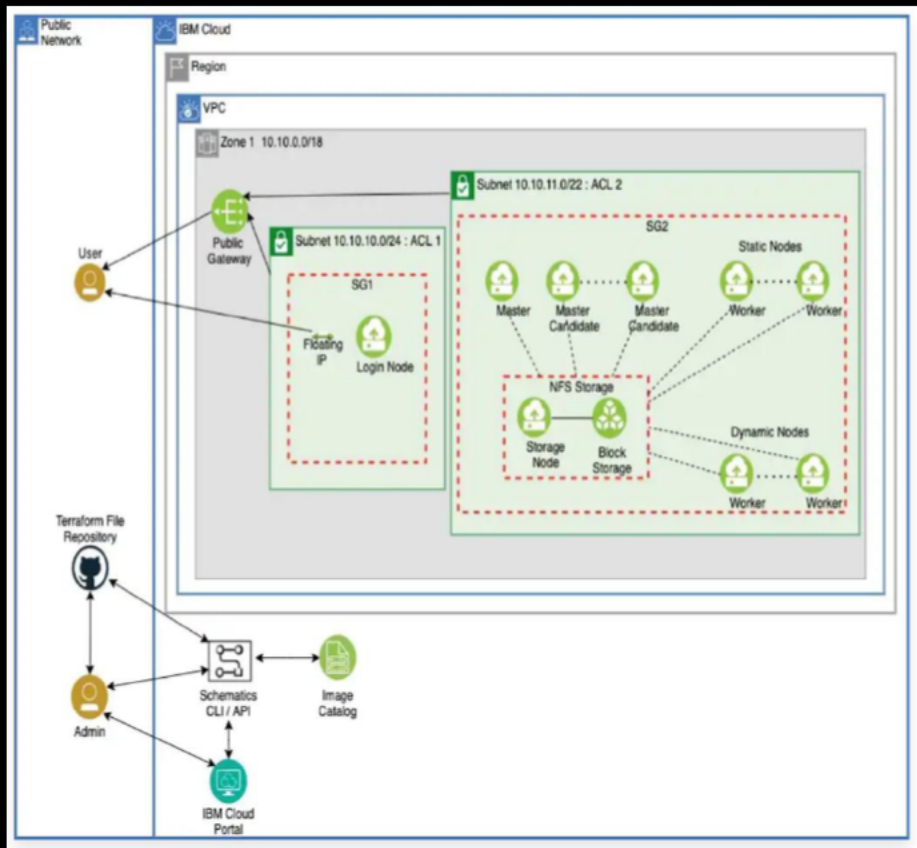
Quick Deployment of LSF on IBM Cloud



Quick Deployment of LSF on IBM Cloud

- Deploy LSF with Resource Connector on IBM Cloud with a few clicks
- Automated deployment for compute, network and storage resources
- Easy to auto-scale up and down based on demand
- Allows seamless bursting to IBM Cloud from on-premise clusters with LSF Multi-Cluster
- Transfer data between clusters intelligently with LSF Data Manager

Quick Deployment of LSF on IBM Cloud



IBM Cloud

Search resources and offerings...

Catalog

Q Spectrum LSF

IBM Spectrum LSF

Services

Explore our broad portfolio of managed services for infrastructure, developer tools, and

Internal IBM pricing displayed
Classic infrastructure services might reflect internal IBM pricing, which should not be shared with external clients. Log out of your internal IBM Cloud account to view external pricing.

Parameter	Description	Value
api_key	This is the API key for IBM Cloud account in which the Spectrum LSF cluster needs to be deployed. Learn more.	<input type="text" value="Enter api_key"/>
ssh_key_name	Name of the SSH key configured in your IBM Cloud account that is used to establish a connection to the LSF master node. If you do not have an SSH key in your IBM Cloud account, create one by using the instructions given here .	<input type="text" value="Enter ssh_key_name"/>
lsf_license_confirmation	Please confirm your use of IBM Spectrum LSF licenses. By entering "true" for the property you have agreed to one of the two conditions. 1. You are using the software in production and confirm you have sufficient licenses to cover your use under the International Program License Agreement (IPLA). 2. You are evaluating the software and agree to abide by the International License Agreement for	<input type="text" value="Enter lsf_license_confirmation"/>

Deprecated LSF Features in FP12



Deprecated LSF Features in FP12

Deprecated feature	Corresponding deprecated parameters and commands	Alternative feature
LSF multicluster capability resource leasing model	<ul style="list-style-type: none">• <code>lsb.resources</code> file:<ul style="list-style-type: none">• HostExport section• SharedResourceExport section• <code>lsb.queues</code> file:<ul style="list-style-type: none">• HOSTS: <code>allremote</code> and <code>all@cluster_name</code> keywords	Updated LSF default behavior.
LSF/XL feature and LSF Advanced Edition.	<ul style="list-style-type: none">• -cname option for the following commands: bacct, bhosts, bjobs, bmgroup, lshosts, lsload.	Updated LSF default behavior.

Deprecated LSF Features in FP12

Deprecated feature	Corresponding deprecated parameters and commands	Alternative feature
Job slot pools for fairshare scheduling	<ul style="list-style-type: none">lsb.queues file:<ul style="list-style-type: none">MAX_SLOTS_IN_POOLSLOT_POOLSLOT_SHAREUSE_PRIORITY_IN_POOL	Guarantee SLA (guarantee service class).
toplib integrations	<ul style="list-style-type: none">lsf.conf file:<ul style="list-style-type: none">LSF_TOPD_PORTLSF_TOPD_TIMEOUT	
Chunk job scheduling	<ul style="list-style-type: none">lsf.conf file:<ul style="list-style-type: none">LSB_CHUNK_RUSAGElsb.params file:<ul style="list-style-type: none">CHUNK_JOB_DURATIONlsb.applications and lsb.queues files:<ul style="list-style-type: none">CHUNK_JOB_SIZE	Use the RELAX_JOB_DISPATCH_ORDER parameter in the lsb.params file to enable multiple jobs with common resource requirements to run consecutively on the same allocation.

Deprecated LSF Features in FP12

Deprecated feature	Corresponding deprecated parameters and commands		Alternative feature
<p>HPC integrations for the following environments:</p> <ul style="list-style-type: none">- Parallel job support using PAM- IBM Parallel Environment (IBM PE) (See Running MPI workload through IBM Parallel Environment Runtime Edition)- SGI CPuset	<ul style="list-style-type: none">• <code>lsf.conf</code> file:<ul style="list-style-type: none">• <code>LSF_PAM_APPL_CHKPNT</code>• <code>LSF_PAM_CLEAN_JOB_DELAY</code>• <code>LSF_PAM_HOSTLIST_USE</code>• <code>LSF_PAM_PLUGINDIR</code>• <code>LSF_PAM_USE_ASH</code>• <code>LSF_PE_NETWORK_NUM</code>• <code>LSF_PE_NETWORK_UPDATE_INTERVAL</code>• <code>LSB_CPuset_BESTCPUS</code>• <code>LSB_CPuset_DISPLAY_CPULIST</code>• <code>LSF_CPusetLIB</code>• <code>lsb.applications</code> file:<ul style="list-style-type: none">• <code>NETWORK_REQ</code>	<ul style="list-style-type: none">• <code>lsb.params</code> file:<ul style="list-style-type: none">• <code>MAX_PROTOCOL_INSTANCES</code>• <code>NETWORK_REQ</code>• <code>STRIPING_WITH_MINIMUM_NETWORK</code>• <code>lsb.queues</code> file:<ul style="list-style-type: none">• <code>MAX_PROTOCOL_INSTANCES</code>• <code>NETWORK_REQ</code>• <code>STRIPING_WITH_MINIMUM_NETWORK</code>• <code>bsub -network</code> command option	<p>Use the <code>blaunch</code> command for parallel jobs.</p>

Deprecated LSF Features in FP12

Deprecated feature	Corresponding deprecated parameters and commands	Alternative feature
Automatic CPU frequency selection in energy aware scheduling (See Automatic CPU frequency selection)		
Relaxed syntax for resource requirement selection strings	<ul style="list-style-type: none">lsf.conf file:<ul style="list-style-type: none">LSF_STRICT_RESREQ: Now fixed to Y.	Updated LSF default behavior (strict syntax for resource requirement selection strings).
Task list files and related commands	<ul style="list-style-type: none">lsf.task filelsf.conf file:<ul style="list-style-type: none">LSF_SHELL_AT_USERSch commandlsfmon commandlslogin commandlsmon commandlstcsh commandlsltasks commandlsrtasks commandlselectible command	

LSF Features to Be Deprecated in the Next Release

Feature to be deprecated in the next LSF release	Corresponding deprecated parameters and commands	Alternative feature
SLA scheduling except guarantee SLA (See Using goal-oriented SLA scheduling).	<ul style="list-style-type: none">lsb.params file:<ul style="list-style-type: none">ENABLE_DEFAULT_EGO_SLAlsb.serviceclasses file:<ul style="list-style-type: none">GOALS: THROUGHPUT, VELOCITY, and DEADLINE keywords.	
Slot and package guarantee policy	<ul style="list-style-type: none">lsb.params file:<ul style="list-style-type: none">SIMPLIFIED_GUARANTEES: Will be fixed to Y.	Guarantee SLA (guarantee service class), which will be the default behavior.
GPU scheduling and features that use ELIM	<ul style="list-style-type: none">lsf.conf file:<ul style="list-style-type: none">LSB_GPU_NEW_SYNTAX: Will be fixed to extend.LSF_GPU_RESOURCE_IGNORE: Will be fixed to Y.LSF_GPU_AUTOCONFIG: Will be fixed to Y.All elim.gpu.* ELIMS:	Updated LSF default behavior (GPU autoconfiguration).

References

Release Notes of FP12 -

<https://www.ibm.com/docs/en/spectrum-lsf/10.1.0?topic=wn-whats-new-in-lsf-101-fix-pack-12>

Download Link of FP12 -

<http://www.ibm.com/support/fixcentral/swg/selectFixes?product=ibm/Other+software/IBM+Spectrum+LSF&release=All&platform=All&function=fixId&fixids=lsf-10.1.0.12-spk-2021-Jun-build600488&includeSupersedes=0>

Detailed Deprecated Features in FP12 -

<https://www.ibm.com/docs/en/spectrum-lsf/10.1.0?topic=12-deprecated-features>

LSF Deployment on IBM Cloud -

<https://cloud.ibm.com/docs/ibm-spectrum-lsf?topic=ibm-spectrum-lsf-about-spectrum-lsf>

Nvidia A100 MIG user guide -

<https://docs.nvidia.com/datacenter/tesla/mig-user-guide/index.html>

Thank you

Yun Gao
LSF Support Professional

—

bjyungao@cn.ibm.com

© Copyright IBM Corporation 2021. All rights reserved. The information contained in these materials is provided for informational purposes only, and is provided AS IS without warranty of any kind, express or implied. Any statement of direction represents IBM's current intent, is subject to change or withdrawal, and represent only goals and objectives. IBM, the IBM logo, and ibm.com are trademarks of IBM Corp., registered in many jurisdictions worldwide. Other product and service names might be trademarks of IBM or other companies. A current list of IBM trademarks is available at [Copyright and trademark information](#).

