

IBM Platform Computing Cloud Services

Dino Quintero

Daniel de Souza Casali

Marcelo Correia Lima

Istvan Gabor Szabo

Maciej Olejniczak

Tiago Rodrigues de Mello

Nilton Carlos dos Santos



 **Cloud**

Infrastructure Solutions



Introduction

This IBM® Redpaper™ publication provides a description of the IBM Platform Computing Cloud Services and describes a scenario about how IBM Platform LSF® multicluster and IBM Spectrum™ Scale Active File Manager help manage the usage of cloud services in an efficient way.

This paper provides factual and comparative economics of deploying a solution on-premises versus in the cloud, and includes information about the benefits of on-premises versus cloud solution deployments.

This paper covers the following topics:

- ▶ IBM Platform Computing Cloud Services: Purpose and benefits
- ▶ IBM Platform Computing Cloud Services architecture
- ▶ IBM Spectrum Scale high performance services
- ▶ IBM Platform Symphony services
- ▶ IBM High Performance Services for Hadoop
- ▶ Platform LSF
- ▶ Hybrid Platform LSF on-premises with a cloud service scenario
- ▶ Data management on hybrid clouds

IBM Platform Computing Cloud Services: Purpose and benefits

Engineering, scientific, financial, or research workloads are not the only demanding workloads for technical and High Performance Computing (HPC) infrastructures. Big data challenges are solved by using the same method, distributing the workload across multiple machines within a technical computing cluster.

Meeting all these demands can be especially challenging for organizations that have seasonal or unpredictable demand spikes, need access to additional compute or storage resources to support a growing business, or are starting to use these technologies. The time that it takes to respond to a critical market analysis, a product release, or a research study can be impacted by resource availability, which affects competitiveness and profitability.

Organizations can quickly and efficiently overcome these challenges by combining market-leading workload management from IBM Platform Computing with the efficiency and cost benefits of cloud computing.

IBM Platform Computing Cloud Services running on the IBM SoftLayer® cloud delivers a versatile, high-performing cloud-based environment to fuel your organization's growth, whether you are engaged in any of the following activities:

- ▶ Seeking to meet variable workload demands
- ▶ In need of clustered resources, but do not have the budget or in-house skills to deploy and use a technical computing infrastructure
- ▶ Running out of data center space and must continue to grow compute and storage capacity
- ▶ Considering to provide applications on a pay-per-use basis, but do not have the infrastructure or time to create a service

If any of these activities are important to you, you can count on the benefits that are delivered by the IBM Platform Computing Cloud Services offering to meet your needs:

- ▶ Ready-to-use IBM Platform LSF and IBM Platform Symphony® clusters in the cloud reduce time to results and accelerate time to market.
- ▶ A high-performance file system with IBM Spectrum Scale that is delivered as a service improves data management and provides seamless transfer between on-premises and cloud infrastructures.
- ▶ Non-shared physical machines, InfiniBand interconnect, the latest processor technology, and your choice of a SoftLayer data center leads to optimal application performance and security.
- ▶ Integrated workload management with both on-premises and on-cloud infrastructures simplifies management and the user experience, and full support from IBM's technical computing experts reduces administrative impact.

IBM Platform Computing Cloud Services architecture

IBM Platform Computing Cloud Services is built on top of SoftLayer. SoftLayer deploys the infrastructure in its data centers in the form of *Points of Delivery (PODs)*, which are groups of thousands of machines, petabytes of storage, and all the networking, firewalls, power distribution, and internet connectivity that is needed to support this infrastructure. Theoretically, a customer might use an entire POD, which is more than 60,000 processor cores on a single cluster. A customer can request other PODs to meet their demands. Usually, a cluster starts much smaller than at the POD scale, and flexes up or down as their needs dictate. All configurations that are delivered by IBM Platform Computing Cloud Services deliver exclusive, non-shared server usage for the client. The IBM Platform Computing Cloud Services solution offers a true cloud-based consumption model: Pay by the hour or by the month for most elements of the service.

IBM Platform Computing Cloud Services is a purpose-built Software as a Service (SaaS) where clients can use ready-to-use clusters that are available either for usage by the hour or by the month. The service is owned and operated by the IBM Platform Computing team, which ensures that the HPC experts are available to provide management and support of your chosen environment. The service uses IBM Platform Computing HPC management and scheduling tools (Platform LSF and Platform Symphony), which provide optimum performance and user experience.

The SaaS architecture counts on the two IBM Platform Computing schedulers, Platform LSF and Platform Symphony, for either traditional HPC clusters or service-oriented architectures (SOA). The offering is provided with or without IBM Spectrum Scale (formerly GPFS™), which can ease and reduce the data transfer that is needed to and from the cloud by using the Spectrum Scale Active File Management (AFM) facility for caching only the needed files to run the workload on the remote site in the case of a hybrid cloud. The IBM Platform Computing Cloud Services high-level architecture is shown in Figure 1.

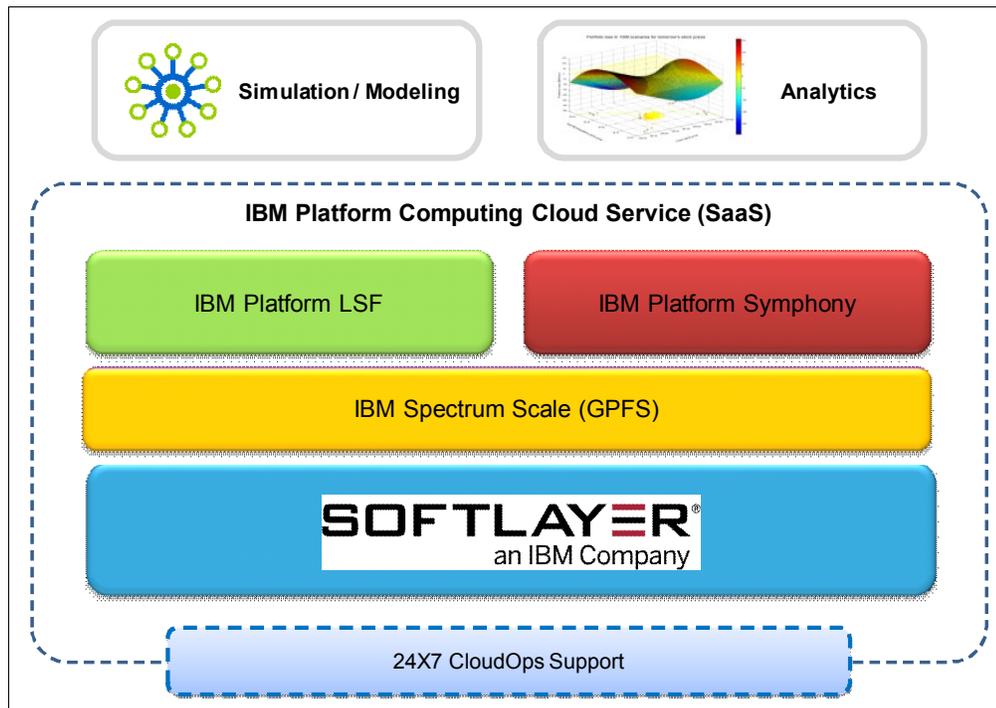


Figure 1 IBM Platform Computing Cloud Services - high-level architecture

IBM Spectrum Scale high performance services

For clients that consider adding storage capacity or who require more performance and scalability than a network file system (NFS) can provide, IBM Spectrum Scale is now available as a service on the SoftLayer cloud as part of the IBM Platform Computing Cloud Services.

Optimized for technical computing and analytics workloads, Spectrum Scale in the cloud enables seamless transfer of files between local and cloud-based resources by using the Spectrum Scale AFM feature.

With the addition of Spectrum Scale in the cloud, the IBM Platform Computing Cloud Services enables speedy deployment of fully supported, ready-to-run technical computing or analytic environments in the cloud.

Organizations that use IBM Platform Computing Cloud Services can easily meet additional resource demands without the cost of purchasing or managing an in-house infrastructure, which minimizes the administrative burden and quickly addresses evolving business needs.

IBM Platform Symphony services

Although a benefit of IBM Platform Symphony is its ability to support diverse applications in a multi-tenant environment while ensuring each application or tenant service level requirement is satisfied, performance tests show that IBM Platform Symphony also helps to provide better performance and efficiency, and superior management and monitoring.

If you do not have a specific application to run on Platform Symphony, for example, but you need a service environment for your Hadoop workload, see “IBM High Performance Services for Hadoop” on page 4.

For more information about how IBM Platform Symphony can help improve your Hadoop workload, see the following website:

<http://www.ibm.com/systems/platformcomputing/products/symphony/highperfhadoop.html>

IBM High Performance Services for Hadoop

IBM High Performance Services for Hadoop is suitable for organizations that are looking for a fully supported, ready-to-run Hadoop environment for production use, or as a development and testing environment. This service enables customers to deploy quickly and easily Hadoop workloads on ready-to-run clusters on the SoftLayer cloud, complete with a bare metal SoftLayer infrastructure, a private network, and your choice of data center to help achieve optimal performance and security.

An experienced and dedicated cloud operations team configures, deploys, and supports the cloud-based infrastructure and the software, which helps minimize the administrative burden on your organization and the need to develop the skills to design and manage a Hadoop environment.

IBM High Performance Services for Hadoop delivers a Hadoop-ready cluster as a service on SoftLayer and helps deliver the following benefits:

- ▶ Rapid access to Hadoop clusters in the cloud for both product use and development testing
- ▶ Optimal performance with bare metal resources
- ▶ Security through physical isolation and choice of data center location
- ▶ Reduced capital expenditure
- ▶ Minimal user and administrator impact
- ▶ Easy adoption of public cloud technology and resources

IBM High Performance Services for Hadoop can deliver the following benefits:

- ▶ More capability and lower costs: Easily meet demand without the upfront costs of purchasing an in-house infrastructure or the ongoing cost of infrastructure management.
- ▶ Match resources to demand, while helping reduce capital expenditures: Help minimize administrative costs by using a skilled cloud operations team with deep Hadoop expertise.
- ▶ Security:
 - Help achieve security through physical isolation with a dedicated virtual local area network (VLAN).
 - Upload data securely through a virtual private network (VPN) or Multi-Protocol Label Switching (MPLS) to gateway servers.
 - Use your SoftLayer data center of choice for regulatory compliance.
- ▶ Faster time to results:
 - Accelerate Hadoop MapReduce workloads with dedicated bare metal servers.
 - Optimize I/O performance with 10-gigabit Ethernet networking.

Platform LSF

Platform LSF is a powerful workload management platform for demanding, distributed HPC environments. It provides a comprehensive set of intelligent, policy-driven scheduling features that enable you to use all of your compute infrastructure resources and ensure optimal application performance.

Platform LSF helps to ensure that all available resources are fully used by enabling you to take full advantage of all technical computing resources in the cloud. Platform LSF helps to ensure that the computing power in the cloud is fully used, and it helps to manage application software licenses usage, which is usually expensive for demanding workloads.

Platform LSF provides the following features:

- ▶ A single source for end-to-end cluster support with access to technical computing experts to eliminate the skills barrier for using clustered resources.
- ▶ Dedicated bare-metal servers and InfiniBand interconnect for applications that require the full capacity of a non-virtualized, and parallel computing environment.

- ▶ Control of data center locality, enabling organizations to choose the location where workloads run to protect their information and meet data regulations.
- ▶ Non-shared physical machines and dedicated network for workloads requiring maximum security.

Hybrid Platform LSF on-premises with a cloud service scenario

A transparent user experience that manages workloads between an on-premises cluster and the cloud can be achieved with IBM Platform LSF Multicluster and IBM Spectrum Scale AFM.

Note: If you already have IBM Platform LSF Standard Edition, skip the section “Upgrading IBM Platform HPC to enable multicluster function”. Otherwise, see *IBM Platform Computing Integration Solutions, SG24-8081* and *IBM Platform Computing Solutions Reference Architectures and Best Practices, SG24-8169* for details about how to implement IBM Platform HPC.

Upgrading IBM Platform HPC to enable multicluster function

To start the upgrade, you need the name of your Platform LSF installation directory (LSF_TOP), the Platform LSF administrators (LSF_ADMINS), and the cluster name (LSF_CLUSTER_NAME) available. If you do not have this information, run the commands that are shown in Example 1 to gather the information.

Example 1 Gather information for Platform LSF standard edition upgrade

```
[root@homecluster etc]# grep LSF_TOP $PCMD_TOP_LOCAL/etc/lsf.install.config
LSF_TOP="/shared/ibm/platform_lsf"
[root@homecluster etc]# grep LSF_ADMINS $PCMD_TOP_LOCAL/etc/lsf.install.config
LSF_ADMINS="phpcadmin root"
[root@homecluster etc]#
[root@homecluster etc]# grep CLUSTER $PCMD_TOP_LOCAL/etc/lsf.install.config
LSF_CLUSTER_NAME="phpc_cluster"
[root@homecluster etc]#
```

The information that is gathered by the commands in Example 1 is necessary to upgrade Platform HPC and install Platform LSF Standard Edition into the cluster.

To start, add the parameters to the `install.config` file from your Platform LSF installation directory, as shown in Example 2. You must add the path for the Platform LSF Standard Edition entitlement file of the installation configuration file, for example:

```
LSF_ENTITLEMENT_FILE="/tmp/phpc/platform_lsf_std_entitlement.dat"
```

Example 2 Configuration file to install Platform LSF (install.config)

```
*****
#                               LSF 9.1.3 INSTALL.CONFIG FILE
*****
#
# Name:      install.config
#
# Purpose:  LSF installation options
#
```

```

# $Id$
#
# File Format:
#   o Options (without # sign) can only appear once in the file.
#   o Blank lines and lines starting with a number sign (#) are ignored.
#
# Option Format:
#   o Each disabled example looks like this:
#     # -----
#     # LSF_OPTION_NAME="EXAMPLE_VALUE"
#     # -----
#
#   o An enabled option looks like this:
#     # -----
#     LSF_OPTION_NAME="ACTUAL_VALUE"
#     # -----
#
# Instructions:
#   1. Edit install.config to specify the options for
#      your cluster. Uncomment the options that you want and
#      replace the EXAMPLE values with your own settings.
#      The sample values shown in this template
#      are EXAMPLES only. They are not always the default
#      installation values.
#
#   2. Run ./lsfinstall -f install.config
#
#
#*****
#           PART 1: REQUIRED PARAMETERS
#   (During an upgrade, specify the existing value.)
#*****
# -----
LSF_TOP="/shared/ibm/platform_lsf"
# -----
# Full path to the top-level installation directory {REQUIRED}
#
# The path to LSF_TOP must be shared and accessible to all hosts
# in the cluster. It cannot be the root directory (/).
# The file system containing LSF_TOP must have enough disk space for
# all host types (approximately 300 MB per host type).
#
# -----
LSF_ADMINS="phpcadmin root"
# -----
# List of LSF administrators {REQUIRED}
#
# The first user account name in the list is the primary LSF
# administrator. It cannot be the root user account.
# Typically, this account is named lsfadmin.
# It owns the LSF configuration files and log files for job events.
# It also has permission to reconfigure LSF and to control batch
# jobs that are submitted by other users. It typically does not have
# authority to start LSF daemons. Usually, only root has
# permission to start LSF daemons.

```

```

# All the LSF administrator accounts must exist on all hosts in the
# cluster before you install LSF.
# Secondary LSF administrators are optional.
#
# -----
LSF_CLUSTER_NAME="phpc_cluster"
# -----
# Name of the LSF cluster {REQUIRED}
#
# It must be 39 characters or less, and cannot contain any
# white spaces. Do not use the name of any host, user, or user group
# as the name of your cluster.
#
#
#*****
#           PART 2: PRIMARY PARAMETERS FOR NEW INSTALLATION
# (These parameters are ignored if they are already defined in the cluster.)
#*****
#
# -----
# LSF_MASTER_LIST="hostm hosta hostc"
# -----
# List of LSF server hosts to be master or master candidate in the
# cluster {REQUIRED when you install for the first time or during
# upgrade if the parameter does not already exist.}
#
# You must specify at least one valid server host to start the
# cluster. The first host that is listed is the LSF master host.
#
# -----
LSF_ENTITLEMENT_FILE="/tmp/phpc/platform_lsf_std_entitlement.dat"
# -----
# You must specify a full path to the LSF entitlement file.
#
...

```

To perform the update after you enter all the environment variables in the configuration file (see Example 2 on page 6), run `lsfinstall`, as shown in Example 3.

Example 3 Run the lsfinstall command

```
[root@homecluster lsf9.1.3_lsfinstall]# ./lsfinstall -f install.config
```

Logging installation sequence in /tmp/phpc/lsf9.1.3_lsfinstall/Install.log

International Program License Agreement

.
.
.

Press Enter to continue viewing the license agreement, or enter "1" to accept the agreement, "2" to decline it, "3" to print it, "4" to read non-IBM terms, or "99" to go back to the previous screen.

Read and accept the license agreement to proceed with the installation. After you finish reading and agreeing to the terms, press the number 1 key, and Platform LSF checks for the prerequisites. If the prerequisites are met, the installer prompts for the distribution TAR file to be used, as shown in Example 4.

Example 4 LSF pre-installation check and distribution selection

LSF pre-installation check ...

Checking the LSF TOP directory /shared/ibm/platform_lsf ...
... Done checking the LSF TOP directory /shared/ibm/platform_lsf ...

You are installing IBM Platform LSF - 9.1.3 Standard Edition.

Checking LSF Administrators ...
LSF administrator(s): " phpcadmin root"
Primary LSF administrator: "phpcadmin"
Checking the configuration template ...
CONFIGURATION_TEMPLATE not defined. Using DEFAULT template.
Done checking configuration template ...
Done checking ENABLE_STREAM ...

Checking the patch history directory ...
... Done checking the patch history directory /shared/ibm/platform_lsf/patch ...

Checking the patch backup directory ...
... Done checking the patch backup directory /shared/ibm/platform_lsf/patch/backup ...

Searching LSF 9.1.3 distribution tar files in /tmp/phpc Please wait ...

1) linux2.6-glibc2.3-x86_64

Press 1 or Enter to install this host type: 1

The installation proceeds without further prompts until a message is displayed similar to the one that is shown in Example 5.

Example 5 Installation completed successfully

You have chosen the following tar file(s):
lsf9.1.3_linux2.6-glibc2.3-x86_64

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

Pre-installation check report saved as text file:
/tmp/phpc/lsf9.1.3_lsfinstall/prechk.rpt.

... Done LSF pre-installation check.

.
. .
.

```
Creating lsf_quick_admin.html ...
... Done creating lsf_quick_admin.html
```

lsfinstall is done.

To complete your LSF installation and get your cluster "phpc_cluster" up and running, follow the steps in "/tmp/phpc/lsf9.1.3_lsfinstall/lsf_getting_started.html".

After setting up your LSF server hosts and verifying your cluster "phpc_cluster" is running correctly, see "/shared/ibm/platform_lsf/9.1/lsf_quick_admin.html" to learn more about your new LSF cluster.

After installation, remember to bring your cluster up to date by applying the latest updates and bug fixes.

Note: For the latest release information about Platform LSF Version 9.1.3, see Platform LSF in the IBM Knowledge Center at the following website:

http://www.ibm.com/support/knowledgecenter/SSETD4_9.1.3/lsf_welcome.html

For the latest information about IBM Platform HPC, see IBM Platform HPC in the IBM Knowledge Center at the following website:

http://www-01.ibm.com/support/knowledgecenter/SSENW4_4.2.0/release_notes/release_notes.dita

Now, restart Platform HPC services to enable the new Platform LSF entitlement, as shown in Example 6.

Example 6 Restart Platform HPC services

```
[root@homecluster platform_lsf]# service phpc stop
Stopping Web Portal services          [ OK ]
Stopping PERF services                [ OK ]
Stopping Rule Engine service          [ OK ]
Stopping PCMD service                 [ OK ]
Stopping Message broker                [ OK ]
Stopping the LSF subsystem             [ OK ]

Stopping Platform HPC Services:       [ OK ]
[root@homecluster platform_lsf]# service phpc start
Checking for xcatd service started     [ OK ]

Starting the LSF subsystem             [ OK ]
- Waiting for EGO service started ..   [ OK ]

Cluster name : phpc_cluster EGO master host name : homecluster EGO master version
: 1.2.10
- Waiting for PCM master node online ..... [ OK ]

Starting PERF services                 [ OK ]
Starting Message broker                 [ OK ]
Starting PCMD service                   [ OK ]
```

```

Starting Rule Engine service           [ OK ]
Starting Web Portal services          [ OK ]

Starting Platform HPC Services:       [ OK ]
[root@homecluster platform_lsf]#

```

Your cluster is ready to be configured as a multicluster server.

Tasks to install Platform LSF in the cloud

The IBM Platform Computing Cloud Services team installs and configures the Platform LSF cluster for you in the cloud. You have to set up only your connection to the new cluster and configure your new multicluster feature.

Note: The connection to the cloud network can be done with a VPN or with MPLS. To connect to the cloud, you need to add to your Domain Name System (DNS) or hosts file only the name of the master and master candidates of the new cloud cluster, and exchange the SSH keys between the hosts.

Configuring the multicluster feature

After you have exchanged the SSH keys, enable the multicluster feature. Copy both of your cluster definitions files (on-premises and in the cloud) on both `$LSF_TOP/conf/` directories, as shown in Example 7.

Example 7 Copy cluster definition files between master nodes

```

[root@homecluster conf]# scp
softlayer:/usr/share/lsf/conf/lsf.cluster.HPC_Services
/shared/ibm/platform_lsf/conf/
lsf.cluster.HPC_Services
100% 1801 1.8KB/s 00:00
[root@homecluster conf]# scp
/shared/ibm/platform_lsf/conf/lsf.cluster.phpc_cluster
softlayer:/usr/share/lsf/conf/
lsf.cluster.phpc_cluster 100% 2897 2.8KB/s 00:00

```

Now edit the `$LSF_TOP/conf/lsf.shared` file and check that all the clusters are defined in the cluster stanza file, as shown in Example 8.

Example 8 LSF shared configuration file containing both clusters

```

# $Revision$Date$
# -----
# T H I S   F I L E:  Is shared by all clusters in the LSF system.
#
# This file contains all definitions that are referenced by individual
# lsf.cluster.<clustername> files. The definitions in this file can be
# a superset, i.e., not all definitions in this file need to be used in
# other files.
#
# See lsf.cluster(5) and "LSF User's and Administrator's Guide".
# -----

```

```

Begin Cluster
ClusterName                # Keyword
phpc_cluster
HPC_Services
End Cluster

```

Note: Make the `lsf.shared` file the same on both clusters.

Now, as shown in Example 9, add a module in the `lsb.modules` file in the local cluster to see resources in the remote cluster. In this case, the file is in the homecluster server at the following path:

```
/install/shared/ibm/platform_lsf/conf/lsbatch/phpc_cluster/configdir/lsb.modules
```

Example 9 Add `schmod_mc` on `lsb.modules`

```

# $Revision$Date$

# Define plug-ins for Scheduler and Resource Broker.
# SCH_PLUGIN column specifies the share module name for Scheduler, while
# RB_PLUGIN specifies the share module name for Resource Broker
# A Scheduler plug-in can have one, multiple, or none RB plug-ins
# corresponding to it.
# SCH_DISABLE_PHASES specifies which phases of that scheduler plug-in
# should be disabled, i.e., deactivated. A scheduler plug-in has four phases:
# preprocessing, match/limit, order/alloc, post processing. Scheduler
# won't start disabled phases over jobs
# Note all share modules should be put under LSF_LIBDIR

Begin PluginModule
SCH_PLUGIN                RB_PLUGIN                SCH_DISABLE_PHASES
schmod_default            ()                    ()
schmod_fcfs               ()                    ()
schmod_fairshare          ()                    ()
schmod_limit              ()                    ()
schmod_mc                 ()                    ()
schmod_parallel           ()                    ()
schmod_reserve            ()                    ()
schmod_preemption         ()                    ()
schmod_advrsv             ()                    ()
schmod_ps                 ()                    ()
#schmod_dc                ()                    ()
End PluginModule

```

Restart the Platform LSF services on both the on-premises cluster and the cloud cluster, as shown in Example 10.

Example 10 Restart Platform LSF services

```

[root@homecluster conf]# lssadmin limrestart all

Checking configuration files ...
No errors found.

Do you really want to restart LIMs on all hosts? [y/n] y
Restart LIM on <homecluster> ..... done

```

```

[root@homecluster conf]# badmin mbdrestart

Checking configuration files ...

No errors found.

MBD restart initiated

[root@softlayer ~]# lsadmin limrestart all

Checking configuration files ...

No errors found.

Do you really want to restart LIMs on all hosts? [y/n] y
Restart LIM on <softlayer> ..... done
[root@softlayer ~]# badmin mbdrestart

Checking configuration files ...

No errors found.

MBD restart initiated
[root@softlayer ~]#

```

To check whether the multicluster feature is correctly configured and these clusters are enabled to access each other, run **lsclusters** and **bclusters** at the prompt to get OK status responses from both clusters, as shown in Example 11.

Example 11 Check the configuration

```

[root@homecluster ~]# lsclusters
CLUSTER_NAME  STATUS  MASTER_HOST  ADMIN  HOSTS  SERVERS
phpc_cluster  ok      homecluster  phpcadmin  1      1
HPC_Services  ok      softlayer    lsfadmin  1      1
[root@homecluster ~]# bclusters
[Job Forwarding Information ]
No local queue sending/receiving jobs from remote clusters

[Resource Lease Information ]
No resources have been exported or borrowed
[root@homecluster ~]#

```

Configuring job forwarding

This scenario shows how to change the high priority queue to send jobs to the IBM Platform Computing Cloud Services cluster. To do this task, change the `high_priority` stanza in the `lsb.queues` file at the local cluster (`phpc_cluster`). In this scenario, the master node `homecluster` path is the following one:

```
/install/shared/ibm/platform_lsf/conf/lsbatch/phpc_cluster/configdir/lsb.queues
```

This scenario does not preempt running jobs because the idea is to show how to send jobs to the cloud instead of interrupting a running job, so comment the PREEMPTION line and add SNDJOBS_T0 to point to the remote cluster. Then, change the description to state the usage of the queue. All changes are shown in bold in Example 12.

Example 12 File lsb.queues on the local cluster

```
Begin Queue
QUEUE_NAME    = high_priority
PRIORITY      = 43
NICE          = 10
SNDJOBS_T0   = receive@HPC_Services
#PREEMPTION   = PREEMPTIVE
#RUN_WINDOW
#CPULIMIT     = 8:0/SunIPC      # 8 hours of host model SunIPC
#FILELIMIT    = 20000
#DATA LIMIT   = 20000          # jobs data segment limit
#CORELIMIT    = 20000
#PROCLIMIT    = 5              # job processor limit
#USERS        = user1 user2 user3
#HOSTS        = all
#ADMINISTRATORS = user1 user3
#EXCLUSIVE    = N
#PRE_EXEC     = /usr/local/lsf/misc/testq_pre >> /tmp/pre.out
#POST_EXEC    = /usr/local/lsf/misc/testq_post |grep -v "Hey"
#REQUEUE_EXIT_VALUES = 55 255 78
DESCRIPTION  = Jobs submitted for this queue are scheduled as urgent\
jobs. Jobs in this queue can jobs in this queue can be forwarded to the Cloud
Services Cluster.
End Queue
```

In similar fashion, configure the receiving side to handle the jobs coming from the high priority queue. The lsb.queues file in the remote cluster, for this scenario, can be found in the SoftLayer host of the HPC_Services Cluster at the following path:

```
/usr/share/lsf/conf/lsbatch/HPC_Services/configdir/lsb.queues
```

In lsb.queues, add a stanza at the end of the file, as shown in Example 13.

Example 13 lsb.queues file on the remote cluster

```
Begin Queue
QUEUE_NAME=receive
RCVJOBS_FROM=high_priority@phpc_cluster
PRIORITY=70
NICE=20
End Queue
```

Now, reconfigure the queues on both sides, as shown in Example 14.

Example 14 Reconfigure the queues

```
[root@softlayer ~]# badmin mbdrestart
```

```
Checking configuration files ...
```

```
No errors found.
```

```
MBD restart initiated
[root@softlayer ~]#
```

```
[root@homecluster ~]# badmin mbdrestart
```

```
Checking configuration files ...
```

```
No errors found.
```

```
MBD restart initiated
[root@homecluster ~]#
```

Example 15 shows how to check the job forwarding status configuration for the local and remote queues.

Example 15 Check job forwarding

```
[root@homecluster ~]# bclusters
[Job Forwarding Information ]
LOCAL_QUEUE  JOB_FLOW  REMOTE    CLUSTER  STATUS
high_priority send      receive   HPC_Servic ok
```

```
[Resource Lease Information ]
No resources have been exported or borrowed
[root@homecluster ~]#
```

```
[root@softlayer ~]# bclusters
[Job Forwarding Information ]
LOCAL_QUEUE  JOB_FLOW  REMOTE    CLUSTER  STATUS
receive      recv      -         phpc_clust ok
```

```
[Resource Lease Information ]
No resources have been exported or borrowed
[root@softlayer ~]#
```

Testing your configuration

Now, test the new configuration by using the command-line interface (CLI) or the graphical user interface (GUI) from Platform HPC. If you have Platform LSF with Platform Application Center, you can use this interface as well.

Note: To submit the job to the cloud, the user must have authority to run jobs on the receiving queue.

This scenario uses the CLI to submit the jobs. Example 16 shows how to issue the **bsub** command to a dummy sleep job.

Example 16 Submit jobs to the respective queues

```
[root@homecluster ~]# bsub -q high_priority sleep 50
Job <857> is submitted to queue <high_priority>.
[root@homecluster ~]#
[root@homecluster ~]# bsub -q medium_priority sleep 50
Job <858> is submitted to queue <medium_priority>.
[root@homecluster ~]#
[root@homecluster ~]# bsub -q medium_priority sleep 50
Job <859> is submitted to queue <medium_priority>.
[root@homecluster ~]#
[root@homecluster ~]# bsub -q high_priority sleep 50
Job <860> is submitted to queue <high_priority>.
[root@homecluster ~]#
```

In this case, four jobs are submitted in a row, but the only queue that can forward jobs is the `high_priority` one. There are only two slots in the on-premises environment, so submit three jobs in the medium priority queue and the last job in the high priority queue. Only the job with the high priority runs in the cloud.

Example 17 shows the running jobs and the pending jobs in their respective queues.

Example 17 Jobs running in the cluster

```
[root@homecluster ~]# bjobs
JOBID  USER  STAT  QUEUE          FROM_HOST  EXEC_HOST  JOB_NAME    SUBMIT_TIME
857    root   RUN   high_priori    homecluster  homecluster *813545588 Mar 31 10:52
858    root   RUN   medium_prio    homecluster  homecluster *813554430 Mar 31 10:52
860    root   RUN   high_priori    homecluster  softlayer@H *813563747 Mar 31 10:53
859    root   PEND  medium_prio    homecluster  homecluster *813557373 Mar 31 10:52
```

As you can see, the first preference is to use available slots in the home cluster, and after there is no resource available, then only the high priority job goes to the cloud, even after being submitted after the last medium priority job.

Note: This is an example on how Platform LSF sends jobs from only a configured queue. Platform LSF is a powerful tool that helps you do advanced scheduling, and provides the best policies to suit your business needs.

Hybrid cloud is ready

The previous sections described the steps to configure a hybrid cloud in a few steps, although with the help of IBM Platform Computing Cloud Services, customers do not need to worry about configuring and managing a cloud infrastructure.

After following the five simple steps that were described in the previous sections, you can have extra capacity added to receive jobs from your existing environment. If you need assistance to configure a hybrid cloud environment, contact the IBM Platform Computing Services team for help and assistance.

Data management on hybrid clouds

Two easy ways to manage data across hybrid clouds are by implementing IBM Platform Data Manager for LSF and IBM Spectrum Scale AFM. Both technologies optimize data transfer needs to reduce costs and time to results because only the required data is moved at the correct time.

IBM Platform Data Manager for LSF

Platform Data Manager for LSF automates the transfer of data that is used by application workloads running on Platform LSF clusters and in the cloud. Frequently used data that is transferred between multiple data centers and the cloud can be stored in a smart, managed cache closer to compute resources. This smart data management helps to improve data throughput and minimizes wasted compute cycles, which helps you lower storage costs in the cloud.

With Platform Data Manager, the following actions occur:

- ▶ Data is staged in and out independently of workloads, freeing compute resources while data is transferred behind the scenes.
- ▶ A smart, managed cache reuses transferred data and avoids duplication of data transfers, sharing cached copies with all workloads that need access to the data, and among multiple users where appropriate.
- ▶ Data transfers are scheduled as jobs in Platform LSF and are subject to Platform LSF scheduling policies that are established by administrators, including priority.

For more information about IBM Platform Data Manager for LSF, see the following website:

<http://www.ibm.com/systems/platformcomputing/products/lsf/datamanager.html>

IBM Spectrum Scale Active File Management

AFM is a scalable, high-performance, file system caching layer that is integrated with Spectrum Scale. You can use AFM to create associations from a local cluster to a remote cluster or storage, and to define the location and flow of file data to automate the management of the data to implement a single namespace view across sites around the world.

AFM masks wide area network (WAN) latencies and outages by using Spectrum Scale to cache massive data sets, allowing data access and modifications even when a remote storage cluster is unavailable. In addition, AFM performs updates to the remote cluster asynchronously, which allows applications to continue operating while not being constrained by limited outgoing network bandwidth.

The AFM implementation uses the inherent scalability of Spectrum Scale to provide a multinode, consistent cache of data that is in a home cluster. By integrating it with the file system, AFM provides a Portable Operating System Interface (POSIX)-compliant interface, making the cache transparent to applications. AFM is easy to deploy, as it relies on open standards for high-performance file serving and does not require any proprietary hardware or software to be installed at the home cluster.

For a step-by-step configuration information, see the following website:

<http://ibm.co/1bPKBfY>

Authors

This paper was produced by a team of specialists from around the world working at the International Technical Support Organization, Poughkeepsie Center.

Dino Quintero is a Complex Solutions Project Leader and an IBM Senior Certified IT Specialist with the ITSO in Poughkeepsie, NY. His areas of knowledge include enterprise continuous availability, enterprise systems management, system virtualization, technical computing, and clustering solutions. He is an Open Group Distinguished IT Specialist. Dino holds a Master of Computing Information Systems degree and a Bachelor of Science degree in Computer Science from Marist College.

Daniel de Souza Casali is an IBM Cross Systems Senior Certified and has been working at IBM for 11 years. Daniel works for Systems and Technology Group in Latin America as a Software Defined Infrastructure IT Specialist. Daniel holds an Engineering degree in Physics from the Federal University of São Carlos (UFSCar). His areas of expertise include UNIX, SAN networks, IBM Disk Subsystems, clustering cloud, and analytics solutions.

Marcelo Correia Lima is a Business Intelligence Architect at IBM. He has 17 years of experience in development and integration of Enterprise Applications. His current area of expertise is Business Analytics Optimization (BAO) Solutions. He has been planning and managing Solutions lifecycle implementation, involving Multidimensional Modeling, IBM InfoSphere® Data Architect, IBM InfoSphere DataStage®, IBM Cognos® Business Intelligence, and IBM DB2®. In addition, Marcelo has added Hadoop, big data, IBM InfoSphere BigInsights™, cloud computing, and IBM Platform Computing to his background. Before working as a Business Intelligence Architect, he was involved in the design and implementation of IBM WebSphere® and Java Enterprise Edition Applications for IBM Data Preparation/Data Services.

Istvan Gabor Szabo is an Infrastructure Architect and Linux Subject Matter Expert at IBM Hungary (IBM DCCE SFV). He joined IBM in 2010 after receiving his bachelor degree in Engineering Information Technology from the University of Obudai - John von Neumann Faculty of Informatics. Most of the time, he works on projects as a Linux technical lead. His areas of expertise are configuring and troubleshooting complex environments, and building automation methodologies for server builds. In his role as an Infrastructure Architect, he works on the IBM Standard Software Installer (ISSI) environment, where he designs new environments based on customer requirements.

Maciej Olejniczak is a Cross-functional Software Support Team Leader in a collaborative environment. He works internationally with external and internal clients, IBM Business Partners, services, labs, and research teams. He is a dedicated account advocate for large customers in Poland. Maciej is an IBM Certified Expert in Actualizing IT Solutions: Software Enablement. He achieved a master level in implementing all activities that transform information technology from a vision to an actual working solution. Maciej is an Open Group Master Certified IT Specialist.

Tiago Rodrigues de Mello is a Staff Software Engineer in Brazil with more than 10 years of experience. Tiago's areas of expertise include Linux system administration, software development, and cloud computing. He is an OpenStack developer and a Continuous Integration engineer at the IBM Linux Technology Center. Tiago holds a Bachelor in Computer Science degree from the Federal University of Sao Carlos, Brazil.

Nilton Carlos dos Santos is an IT Architect and a Certified IT Specialist and has been with IBM since 2007 and has 18 years experience in the IT industry. Before joining IBM, he worked in several different areas of technology, including Linux and UNIX administration, database management, development in many different languages, and network administration. Nilton Carlos also has deep expertise in messaging, automation, monitoring, and reporting system tools. He enjoys working with open source software.

Thanks to the following people for their contributions to this project:

David Bennin and Richard Conway
International Technical Support Organization, Poughkeepsie Center

Now you can become a published author, too!

Here's an opportunity to spotlight your skills, grow your career, and become a published author—all at the same time! Join an ITSO residency project and help write a book in your area of expertise, while honing your experience using leading-edge technologies. Your efforts will help to increase product acceptance and customer satisfaction, as you expand your network of technical contacts and relationships. Residencies run from two to six weeks in length, and you can participate either in person or as a remote resident working from your home base.

Find out more about the residency program, browse the residency index, and apply online at:

ibm.com/redbooks/residencies.html

Stay connected to IBM Redbooks

- ▶ Find us on Facebook:
<http://www.facebook.com/IBMRedbooks>
- ▶ Follow us on Twitter:
<http://twitter.com/ibmredbooks>
- ▶ Look for us on LinkedIn:
<http://www.linkedin.com/groups?home=&gid=2130806>
- ▶ Explore new IBM Redbooks® publications, residencies, and workshops with the IBM Redbooks weekly newsletter:
<https://www.redbooks.ibm.com/Redbooks.nsf/subscribe?OpenForm>
- ▶ Stay current on recent Redbooks publications with RSS Feeds:
<http://www.redbooks.ibm.com/rss.html>

Notices

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

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

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

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

The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law: INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some states do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

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

Any references in this information to non-IBM websites are provided for convenience only and do not in any manner serve as an endorsement of those websites. The materials at those websites are not part of the materials for this IBM product and use of those websites is at your own risk.

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

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

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

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

COPYRIGHT LICENSE:

This information contains sample application programs in source language, which illustrate programming techniques on various operating platforms. You may copy, modify, and distribute these sample programs in any form without payment to IBM, for the purposes of developing, using, marketing or distributing application programs conforming to the application programming interface for the operating platform for which the sample programs are written. These examples have not been thoroughly tested under all conditions. IBM, therefore, cannot guarantee or imply reliability, serviceability, or function of these programs.

© Copyright International Business Machines Corporation 2015. All rights reserved.

Note to U.S. Government Users Restricted Rights -- Use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM Corp.

This document REDP-5214-00 was created or updated on June 17, 2015.



Send us your comments in one of the following ways:

- ▶ Use the online **Contact us** review Redbooks form found at: ibm.com/redbooks
- ▶ Send your comments in an email to: redbooks@us.ibm.com
- ▶ Mail your comments to:
IBM Corporation, International Technical Support Organization
Dept. HYTD Mail Station P099
2455 South Road
Poughkeepsie, NY 12601-5400 U.S.A.



Trademarks

IBM, the IBM logo, and ibm.com are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries, or both. These and other IBM trademarked terms are marked on their first occurrence in this information with the appropriate symbol (® or ™), indicating US registered or common law trademarks owned by IBM at the time this information was published. Such trademarks may also be registered or common law trademarks in other countries. A current list of IBM trademarks is available on the Web at <http://www.ibm.com/legal/copytrade.shtml>

The following terms are trademarks of the International Business Machines Corporation in the United States, other countries, or both:

BigInsights™
Cognos®
DataStage®
DB2®
GPFS™

IBM®
IBM Spectrum™
InfoSphere®
LSF®
Redbooks®

Redpaper™
Redbooks (logo) ®
Symphony®
WebSphere®

The following terms are trademarks of other companies:

SoftLayer, and SoftLayer device are trademarks or registered trademarks of SoftLayer, Inc., an IBM Company.

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

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

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

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