



Introduction to IBM Technical Computing Clouds

IBM Redbooks Solution Guide

The IT Industry has tried to maintain a balance between more demands from business to deliver services against cost considerations of hardware and software assets. Business growth depends on information technology (IT) being able to provide accurate, timely, and reliable services, but there is a cost that is associated with running IT services. The main challenge of IT is to provide timely, reliable, and accurate services for business needs, which have led to the growth and development of high performance computing (HPC). HPC traditionally has been the domain of powerful computers (called “supercomputers”) that are owned by governments and large multinationals. Existing hardware was used to process data and provide meaningful information by single systems working with multiple parallel processing units. Limitations were based on hardware and software processing capabilities. Because of the cost that was associated with such intensive hardware, the usage was limited to a few nations and corporate entities.

The advent of workflow-based processing model and virtualization and the high availability concepts of clustering and parallel processing have enabled existing hardware to provide the performance of the traditional supercomputers. New technologies, such as graphics processing units (GPUs), has pushed existing hardware to perform more complicated functions faster than previously possible on the same hardware, while virtualization and clustering have made it possible to provide greater level of complexity and availability of IT services. Sharing of resources to reduce cost also is possible because virtualization moved from a traditionally static IT model that is based on maximum load sizing to a leaner IT model that is based on workflow-based resource allocation through smart clusters. With the introduction of cloud technology, the resource requirement is becoming more on demand as compared to the traditional forecasted demand, thus optimizing cost considerations.

These technological innovations have made it possible to push the performance limits of existing IT resources to provide high performance output. The technical power to achieve computing results can be achieved with much cheaper hardware by using smart clusters and grids of shared hardware. The workflow-based resource allocation has made it possible to achieve high performance from a set of relatively inexpensive hardware working as a cluster. Performance can be enhanced by breaking across silos of IT resources, lying dormant to provide on-demand computing power wherever required. Data-intensive industries, such as engineering and life sciences, can now use the computing power on demand that is provided by the workflow-based technology. Using parallel processing by heterogeneous resources working as one unit under smart clusters, complex unstructured data can be processed to feed usable information into the system.

With the reduction in the cost of hardware resources, the demand for HPC has moved technical computing from scientific labs to mainstream commercial applications, as shown in Figure 1.

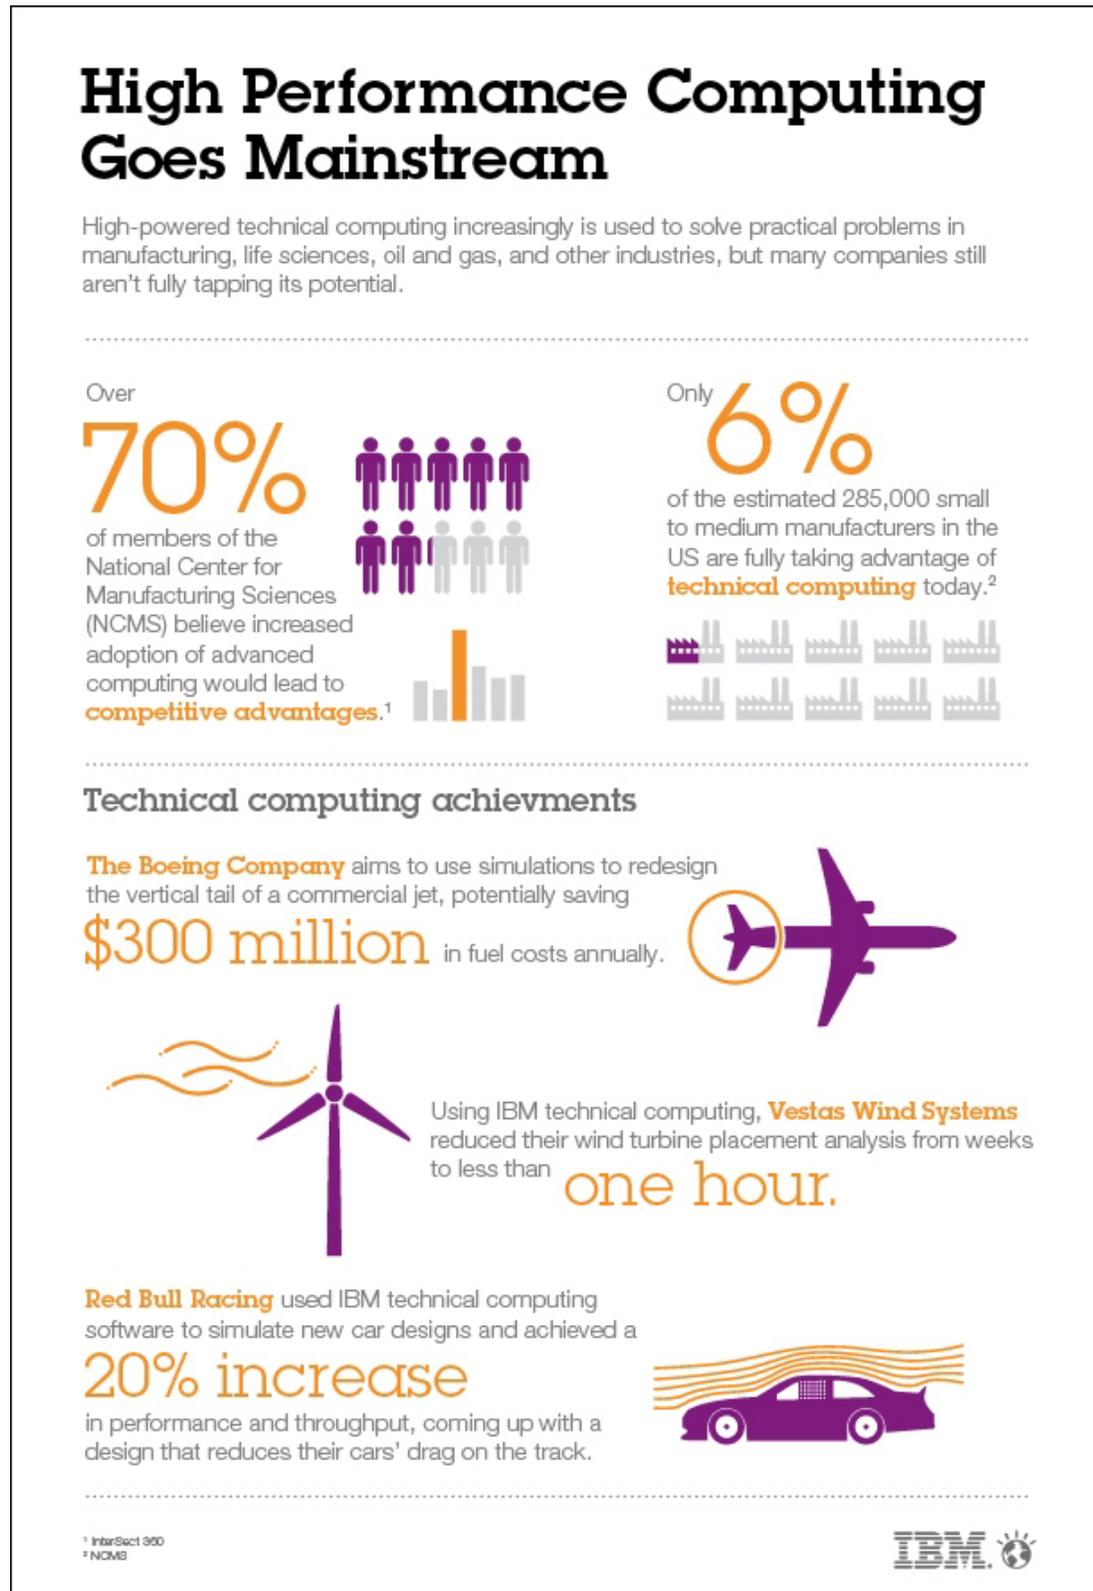


Figure 1. Technical computing goes mainstream

Technical computing has seen demand from sectors such as aerodynamics, automobile design, engineering, financial services, and oil and gas Industries. Improvement of cooling technology and power management of these superfast computing grids have allowed more efficiency and performance from existing hardware. Increased complexity of applications and demand for faster analysis of data has led technical computing to become mainstream or widely available. Thus, IBM® Technical Computing is focused on helping clients to transform their IT infrastructure to accelerate results. The goal of having technical computing in mainstream industries is to meet the challenges of applications that require high performance computing, faster access to data, and intelligent workload management, and make them easy to use and access from different applications with high performance outputs per business demands.

Did you know?

A cluster is typically an application or set of applications whose primary aim is to provide improved performance and availability at a lower cost compared to a single computing system with similar capabilities. A grid is typically a distributed system of homogeneous or heterogeneous computer resources for general parallel processing of related workflow, which is usually scheduled by using advanced management policies. A computing cloud is a system (private or public) that allows on-demand self-service, such as resource creation on demand, dynamic sharing of resources, and elasticity of resource sizing that is based on advanced workflow models.

IBM Platform Computing solutions have evolved from cluster to grid to cloud because of their abilities to manage heterogeneous complexities of distributed computing resources. IBM Platform Computing provides solutions for mission-critical applications that require complex workload management across heterogeneous environment for diverse industries, from life sciences to engineering to financial sectors involving complex risk analysis. IBM Platform Computing has a 20-year history of working on highly complex solutions for some of the largest multinational companies, where there are proven examples of robust management of highly complex workflow across large distributed environments to deliver results.

Business value

IBM HPC clouds can help enable the transformation of both the client's IT infrastructure and its business. Based on an HPC cloud's potential impact, clients are actively changing their infrastructure toward private clouds, and beginning to consider public and hybrid clouds. Clients are transforming their existing infrastructure to HPC clouds to enhance the responsiveness, flexibility, and cost effectiveness of their environment. This transformation helps a client enable an integrated approach to improve computing resource capacity and to preserve capital.

In a public cloud environment, HPC still must overcome a number of significant challenges, as shown in Table 1.

Table 1. Challenges of HPC in a public cloud

Challenges in a public cloud	Details
Security	<ul style="list-style-type: none"> • Cloud providers do not provide guarantees for data protection. • IP is in-flight outside the firewall and on storage devices.
Application licenses	<ul style="list-style-type: none"> • Legal agreements (LTUs) can limit licenses to geographies or corporate sites. • Unlimited licenses can be more expensive.
Business advantage	<ul style="list-style-type: none"> • Cloud resources are expensive compared to local resources if they are used incorrectly. • Building and automating a business policy for using cloud can be difficult.
Performance	<ul style="list-style-type: none"> • If applications run poorly in a private cloud, the applications do not improve in public clouds.
Data movement	<ul style="list-style-type: none"> • Data must be replicated in the cloud before jobs can run. • Providers charge for data I/O and storage.

If you use private clouds, HPC might not suffer from the challenges of public clouds, but can have a common set of existing issues, as shown in Table 2.

Table 2. Issues that a private cloud can have in HPC

Issues	Details
Inefficiency	<ul style="list-style-type: none"> • Less than fully utilized hardware • High labor cost to install, monitor, and manage HPC environments • Constrained space, power, and cooling
Lack of flexibility	<ul style="list-style-type: none"> • Resource silos tied to a specific project, department, or location • Dependency on specific individuals to run technical tasks
Delayed time to value	<ul style="list-style-type: none"> • Long provisioning times • Limited ability to fulfill peak demand • Constrained access to special purposes devices (for example, GPUs)

Solution overview

There are three different cloud-computing architectures:

- Private clouds
- Public clouds
- Hybrid clouds

A private cloud is an architecture where the client encapsulates its IT capacities “as a service” over an intranet for their exclusive use. The cloud is owned by the client, and managed and hosted by the client or a third party. The client defines the ways to access the cloud. The advantage is that the client controls the cloud so that security and privacy can be ensured. The client can customize the cloud infrastructure based on its business needs. A private cloud can be cost effective for a company that owns many computing resources.

A public cloud provides standardized services for public use over the Internet. Usually, it is built on standard and open technologies, providing web page, API, or SDK for the consumers to use the services, which include standardization, capital preservation, flexibility, and time to deploy.

The clients may integrate a private cloud and a public cloud to deliver computing services. This architecture is called hybrid cloud computing. Figure 2 highlights the differences and relationships of these three types of clouds.

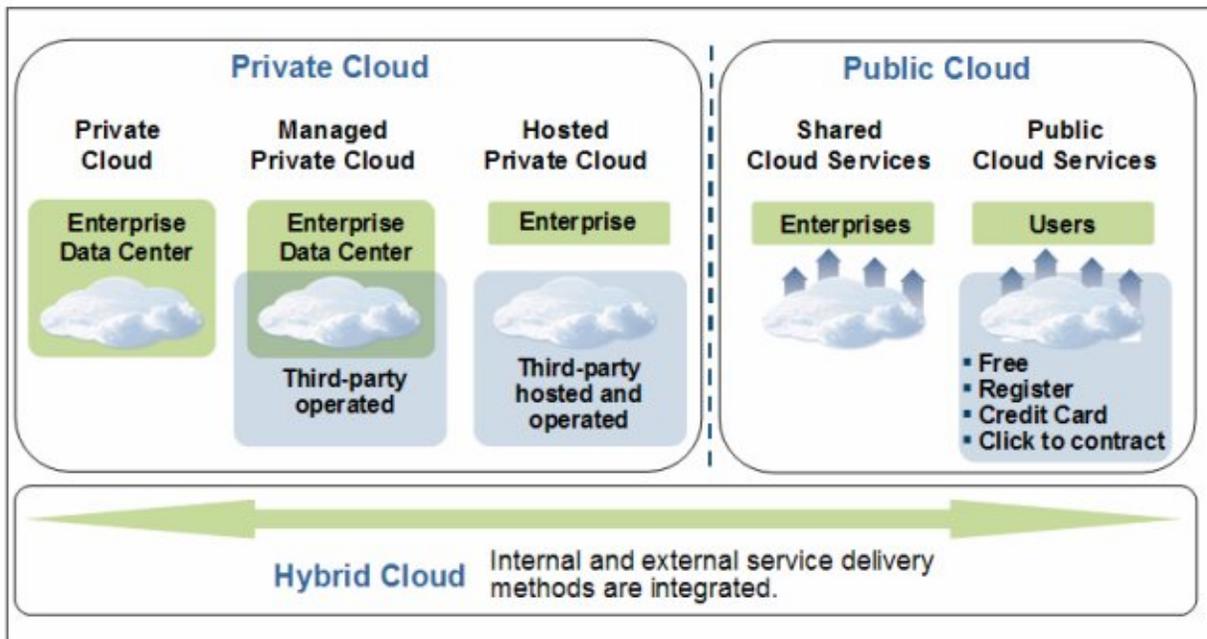


Figure 2. Types of clouds

Figure 3 shows the IBM HPC cloud reference model.

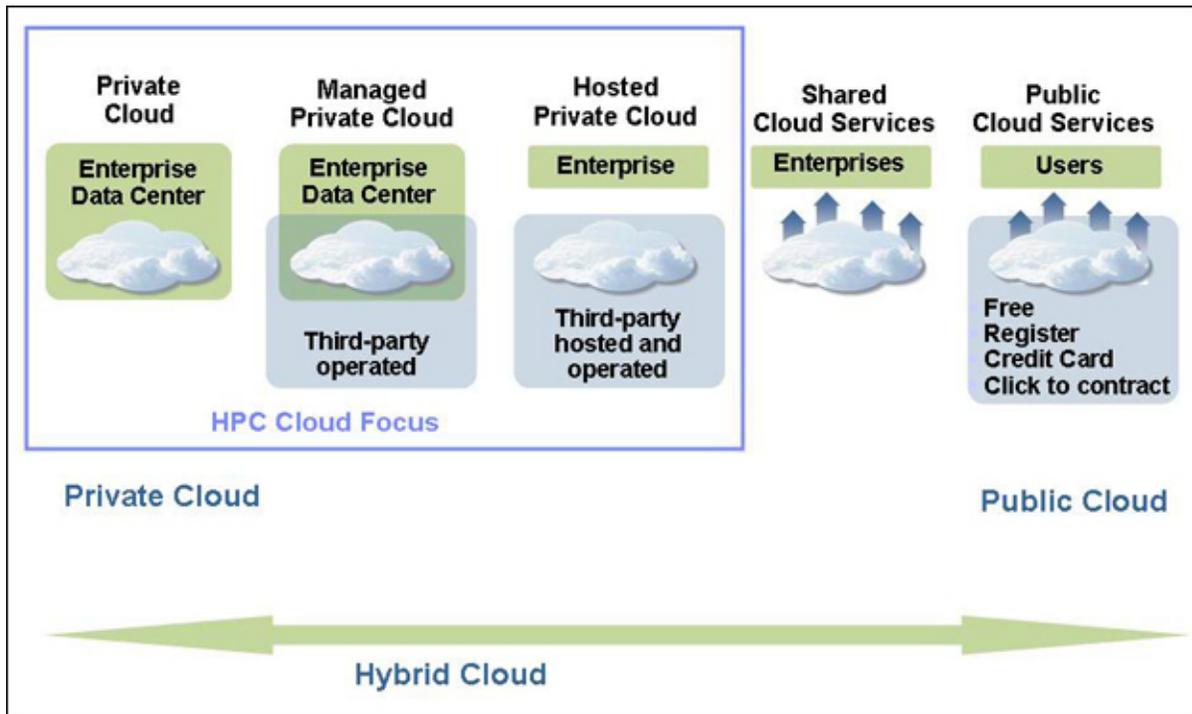


Figure 3. IBM HPC cloud

The HPC private cloud has three hosting models: private cloud, managed private cloud, and hosted private cloud. Table 3 describes the characteristics of these models.

Table 3. Private cloud models

Private cloud model	Characteristics
Private cloud	Client self-hosted and managed
Managed private cloud	Client self-hosted, but third-party managed
Hosted private cloud	Hosted and managed by a third party

Solution architecture

IBM HPC and IBM Technical Computing (TC) is about being flexible with the type of hardware and software that is available to implement the solution, such as the following hardware and software:

- IBM System x®
- IBM Power Systems™
- IBM General Parallel File Systems (GPFS™)
- Virtual infrastructure OpenStack

In addition to the above list, IBM Platform Computing provides support to heterogeneous cluster environments with additional IBM software or third-party software (see Figure 4):

- IBM Platform LSF®
- IBM Platform Symphony®
- IBM Platform Computing Management Advanced Edition (PCMAE)

- IBM InfoSphere® BigInsights™
- IBM General Parallel File System (GPFS)
- Bare Metal Provisioning through xCAT
- Solaris Grid Engine
- Open Source Apache Hadoop
- Third-party schedulers

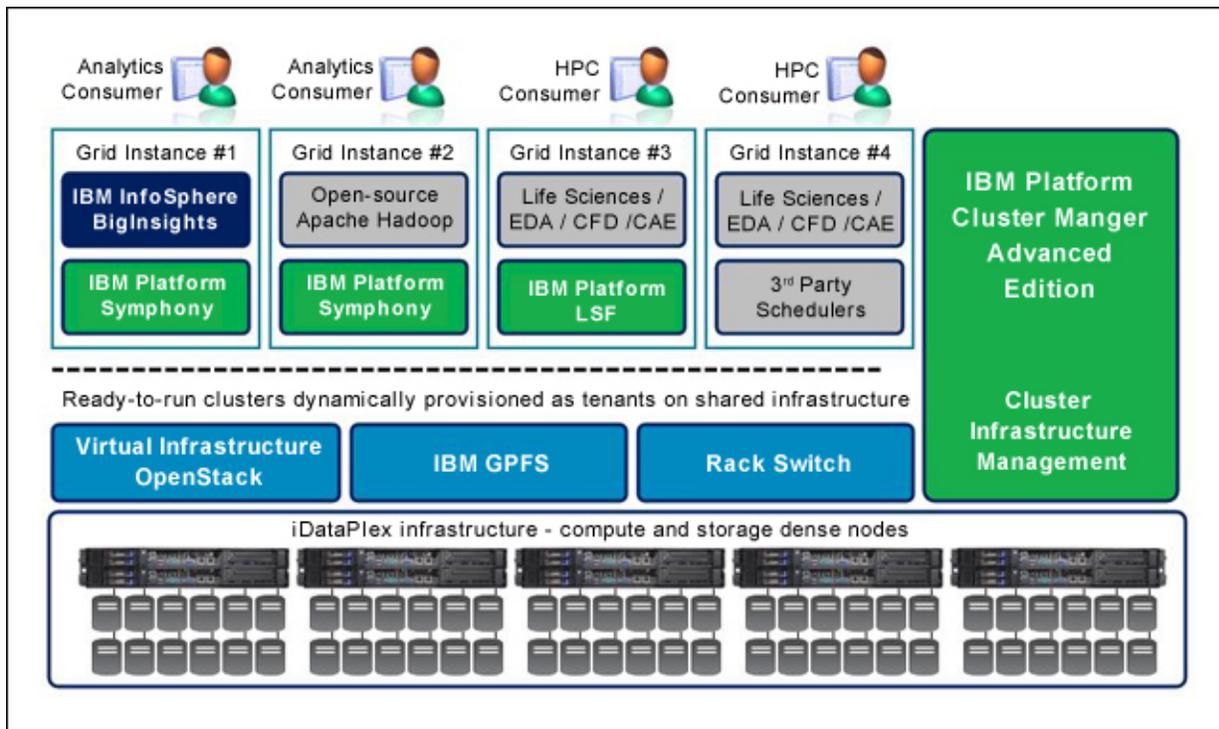


Figure 4. Overview of IBM Technical Computing and analytics clouds solution reference architecture

IBM Cluster Manager tools use the bandwidth of the network devices to lower the latency levels. Here are some of the supported devices:

- IBM RackSwitch™ G8000, G8052, G8124, and G8264
- Mellanox InfiniBand Switch System IS5030, SX6036, and SX6512
- Cisco Catalyst 2960 and 3750 switches

IBM Cluster Manager tools use storage devices that are capable of high parallel I/O to help provide efficient I/O related operations in the cloud environment. Here are some of the storage devices that are used:

- IBM DCS3700
- IBM System x GPFS Storage Server

Usage scenarios

Cloud computing provides the flexibility to use resources when they are required. In terms of a technical computing cloud environment, cloud computing provides not only the flexibility to use the resources on demand, but helps to provision the computing nodes, as per the application requirement, to help manage the workload. By implementing and using IBM Platform Computing Manager (PCM), dynamic provisioning of the computing nodes with operating systems is easily achieved. This dynamic provisioning solution helps you to better use the hardware resources and fulfill various technical computing requirements for managing the workloads. Figure 5 shows the infrastructure of an HPC cloud.

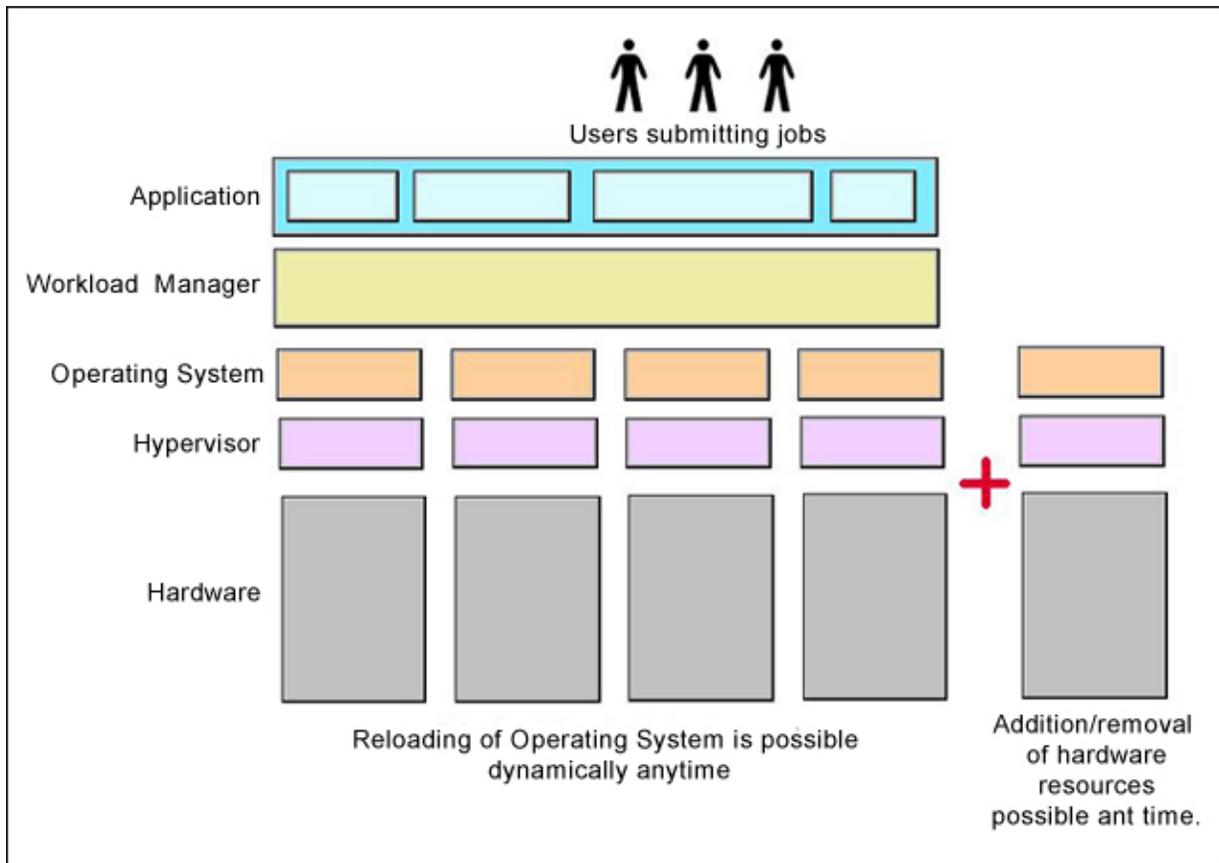


Figure 5. Flexible infrastructure with cloud

Cloud computing can reduce the manual effort of installation, provisioning, configuration, and other tasks that might have been performed manually before. When these computing resource management steps are done manually, they can take a significant amount of time. A cloud-computing environment can reduce system management complexity by implementing automation, business workflows, and resource abstractions.

IBM Platform Cluster Manager - Advanced Edition (PCM-AE) provides many automation features to help reduce the complexity of managing a cloud-computing environment. Using PCM-AE, you can perform the following actions:

- Rapidly deploy multiple HPC heterogeneous clusters in a shared hardware pool.
- Perform self-service, which allows you to request a custom cluster and specify the size, type, and time frame.

- Dynamically grow and shrink (flex up and down) the size of a deployed cluster based on workload demand, calendar, and sharing policies.
- Share hardware across clusters by rapidly reprovisioning the resources to meet the infrastructure's needs (for example, Windows and Linux, or a different version of Linux).

These automation features reduce the time that is required to make the resources available to the clients quickly.

In cloud computing, there are many computers, network devices, storage devices, and applications that are running. To achieve high availability, throughput, and resource utilization, cloud-computing infrastructures use monitoring mechanisms, which are also used to measure the service and resource usage, which is key for chargeback to users. The system statistics are collected and reported to the cloud provider or user, and based on these figures, dashboards can be generated.

Monitoring brings the following benefits:

- Avoid outages by checking the health of the cloud-computing environment
- Improve resource usage to help lower costs
- Identify performance bottlenecks and optimizes workloads
- Predict usage trend

IBM SmartCloud® Monitoring V7.1 is a bundle of established IBM Tivoli® infrastructure management products, including IBM Tivoli Monitoring and IBM Tivoli Monitoring for Virtual Environments. The software delivers dynamic usage trending and health alerts for pooled hardware resources in the cloud infrastructure. The software includes sophisticated analytics, and capacity reporting and planning tools. You use these tools to ensure that the cloud is handling workloads quickly and efficiently.

For more information about IBM SmartCloud Monitoring, see the following website:

<http://www-01.ibm.com/software/tivoli/products/smartcloud-monitoring/>

Integration

Technical computing workloads have the following characteristics:

- Large number of machines
- Heavy resource usage, including I/O
- Long running workloads
- Dependent on parallel storage
- Dependent on attached storage
- High bandwidth, low latency networks
- Compute intensive
- Data intensive

HPC clusters frequently employ a distributed memory model to divide a computational problem into elements that can be simultaneously run in parallel on the hosts of a cluster. This often involves the requirement that the hosts share progress information and partial results by using the cluster's interconnect fabric. This is most commonly accomplished through the usage of a message passing mechanism. The most widely adopted standard for this type of message passing is the message passing interface (MPI) standard, which is described at the following website:

<http://www.mpi-forum.org>

The Technical Computing cloud solutions also include integrated IBM Platform Computing software that addresses technical computing challenges:

- IBM Platform HPC is a complete technical computing management solution in a single product, with a range of features that improve time-to-results and help researchers focus on their work rather than on managing workloads.
- IBM Platform LSF provides a comprehensive set of tools for intelligently scheduling workloads and dynamically allocating resources to help ensure optimal job throughput.
- IBM Platform Symphony delivers powerful enterprise-class management for running big data, analytics, and compute-intensive applications.
- IBM Platform Cluster Manager-Standard Edition provides easy-to-use yet powerful cluster management for technical computing clusters that simplifies the entire process, from initial deployment through provisioning and ongoing maintenance.
- IBM General Purpose File System (GPFS) is a high-performance enterprise file management platform for optimizing data management.

IBM Platform MPI is a high-performance and production-quality implementation of the MPI standard. It fully complies with the MPI-2.2 standard and provides enhancements, such as low latency and high-bandwidth, point-to-point, and collective communication routines over other implementations.

For more information about IBM Platform MPI, see the IBM Platform MPI User's Guide, SC27-4758-00:

<http://www-01.ibm.com/support/docview.wss?uid=pub1sc27475800>

SOA is a software architecture in which the business logic is encapsulated and defined as services. These services can be used and reused by one or multiple systems that participate in the architecture. SOA implementations are generally platform-independent, meaning that infrastructure considerations do not hinder the deployment of new systems or the enhancement of existing systems. Many financial institutions deploy a range of technologies, so the heterogeneous nature of SOA is important.

IBM Platform Symphony combines a fast service-oriented application middleware (SOAM) component, with a highly scalable grid management infrastructure. Its design delivers reliability and flexibility, while also ensuring low levels of latency and high throughput between all system components.

For more information about SOA, see the following website:

<http://www-01.ibm.com/software/solutions/soa/>

MapReduce is a programming model for applications that process large volumes of data in parallel by dividing the work into a set of independent tasks across a large number of machines. MapReduce programs in general transform lists of input data elements into lists of output data elements in two phases: map and reduce. MapReduce is used in data-intensive computing such as business analytics and life science. Within IBM Platform Symphony, the MapReduce framework supports data-intensive workload management by using a special implementation of SOAM to manage MapReduce workloads.

Workflow is a task that is composed of a sequence of connected steps. In HPC clusters, many workflows are running in parallel to finish a job or to respond to a batch of requests. As the complexity increases, workflows are more complicated. Workflow automation is becoming increasingly important because of the following reasons:

- Jobs must run at the correct time and in the correct order.
- Mission-critical processes have no tolerance for failure.
- There are inter-dependencies between steps across machines.

Clients need an easy-to-use and cost-efficient way to develop and maintain the workflows.

Visualization is a typical workload in engineering for airplanes and automobiles designers. The designers create large computer aided design (CAD) environments to run their 2D/3D graphic calculations and simulations for the products. These workloads demand a large hardware environment that includes graphic workstations, storage, software tools, and so on. In addition to the hardware, the software licenses are also expensive. Thus, the designers are looking to reduce costs, and expect to share the infrastructure between computing aided engineering (CAE) and CAD.

Supported platforms

The reference architecture for the Technical Computing cloud solution is based on IBM Flex System™, IBM System x, and IBM Platform Computing.

IBM has created reference architectures for target workloads and applications. Each of these reference architectures includes recommended small, medium, and large configurations that ensure optimal performance at entry-level prices. These reference architectures are based on powerful, predefined, and tested infrastructure with a choice of the following systems:

- IBM Flex System combines leading-edge IBM POWER7®, IBM POWER7+™ and x86 compute nodes with integrated storage and networking in a highly dense, scalable blade system. The IBM Application Ready solution supports IBM Flex System x240 (x86), IBM Flex System p260, and p460 (IBM Power Systems) compute nodes.
- IBM System x helps organizations address their most challenging and complex problems. The Application Ready Solution supports a rack-mounted, cost-optimized System x 3550 M4 to System x IBM iDataPlex® dx360 M4 systems to optimize density, performance, and graphics acceleration for remote 3D visualization.
- IBM System Storage® DS3524 is an entry-level disk system that delivers an ideal price/performance ratio and scalability. You also can choose the optional IBM Storwize® V7000 Unified for enterprise-class, midrange storage to consolidate block-and-file workloads into a single system.
- IBM Intelligent Cluster™ is a factory-integrated, fully tested solution™ that helps simplify and expedite deployment of IBM Flex System x86 or System x -based Application Ready Solutions.

For more information about these supported reference architectures, see the following websites:

- IBM System x:
<http://www.ibm.com/systems/x/?lnk=mprSY-sysx-usen>
- IBM PureSystems:
<http://www.ibm.com/ibm/puresystems/us/en/?ad=messagedetect-6>
- IBM Platform Computing:
ibm.com/platformcomputing

Ordering information

Table 4 shows a list of IBM Platform Computing solutions for HPC Cloud, workload management, big data analytics, and cluster management, and the available products.

Table 4. IBM Platform Computing Solutions

Product family	Offering name - chargeable component
IBM Platform LSF V9.1	IBM Platform LSF - Express Edition IBM Platform LSF - Standard Edition (includes Power support) IBM Platform LSF - Express to Standard Edition Upgrade IBM Platform Process Manager IBM Platform License Scheduler IBM Platform RTM IBM Platform Application Center IBM Platform MPI IBM Platform Dynamic Cluster IBM Platform Sessions Scheduler
	IBM Platform Analytics - Express Edition IBM Platform Analytics - Express to Standard Upgrade IBM Platform Analytics - Standard Edition IBM Platform Analytics - Standard to Advanced Upgrade IBM Platform Analytics - Advanced Edition IBM Platform Analytics Data Collectors
	IBM Platform LSF - Advanced Edition
IBM Platform Symphony V6.1	IBM Platform Symphony - Express Edition IBM Platform Symphony - Standard Edition IBM Platform Symphony - Advanced Edition IBM Platform Symphony - Developer Edition IBM Platform Symphony - Desktop Harvesting IBM Platform Symphony - GPU Harvesting IBM Platform Symphony - Server and VM Harvesting IBM Platform Analytics IBM Platform Symphony - Express to Standard Upgrade IBM Platform Symphony - Standard to Advanced Upgrade
IBM Platform HPC V3.2	IBM Platform HPC V3.2 - Express Edition for System x
	IBM Platform HPC - x86 Nodes (other equipment manufacturers (OEM) only)
IBM Platform Cluster Manager V4.1	IBM Platform Cluster Manager - Standard Edition IBM Platform Cluster Manager - Advanced Edition

Ordering information, including the program number, version, and program name, are shown in Table 5.

Table 5. Ordering program number, version, and program name

Program number	VRM	Program name
5725-G82	9.1.1	IBM Platform LSF
5725-G82	9.1.0	IBM Platform Process Manager
5725-G82	9.1.0	IBM Platform RTM
5725-G82	9.1.0	IBM Platform License Scheduler
5725-G82	9.1.0	IBM Platform Application Center
5725-G88	4.1.0	IBM Platform Cluster Manager
5641-CM6	4.1.0	IBM Platform Cluster Manager, V4.1 with one-year S&S
5641-CM7	4.1.0	IBM Platform Cluster Manager, V4.1 with three-year S&S
5641-CM8	4.1.0	IBM Platform Cluster Manager, V4.1 with five-year S&S
5641-CMG	4.1.0	IBM Platform Cluster Manager, V4.1 Term License with one-year S&S
5725-G86	6.1.0	IBM Platform Symphony
5725-G84	8.3.0	IBM Platform Analytics
5725-K71	3.2.0	IBM Platform HPC - Express Edition

Related information

For more information, see the following documents:

- *IBM Technical Computing Clouds*, SG24-8144
<http://www.redbooks.ibm.com/abstracts/sg248144.html>
- *IBM Platform Computing Solutions*, SG24-8073
<http://www.redbooks.ibm.com/sg248073.html>
- *IBM Platform Computing Integration Solutions*, SG24-8081
<http://www.redbooks.ibm.com/sg248081.html>
- IBM Platform Computing page
<http://ibm.com/systems/technicalcomputing/platformcomputing/index.html>
- IBM Technical Computing page
<http://www-03.ibm.com/systems/technicalcomputing/>
- High performance computing cloud offerings from IBM
<http://public.dhe.ibm.com/common/ssi/ecm/en/dcs03006usen/DCS03006USEN.PDF>
- IBM Engineering Solutions for Cloud: Aerospace and Defense, and Automotive
<http://public.dhe.ibm.com/common/ssi/ecm/en/dcs03009usen/DCS03009USEN.PDF>
- IBM Application Ready Solutions for Technical Computing
<http://www-03.ibm.com/support/techdocs/atmastr.nsf/WebIndex/WP102288>

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 give 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 Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this IBM product and use of those Web sites 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. 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.

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 measurement may have been estimated through extrapolation. Actual results may vary. Users of this document should verify the applicable data for their specific environment.

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 2013. 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 was created or updated on August 30, 2013.

Send us your comments in one of the following ways:

- Use the online **Contact us** review form found at:
ibm.com/redbooks
- Send your comments in an e-mail to:
redbook@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.

This document is available online at <http://www.ibm.com/redbooks/abstracts/tips1023.html> .

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™
GPFS™
IBM®
IBM Flex System™
IBM SmartCloud®
iDataPlex®
InfoSphere®
Intelligent Cluster™
LSF®
POWER®
Power Systems™
POWER7®
POWER7+™
RackSwitch™
Redbooks®
Redbooks (logo)®
Storwize®
Symphony®
System Storage®
System x®
Tivoli®

The following terms are trademarks of other companies:

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

Windows, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.

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