

White Paper February 2012

By Robert Miggins, Senior VP of Business Development and

Rich Rivera, HPC Specialist

www.peer1.com

Outsourced High Performance Computing

On-demand Supercomputing Resources
That Fit Budget and Technical Requirements

Table of Contents

Introduction: In-House HPC Cannot Keep Pace	1
The Challenge: Overcoming the Slow Pace and High Costs of HPC Project Initiation	2
Time Delays in Starting HPC Projects	2
High Costs of In-House HPC	2
Market Leaders Turn to Outsourced HPC	3
Instant Infrastructure on Demand	4
Lowering Capital Expenditure While Building to Scale	4
Secure and Reliable to the Most Stringent HPC Standards	5
Comparing Costs: Outsourced HPCC Versus In-House HPC Infrastructure	5
Managed HPC Services are Not All the Same	6
Case Studies	8
Occipital: Delivering Innovative Views to Mobile Users	8
Case Study: Media Software and Services Provider Speeds and Secures Hollywood Studios' IP	9
Conclusion	10

Introduction: In-House HPC Cannot Keep Pace

Today, many companies rely on High Performance Computing, or HPC, as a major component of competitive advantage. As such, CIOs across multiple industries continuously seek to add HPC resources quickly and easily in response to new business demands.

Despite HPC's importance, CIOs still find themselves locked in budget battles and arguments over IT and business priorities when trying to support evolving HPC needs with additional compute resources. To overcome this challenge and keep their business as agile as possible, leading companies are streamlining their HPC growth through strategic outsourcing.

Organizations once considered HPC as hallowed ground which they would never outsource, particularly because of intellectual property and security concerns. At the same time, a market disconnect existed wherein service providers failed to address CIOs' needs with solutions that provided fast ramp-up as well as security and service flexibility.

Fortunately, the situation has changed. New HPC private cloud solutions, or HPCC, provide the benefits of a traditional outsourced service with the IP protection, security, flexibility, and speed that CIOs require for their HPC initiatives. As a result, CIOs now outsource HPC just as they outsource other core portions of their business. They view HPCC in the same manner as outsourced Microsoft Exchange email, CRM applications, and product development.

The move to HPCC has opened the floodgates to innovation. Instead of long project delays caused by budget wrangling, leading companies are seizing on new market opportunities with superior agility and setting themselves apart from the competition. It's not just large companies. Midsized and small organizations are also ramping up their capabilities to respond to unforeseen demands and bring the highest quality products and services to markets in the shortest timeframe possible, all while reducing CapEx.

How Market Leaders Leverage Outsourced HPC Infrastructure

- Oil & Gas: Analysis of geophysics maps, producing a final image map to determine drilling locations.
 - Digital Content Creation: Processing data sets to produce animations/CGI, architectural drawings, or 3-D applications.
 - Banking & Finance: Derivatives analysis, analytics of large data sets (i.e., numbers crunching).
 - Government & Scientific Research: Quantitative modeling.
-

This paper defines HPC outsourced infrastructure services, the reasons behind the growing trend of outsourcing HPC infrastructure, and the benefits organizations can realize from doing so. It also provides an overview of what to seek in an ideal solution for managed HPC hosting services, known as High Performance Cloud Computing (HPCC) services.

Note: While some refer to outsourced HPC as outsourced graphics processing units (GPUs), this paper uses the "HPC" acronym to stand for HPC, with the understanding that GPUs underpin HPC.

The Challenge: Overcoming the Slow Pace and High Costs of HPC Project Initiation

The current challenging economic environment and increasingly competitive landscape compel organizations to act fast when business opportunities arise. As such, agile companies with flexible infrastructure and computing capabilities realize significant competitive advantage in today's markets.

For companies that produce and analyze vast quantities of data, HPC has streamlined computing tasks, in many cases shrinking project deliverable times from weeks or months to just hours or days. But many companies have outgrown their datacenter space and resources allocated to HPC, and adding more is both slow and costly.

GPUs Form The Heart of HPC.

At the core of HPC infrastructure are GPUs, or graphics processing units. GPUs rapidly manipulate and alter memory to accelerate the building of images in a frame buffer intended for output to a display. GPUs are very efficient at manipulating computer graphics, and their highly parallel structure makes them more effective than general-purpose CPUs for algorithms where processing of large blocks of data is done in parallel.¹

Time Delays in Starting HPC Projects

Many organizations do not maintain significant excess HPC resources in their datacenter in the event that a new need for HPC arises. With IT budgets under extreme scrutiny today, management reviews every major IT expenditure thoroughly, often vetting decisions through multiple levels of the organization. Even if CIOs successfully achieve management buy-in for additional HPC resources, they often incur lengthy timelines for datacenter preparation, and infrastructure procurement, delivery, installation, and setup.

High Costs of In-House HPC

Running and maintaining HPC in the datacenter is very expensive, requiring IT departments to spend heavily in several areas: datacenter space, computing infrastructure, system administration, and energy.

Unlike traditional computing environments, HPC computing infrastructure quickly consumes energy and datacenter space, mostly due to GPUs' high rate of power consumption. This power consumption forces organizations to restrict the number of HPC servers to less than 10 per rack, whereas normal compute environments often accommodate 42 servers in a rack. This means an in-house HPC system requires far more racks and a similarly higher proportion of cooling equipment to prevent overheating in the datacenter. The combination of more racks and more cooling equipment quickly consumes floor space when accommodating HPC demands.

The server costs of HPC systems can escalate rapidly. HPC system prices normally range from \$10,000-\$60,000 for just one GPU-based server. The exact cost will depend upon several factors, including the amount of GPUs, the model of GPU, and CPU type as well as varying requirements for memory and disk subsystems. At the same time, IT incurs high development and system administrator costs caused by a disconnect between HPC hardware capabilities and HPC application needs.

IDC explains the problem this way in a recent study: "The underlying problem is that modern HPC hardware with large numbers of CPU cores, each with decreasing levels of memory and memory bandwidth, is causing a mismatch with existing application software, driving a need to fundamentally redesign and rewrite HPC application software for greater parallelism, in order to perform well on future HPC systems."²

The above statement from IDC points to two additional ongoing challenges of in-house HPC systems: upgrades and maintenance. Upgrades to HPC systems must occur on a regular basis. Newer, more advanced systems consistently outperform older ones, with GPUs increasing in performance while reducing power consumption. For example, the M2050 series NVIDIA GPU can consume up to 250 watts of power while its newer counterpart, the NVIDIA M2090 GPU, cuts the maximum wattage down to 225 and delivers faster processing times.

Lagging behind in computing power in fast-paced industries can drastically weaken a company's competitive position. This adds another cost burden to IT, as it must continuously dedicate more budget and personnel to maintaining and upgrading the HPC environment. As an organization adopts more HPC resources, IT must add personnel to provide end-user support as well as maintain and troubleshoot systems.

Regarding energy consumption, with HPC systems consuming far more energy in the datacenter, it's little wonder that HPC user conferences continue to focus on discussions surrounding best practices for reducing HPC energy usage.

Power-Hungry HPC Infrastructure.

The largest supercomputers can consume from 1 megawatt to 10Mw of energy. This alone adds up to \$20 million per year in OpEx.

The costs and time associated with adding in-house HPC resources puts CIOs and IT departments in a tough bind. They know they must be responsive to the engineers, product developers, scientists, and others who rely on HPC to perform the work that is the lifeblood of their company. But they also recognize that their budgets cannot sustain the levels of HPC requested.

IT leaders simply do not have the luxury of purchasing more datacenter space and energy to fulfill each new HPC request from the business. To make matters worse, there's always a competitor who has just purchased the next generation of HPC infrastructure, eliminating the competitive advantage that was so hard to achieve in the first place.

Today's market reality dictates that the most nimble companies are winning the day. This is true across industries, which becomes evident when one considers such high-profile advances as those in 3-D animation and imaging. Or the pace at which financial services firms analyze minute-by-minute changes in world markets, and the vast research findings that biotech companies cull from the human genome. As such, many IT leaders are seeking alternative ways to perform HPC without incurring the high costs and long ramp-up times associated with HPC projects.

If they can minimize the resource, cost, and time challenges presented by HPC infrastructure needs, organizations can spend more time focusing on their core competencies and winning business.

Market Leaders Turn to Outsourced HPC

Outsourced HPC is a service that helps organizations with an acute need for processing large data sets and producing results from them. Referred to as High Performance Cloud Computing, or HPCC, it provides access to a cloud of HPC services via a service provider. This new model of HPC hosting services is designed to offload high-CapEx datacenter infrastructure in exchange for a fast-turn, pay-per-use infrastructure that IT teams can utilize on demand.

Leading users of HPCC services include digital content creators, oil & gas companies, financial services firms, and government & research organizations. Managed HPCC services benefit these organizations in the same way that managed hosting services have for over a decade: they eliminate the cost, time, resource, and planning associated with installing and managing in-house technology infrastructure. This allows each organization to focus on its core competency. Rather than wasting valuable resources on datacenter infrastructure design and HPC system maintenance, they produce more end results faster – more 3-D images, more accurate geophysics maps, and more exciting digital content.

Instant Infrastructure on Demand

Rather than waiting weeks or months for project approval and to acquire the necessary hardware and software, an organization can utilize HPCC services to instantly initiate any size project within days. The reasonable cost and on-demand nature of HPCC, even at extremely high scale, means that IT teams no longer need to predict their HPC needs or battle to decide which projects should take priority. When requests come to IT from engineers, scientists, and others who are working on the next “hot” initiative, IT responds as an enabler rather than a roadblock.

This truly transforms IT from a roadblock into a business enabler. In the overview of its special three-part report, *The New Enterprise Architecture*, Gartner states, “Clearly, the design points of new architectures must be different from the old. Architectures must enable business transformation, rather than stand in the way.”³ The report goes on to explore what Gartner views as the four elements of a new foundation in enterprise architecture:

- The “grid”
- Architectural styles that represent key business processes
- Patterns, which are logical technology models
- Core technology building blocks, or “bricks,” that link to the grid⁴

HPCC aligns substantially to the concept that Gartner outlines in its report. Leveraging HPCC services, IT leaders can utilize the latest grid-style architecture to ensure fast results, regardless of HPC project size and timing.

Lowering Capital Expenditure While Building to Scale

HPCC offers a significant opportunity for CIOs and CFOs to shed CapEx from their slimmed-down budgets. As stated previously, HPC requires nearly seven times the racks as well as other heavy duty equipment. With HPCC, there is no need to amortize the costs of both expensive servers and the surrounding infrastructure.

While leasing of in-house HPC equipment once appeared to address the CapEx issue, organizations still find themselves strapped with monthly lease and energy payments during times when HPC resources remain underutilized. With HPCC’s pay-as-you-go model, organizations save money by opting out of HPC resource usage when not needed, owing nothing until they resume usage. This low CapEx model alleviates

much of the angst that traditionally builds up between IT and the business. IT can now green light projects that the business deems necessary, while business and innovation leaders can obtain high-scale resources without worrying about spinning down demand at a later date.

HPCC services quickly obviate the scaling issues that complicate in-house HPC. Leading HPCC service providers have grown out of the long-standing managed hosting service provider market. As a result, they have years of experience forecasting and allocating resources to suit market demand. Such service providers own massive amounts of HPC infrastructure and maintain close partnerships with leading HPC technology OEMs, such as nVidia, for GPUs. Even as projects arise quickly or unexpectedly, customers will always have the right amount and latest generation of HPC resources available to them.

Secure and Reliable to the Most Stringent HPC Standards

The security features of HPCC services cannot be discussed without stressing the fact that true cloud security can only be achieved through a private cloud. In her December 2011 blog post entitled *Beware misleading marketing of "private clouds,"* Gartner analyst Lydia Leong states the following: "The NIST definition of private cloud is clear: 'Private cloud. The cloud infrastructure is provisioned for exclusive use by a single organization comprising multiple consumers (e.g., business units). It may be owned, managed, and operated by the organization, a third party, or some combination of them, and it may exist on or off premises.' In other words, NIST defines private cloud as single-tenant."⁵

Gartner: Organizations Need to Understand Cloud Security More Thoroughly.

In Gartner's top predictions for 2012 and beyond, they state, "While enterprises are evaluating the potential cloud benefits in terms of management simplicity, economies of scale and workforce optimisation, it is equally critical that they carefully evaluate cloud services for their ability to resist security threats and attacks."⁶

The most robust managed HPCC services provide organizations with a private cloud environment, with no semblance of multi-tenancy. This means a company's valuable data and IP will never come along with other organizations' data.

Even many conservative CIOs, traditionally late adopters of IT outsourcing services, find themselves drawn to the complementary and secure aspects of private cloud-based HPCC services. Because these CIOs can start HPC outsourcing in a limited fashion, they can sample a test project, compare costs and results to internally managed HPC, and outsource as much as they feel comfortable with going forward.

The Outsourcing of Everything.

According to Deloitte, 86% of the Fortune 500 have employed a shared services model, and over 40% of the Fortune 500 have offshored at least one function/process.⁷

Comparing Costs: Outsourced HPCC Versus In-House HPC Infrastructure

Table 1 below demonstrates the typical cost savings that outsourced HPCC provides to an organization in each year of a three-year project period. The cost comparison data is based upon a HPC project scenario with the following criteria:

- 8 servers (16 GPUs) operate steady state for 3 full years.
- Each server incorporates 2 Intel® Xeon® Quad Core Processors (CPUs), 2 NVIDIA® "Fermi" GPUs, 24GB of memory, and 900GB of local storage.

- All 8 servers fit into a 42U server cabinet, consuming an equivalent amount of energy
- Each server costs \$10,000, and a 9th server is purchased as a spare. The \$90,000 total server cost is amortized over the three-year period, resulting in a yearly cost of \$30,000
- Data transfer charges include the additional \$800 per month for a dedicated 100MB feed
- Personnel costs include partial salary and benefits of system administrator staff to watch over the infrastructure on a 24x7x365 basis.

Table 1. Single-Year Cost Comparison: In-House HPC vs. Outsourced HPCC

Cost Item	In-House	HPCC Service Provider
Usage Costs		\$87,360
Server Hardware (1)	\$30,000	
Server Software (monitoring, patching)	\$9,600	
Hardware Maintenance, Spares	\$8,000	
Data Transfer (2)	\$9,600	
Remote Support	\$2,000	
Personnel (3)	\$34,600	
Datacenter Facility & Power	\$9,600	
Total Over 3 Years	\$103,400	\$87,360

The cost comparison data in Table 1 reflect only the most basic planned expenses of a HPC system. That said, they do account for the frequent unplanned costs that accompany such environments, including consultant fees, administrator overtime, additional space build-outs (i.e. electrical, cabling, racks, etc.), and the costs of other IT personnel needed to assist with change management and IT architecture planning.

From a pure cost perspective, outsourced HPCC clearly lowers the CapEx required to perform HPC. Beyond the hard cost items, HPCC affords ultimate business flexibility. Businesses can now utilize many hundreds of GPUs in contrast to an internal strategy whereby they could only have afforded a few dozen GPUs. The lower cost and added flexibility of HPCC means they can complete tasks faster and more efficiently.

Managed HPC Services are Not All the Same

Managed HPC services can vary greatly, particularly in terms of the infrastructure and levels of security provided, ease of use, customer control, and pricing. IT leaders should gain a solid understanding of the following service elements prior to choosing a HPCC service provider.

- **Professional-grade managed hosting provider.** Managed HPCC services are natural extensions of the services that professional-grade managed hosting providers have offered for years. While Amazon, Google, and several other public cloud companies attempt to develop HPCC revenue by tapping into the unused portions of their datacenters, they lack the core competency of full-service hosting providers: delivering a fully supported, dedicated compute infrastructure. There is simply no comparison between the latter's guaranteed service level agreements (SLAs) for the most stringent hosting requirements and the lackluster SLAs of Amazon and Google. No serious user of HPC would outsource their critical business enabler to an organization that doesn't back the service with 24x7x365 support from highly skilled engineers.
- **Incorporation of GPUs into core hosting infrastructure.** Best-in-class HPCC service providers blend their HPC infrastructure into their core private cloud capabilities. As a matter of practice, professional-grade cloud business customers receive customizable resource pools, with no need to work their applications and processes around pre-defined virtual machines. The customers own the infrastructure, or resource pools, for as long as they need to use them, enabling them to maintain system administrator control. Best-in-class providers achieve this blending of HPC and private cloud by inserting GPUs into an individual customer's private cloud infrastructure.
- **Private cloud with no multi-tenancy.** While all data is valuable, HPC data represents the lifeblood of many companies. Organizations simply cannot risk relying on shared infrastructure where their data and processes become exposed to the shortcomings of other organizations. With multi-tenancy, HPC data and operations can fall prey to Distributed Denial of Service (DDoS) attacks, hacking, and even simple administrator errors of other organizations residing on the same resource pool. HPCC within private cloud environments eliminates these risks. Leading HPCC service providers with single-tenant environments can quickly identify, isolate, and respond to issues faster, because they know exactly who owns, who touched, and who is responsible for everything that happens to a specific resource pool. As a result, they support a customer's entire hosted infrastructure – not just the HPC portion – from end to end.
- **On-demand scalability and flexible pricing.** Organizations are pursuing HPCC services primarily to gain responsiveness to business needs, save time, and save money. For CIOs to guarantee themselves success in each of these areas, they need to clarify the terms of their HPC outsourcing agreement. Given the extreme variability in usage definitions and pricing structures of HPCC, they should push a potential service provider to identify the specific resources they must pay for and how they must pay. Leading service providers enable customers to both scale up their infrastructure during peak workload periods as well as scale down to save costs during lower workload periods. As a result, customers do not need to predict the amount of resources they need.

This flexibility and ability to scale up and down on demand represent significant ROI potential when outsourcing HPC, particularly for organizations that struggle to predict the timing of their next big HPC project. Reputable service providers should have no problem producing price lists that enable customers to plan. These should include definitions for computing instances, as a frequent complaint among Amazon customers is that they incur price gouging for partially and barely used instances.

Case Studies

Outsourced HPC services can take many forms. For that reason, CIOs should tread carefully as they evaluate the capabilities, contractual terms, and flexibility of potential service providers. The following two case studies demonstrate how such services can vary.

Occipital: Delivering Innovative Views to Mobile Users

Occipital provides real-time panorama creation for millions of iPhone, iPod, and iPad users. They just tap their device screen and pan in any direction. As they do, they see their panorama being built in real time as every incoming frame is processed. When they're done, users can instantly save their panorama or share via email, Twitter, or Facebook—no need to wait for it to stitch. They can even view their newly made panoramas as photos or experience them in an immersive 360-degree view.

To provide this unique service, Occipital originally relied upon end-user devices to stitch together and render the images using a unique algorithm. Once images were complete, users would then send them as JPEG files to Occipital. Occipital would then enhance the images using the more powerful CPUs on their own servers and then return the enhanced images to users.

The Challenge

Occipital anticipated competition entering the market. To maintain its strong leadership position, the company wanted to provide users with the highest quality images possible and in less time. To achieve this, they knew they had to leverage GPU-based HPC. “We wanted to deploy a new algorithm to run on our servers to both remove the burden of rendering from the client-side device and improve the quality of the images,” said Jeff Powers, CEO of Occipital. “Massively parallel GPUs raise the bar higher, and we're expected to innovate to remain competitive. People want better quality products, and you just can't pull that off on a CPU in a reasonable time.”

The Solution

Occipital considered purchasing GPUs and running their own HPC servers in-house. However, after analyzing the need for strong network performance, redundancy, and 24x7x365 uptime of their system, Occipital management decided to outsource their HPC to a service provider.

They evaluated Amazon EC2 and several traditional hosting providers, eventually choosing the HPCC service from PEER 1 for two key reasons. “First, we needed full control over the machine to run OpenGL, and we couldn’t do that with any other service where we could get GPU access,” stated Powers. “We played with EC2 cluster compute instances, but we couldn’t get them to do OpenGL for us either.” Today, Occipital leverages their PEER 1 single dedicated-instance GPU server to satisfy growing user demand.

The Benefits

Since deploying on PEER 1’s HPCC, user satisfaction has skyrocketed for Occipital. Rave reviews about image quality top the feedback from users who are astounded by their new-look panoramas. “We are pleased to further differentiate ourselves while simplifying our compute infrastructure,” stated Powers. “It’s hard to believe that we can do this with a system that we control but which acts just like a normal Web service.”

Occipital has integrated its HPC computing and normal hosting needs onto the PEER 1 service. As such, the company gains improved network performance for all its company-wide needs while saving budget by avoiding costs associated with transporting data from its panoramic images.

Case Study: Media Software and Services Provider Speeds and Secures Hollywood Studios’ IP

Established in 1995, this media company is an award-winning provider of high quality, differentiated video encoding solutions. Through their encoding solutions and desktop applications, the company empowers enterprises, video professionals, and small-to-medium sized organizations to easily and affordably encode, manage, and deliver the highest quality video online.

As video rapidly proliferates online and production grows more sophisticated, the company’s clients rely on them to deliver the most leading edge services and highest quality video possible. In many cases, the company manages the video encoding software systems for customers, specifically those concerned with keeping up with the industry’s fast pace of change.

The Challenge

Encoding video is both time and compute intensive. In fact, just one hour of HD video can take more than an hour to convert to H.264/Advanced Video Encoding (AVC), the industry’s standard for compression. “With customers providing us with high-bit-rate video in multiple formats and resolutions, we were spending more than ten times the original video time to convert to H.264,” explained the company’s Technical Director. “Customers hated the slow turn-around. We simply had to get faster converting to H.264, because it’s the sole standard for Apple devices and one of the largest formats for Web delivery.”

The solution

To increase the speed of video conversion, the company decided to leverage HPC infrastructure from a HPC services provider. They evaluated several providers, including Amazon, and eventually chose the High Performance Private Cloud Computing (HPCC) service from PEER 1. “With PEER 1 we get the best combination of performance and flexibility,” continued the Technical Director. “Amazon offered just one CPU-based instance that they defined and we couldn’t change. With PEER 1’s private cloud, we could say ‘put this on a lower end server with GPUs,’ making it faster and less expensive.”

The Benefits

By having remote access to a dedicated HPC machine at PEER 1, the company now remotely installs its customers’ videos and has cut encoding times by 200-300% when compared with CPU-based solutions. “If you’re a Hollywood studio and need to encode, well, time is money,” stated the Technical Director. “We can now speed the conversion process with fewer resources than we used before.”

The company has realized both cost and speed benefits using the PEER 1’s HPCC service. At the same time, they have not had to make any tradeoffs regarding the safety and security of their customers’ video assets. According to the Technical Director, “Customers quickly realized that our security is robust using PEER 1. From a physical server perspective, it’s all locked down and has redundant data. And everything’s encrypted from our customer content side, so no one can do anything without proper user credentials.”

Conclusion

The pace of business is simply too fast today for data-intensive companies to manage their HPC needs in-house. The need to shed CapEx while innovating at a fast and furious pace pits innovation against budgetary needs. Fortunately, HPCC provides a practical solution that satisfies the needs of both camps.

Leveraging HPCC, leading organizations across the industries of oil & gas, digital content creation, banking & finance, and government & scientific research have seized on the new reality – that outsourcing HPC is just like outsourcing any other application today. As a result, they are reaping the benefits of improved agility, innovation, and speed while maintaining a low-cost footprint that can evolve as market dynamics require.

“This notion of disposable computing dramatically reduces the barrier for research and development! Clouds will surely revolutionize how HPC is applied because of its utilitarian usage model. Clouds will make HPC genuinely mainstream.”⁸

**– Wolfgang Gentzch, HPC Consultant
and General Chair of the ISC Cloud
Conference**

Trying outsourced HPC managed services is easy. Sign up today for a test drive of HPCC at www.peer1.com

Footnotes

¹ <http://en.wikipedia.org/wiki/GPU>

² IDC, Financing a Software Infrastructure for Highly Parallelised Codes — IDC FINAL Report for the DG Information Society of the European Commission, July 2011.

³ Gartner, Enterprise Architecture Special Report: Overview: Accessed January 8, 2012 from the following source: <http://www.gartner.com/pages/story.php.id.2230.s.8.jsp>

⁴ Ibid.

⁵ Gartner, Beware Misleading Marketing of "Private Clouds," December 9, 2011. Accessed January 8, 2012 from the following source: http://blogs.gartner.com/lydia_leong/2011/12/09/beware-misleading-marketing-of-private-clouds/

⁶ Gartner, Gartner Reveals Top Predictions for IT Organizations and Users for 2012 and Beyond, December 2011. Accessed January 8, 2012 from the following source: <http://www.gartner.com/it/page.jsp?id=1862714>

⁷ Deloitte, About Strategy Jobs, Accessed January 8, 2012 from the following source: <http://jobs.deloitte.com/new-york/strategy/strategy-%EF%B9%A0-operations-infrastructure-operations-outsourcing-advisory-services-bpo-manager-jobs>

⁸ HPCwire, Upcoming Event to Explore the Role of Clouds in HPV, February 15, 2011. Accessed January 23, 2012 from the following online source: http://www.hpcwire.com/ar/hpccloud/2011-02-15/upcoming_event_to_explore_the_role_of_clouds_in_hpc/