insideHPC

The insideHPC Guide to

Cloud Computing

BROUGHT TO YOU BY

PENGUIN COMPUTING

Introduction

IT organizations are facing increasing pressure to deliver critical services to their users while their budgets are either reduced or maintained at current levels. New technologies have the potential to deliver industry-changing information to users who need data in real time, but only if the IT infrastructure is designed and implemented to do so. While computing power continues to decline in cost, the management of large data centers, together with the associated costs of running these data centers, increases. The server administration over the life of the computer asset will consume about 75 percent of the total cost.

Cost to Manage a Small Server Farm

Number of servers = 200

Loaded cost of system admin = $ 150,000

Useful life of server = Five years

Cost to manage a server for Five years = ($ 150,000/200) * Five years = $ 3,750

Percent of server cost = ($ 3,750/$5,000) = 0.75

Scaling an existing data center in many cases is not possible as demands from users increase, and as the real estate or construction costs become too high or unavailable. Cloud computing can offer IT departments the ability to respond to critical needs while costs remain under control and predictable. With more compute power available to organizations than previously possible, business innovation on many fronts is just a click away when working with a cloud computing provider.

The term, "cloud computing" can mean a variety of things to different people. The National Institute of Science and Technology (NIST), which is part of the U.S. federal government, defined the term, which is used in U.S. federal government activities and broadly encompasses a number of important points.

Contents

The NIST also defines characteristics that are essential to a cloud computing environment. These include:

On-demand self-service
The consumer can acquire computing and storage as needed without requiring human interaction at the service provider.

Broad network access
The services provided can be accessed through standard mechanisms and include a range of client devices.

Resource pooling
The provider's computing or storage environment should allow for multiple uses (tenants) at the same time.

Rapid elasticity
The capabilities can be expanded or contracted depending on the demand of the user.

Measured service
The billing reflects the resources used in some abstracted fashion and may include compute, storage and network use. These resources can be monitored as used.

> "Cloud computing is a model for enabling ubiquitous, convenient, on-demand network access to a shared pool of configurable computing resources (e.g., networks, servers, storage, applications, and services) that can be rapidly provisioned and released with minimal management effort or service provider interaction."[1]

At a fundamental level, GridGain enables promoting data up from residing in slow mechanical storage systems to fast memory. GridGain's IMDG solves many critical IT pain-points at once:

- **Performance**
- **Scalability**
- **High availability**
- **Data consistency and reliability**
- **Detailed insight and management**

Cloud Computing Models

While the concept of cloud computing can be generalized with the previous definitions, different models of use have come to be used, and in some cases providers have focused on one or more of these models.

Software-as-a-Service (Saas)

Users are able to run the cloud provider's applications, which are running on the provider's cloud computing infrastructure. The user interface is available through various client devices. The end user does not have control over the servers, storage, networking, etc.

Platform-as-a-Service (Paas)

Users can deploy their own software to the cloud provider's infrastructure using various development tools, libraries, programming languages and tools that are available. The user does not have control over the underlying hardware systems, but can control how the software is deployed.

Infrastructure-as-a-Service (Iaas)

The capability provided gives the user or developer the most control over the environment. The consumer can provision, through software, the storage, networks and other resources to run arbitrary software. Although the most flexible for the user, the underlying infrastructure still cannot be managed or controlled.

Economic Savings

For companies and organizations considering moving from an in-house data center to an external one (the cloud), financial factors may drive the decision. There are a number of factors to consider when comparing an on-site data center with a cloud-based provider.

Capital expenditure (CAPEX)

As new hardware is developed and released by systems vendors, the cost to purchase servers must be considered. Many data centers have thousands of systems that are accessed continually, so careful planning is needed to ensure that only portions of the computing or storage infrastructure are replaced at a given time. As users demand more interaction with larger and larger data, new computer servers must be acquired. As newer and faster processors are released frequently by CPU providers at similar costs, IT purchasers must refresh systems on a regular basis, just to satisfy their end-user demands. This can lead to significant amounts of CAPEX on a regular, repeating basis.

Operational expenditure (OPEX)

As data centers have become more complex, significant effort is required to plan, maintain and manage the resources. With every new system acquired, it must be provisioned and integrated into the existing environment. Storage systems must be monitored for failure, and backup services implemented. In addition to managing the systems, software and storage, other costs such as power for the hardware and for the cooling and air circulation must be planned for. Each new compute or storage system requires a certain amount of power and cooling. Expert personnel are a significant cost as well. Staff that are knowledgeable in the specifics of complex applications (tuning for certain inputs) are required for an efficient implementation for the desired workloads. Over time, the operational expenses will be more than the original cost of the hardware.

Cloud Computing Components

Servers

The most visible component of most cloud providers' offerings is the compute resource itself. Cloud providers typically will offer various "instances" of computing power. These instances of compute power are what the end user pays for, either by the minute, hour or longer periods, multiplied by the number of cores used. Cloud providers typically will have a range of instances with varying clock speeds, and amount of memory per instance. In many cases, a virtual CPU (vCPU) is the unit of measurement, which is different from a physical core. Using virtualization technologies, a vCPU sits on a physical core, but is not equivalent. A number of vCPUs can share a physical core, thus reducing performance as applications will compete for resources. The more vCPUs reserved at higher clock rates with more memory will cost the consumer more money.

Storage

Storage can be accessed in various ways from cloud providers. Perhaps the most common use is just to keep files stored (or copies made as backup) on a cloud provider's storage infrastructure. Some common and visible examples of this type of use would be Google Drive, Box and Dropbox. These offerings are not tied to a computing environment and act as just another drive that can be accessed by various clients and applications. For cloud providers that offer a complete solution, they may or may not charge for storage, in conjunction with various instances. A cost may occur if the amount of storage exceeds a certain level for an amount of time that the data exists on the cloud provider's disks. In addition, a storage amount may be included as part of the purchase of an instance of compute.

Software

The software accessed from a cloud provider can be from various sources. An independent software vendor (ISV) may have an agreement with the provider to have their software installed in the cloud environment. In other cases, the user may have to upload the software executable in order to run the software, perhaps from a company directory of approved software. When licenses are required to run the ISV software, mechanisms must be put in place to assure that the user has the privileges to run the software and that the accounting is set up for software billing back to the ISV. A number of different mechanisms exist for ISV licensing, depending on the ISV policies and the customer's relationship with the ISV.

Service

As with any type of product that is provided to a customer, the level and expertise of the cloud provider will be very important in the overall customer experience. In the case of cloud providers, this may include upfront studies as to the return on investment of the use of the cloud services, help with uploading input files, tuning of the computer systems available, and close monitoring of the customer's resource use. Although cloud computing is designed (and defined here) as being self-serve, excellent and timely support may be needed for novices and special, more complicated uses.

Three Main Types of Cloud Computing

The deployment method and, thus, the use for cloud computing has emerged with three main types, each with its main user communities:

1 PRIVATE CLOUD

The cloud infrastructure and software to provision and use the computer, storage and network systems are available for use by a single organization. This organization can be an entire company or divisions within a company. The physical infrastructure may be housed on site or off site from the organization that is using the systems. This would typically be in the existing company's data center, but could be housed and implemented elsewhere.

2 PUBLIC CLOUD

In this case, which is the most widely discussed in the IT industry, the cloud infrastructure resides at an entirely different company and is available to the general public. It may be owned and operated by commercial companies, universities, research organizations or even governments that make the systems available to those outside of these organizations. While not solely HPC focused, the more familiar of the public cloud offerings are owned and operated by Amazon (Amazon Web Services) and Microsoft (Azure), in addition to many smaller offerings.

3 HYBRID CLOUD

Hybrid clouds can be thought of as being composed of more than one different cloud infrastructures. This can be a combination of private and public cloud implementations. Through software and communication links, these different cloud implementations can act together. An example is that when an organization's private cloud becomes oversubscribed, the additional work can be sent to another cloud instance, which can be provisioned to run the same software stack. In some cases, this is called "cloud bursting" and has the effect of load balancing between cloud implementations.

Why HPC in a Cloud Environment?

High-performance computing (HPC) generally refers to the practice of aggregating computing power in a way that delivers much higher performance than one could get out of a typical desktop computer or server in order to solve large problems in science, engineering or business. Depending on the problem type and the science and algorithm used, these applications can require just one or thousands of compute cores running together. The range of use for HPC-type technologies runs the gamut from a small company designing components for a larger integrator (think automobile supplier for a major automobile manufacturer) to Department of Energy (DOE) labs that own and maintain the fastest computers in the world that, for example, investigate climate change and nuclear weapon simulations. Although many of these use cases are not suitable for running their HPC and ISV applications in the cloud, a significant number of those running technical programs can benefit from cloud computing environments. There are a number of business reasons to consider HPC in the cloud.

Project deadlines

There are times of the year when extra compute power is needed, but only for a short period of time, such as in order to meet a project deadline. It does not make financial sense to purchase hardware that may only be used for one week per quarter, or perhaps during the final design run before production. In some organizations, the lack of adequate HPC servers can create long wait times in queues, reducing innovation or causing projects to be delayed.

Reduce in-house infrastructure

Running and maintaining a data center for HPC applications is costly and requires power and cooling in addition to the servers and storage. Dedicated parts of a physical building must be built and include proper HVAC systems and/or water cooling systems. This can get very expensive and the square footage cannot be used for other business-related purposes.

Latest software versions

ISV applications, which are used heavily in commercial HPC environments, release new versions periodically. With an in-house data center, someone must be sure to monitor the ISV update and release schedule, and then install either patches or new versions. By using a cloud provider that has relationships with the ISV, this maintenance is taken care of and performed without the end user having to be concerned.

Latest hardware available

New and faster processors and systems incorporating these are released frequently. If a company or organization were to purchase new servers or storage, a process involving many people would have to be followed, and include creating a request for proposal, understanding the bids, evaluating choices, and implementing and integrating the new servers or storage into an existing environment. The entire process would take time and personnel expense. By counting on the cloud provider to house the latest and highest-performing servers and storage systems, organizations can focus on their business and services more and not become burdened by having to purchase servers, storage and networking.

Reduce in-house maintenance

Even when a data center is up and running, time must be spent to monitor and maintain the infrastructure. From downloading patches and new versions and monitoring the system's health, a significant amount of time must be allocated for these tasks. System administrators may have to be on call 24 hours a day to solve critical issues. When working with a cloud provider, these tasks are the responsibility of the cloud provider and are part of a service-level agreement that would include uptime guarantees and network access.

A recent IDC survey indicated that about 25 percent of sites that ran HPC workloads are using some sort of cloud computing[2], and that just over 30 percent of the HPC workloads were being performed at cloud sites.

Critical Cloud Requirements for HPC

Users in HPC environments have requirements for using a cloud provider that are different than typical enterprise applications. When running an application that is designed to return an answer or solution in the fastest time, performance matters. An in-house data center may serve hundreds or thousands of users for a certain class of applications where absolute speed is not critical. However, those involved in HPC-type applications demand the fastest response. There are a number of considerations for ensuring maximum performance for running HPC applications in the cloud.

1 Right type infrastructure

Make sure that an HPC infrastructure is designed in from the beginning, not cobbled together as an afterthought. This would include:

- **Latest servers, adequate number of servers** Make sure that any application that is submitted to the cloud can run, based on the characteristics of the application (some are more scalable than others) and the limits imposed by the user (maximum number of cores to be used).
- **Bare-metal for HPC –** Users demand the fasted performance, which is achieved by running the application directly on the server's OS. This is contrary to using a virtual machine (VM), which is another layer of software and can limit performance of using the computer system capabilities directly. The "hypervisor tax" can reduce performance by up to 10 percent, according to some estimates.[3, 4]
- **Accelerators –** GPUs can speed up the processing of certain types of applications. Incorporating a number of GPUs into the computer systems is increasingly important for reducing the time to completion of an application. For applications that can take advantage of one or more GPUs, it is critical for the cloud provider to offer this capability to the user.
- **Low latency and high-speed interconnect –** One of the most important parts of an HPC cloud offering is the connection between systems. Much of the HPC software ecosystem can take advantage of hundreds to thousands of cores running simultaneously, and requires system-to-system communication at very fast speeds. The current standard for HPC interconnect applications is InfiniBand (IB) and is a requirement for high-performing applications that can use many cores across many separate computer systems. Tightly coupled applications are often written with an API, the messaging passing interface (MPI) or with shared memory programming models. There are a number of MPI libraries available for use. While loosely coupled applications can get by using a GigE (1 or 10) network, closely coupled applications require the low latency and high bandwidth that IB delivers.
- **Geo-distributed resources –** For redundancy, many cloud providers maintain data centers in geographically different areas of the world. This can allow for the work and environments to be transferred to a data center that is up and running, should a significant event bring a data center down.

Bare-Metal Cloud Services

Many cloud providers provide a virtualization layer of software on the servers available for use. While virtualization software makes it easier to manage a range of cloud services, some customers require the faster performance, control and flexibility afforded by installing software directly onto the system. Virtualization is not a requirement for cloud-based computing as some may believe. A bare-metal environment that is an exact copy of the operating environment that may exist within an internal data center can be created.

2 Service

Although all cloud providers offer some sort of service to their customers, those that would use HPC applications in a cloud environment must be able to rely on the expertise that an organization that has been involved with HPC for many years has.

- **HPC experts –** The requirements for those using HPC in a cloud environment will differ from those of an enterprise customer. The details of how to get maximum performance from a given set of hardware will be paramount for the user. Specific libraries and the tuning and setup of the ISV application on the cloud provider's hardware requires the knowledge from dedicated HPC experts.
- **Smaller company, each customer matters –** Although a very large cloud provider can make available many types of instances and could scale into the millions of cores, most likely the provider is concerned with the larger enterprise software environment. Smaller, niche players can dedicate more expert resources to each customer, ensuring excellent customer satisfaction. In addition, if a customer requires a specific setup for their HPC application, a smaller provider will more likely be able to customize the environment for that specific user.
- **Cloud provider admins –** Operating a compute farm specifically for HPC applications that can be run on bare-metal servers requires a different set of administrative expertise than a broad-based virtual offering environment. Providers that have previous experience in selling hardware to the HPC community will be able to work with the consumer to solve their most difficult challenges.

3 ISV licensing agreements

A large portion of the commercial HPC market uses ISV software, which carries licensing costs. In order to provide a smooth experience, licensing with an ISV must be as seamless as possible.

- **License maintenance on cloud provider side –** The cloud provider works with an ISV to hold a license and then bills the user for the job completed or the resources used as defined by the license.
- **For those ISVs that don't provide a cloud license or short-term lease,** it is important that the cloud provider has the flexibility and expertise to incorporate traditional licensing mechanisms, either locally or via secure tunnels.
- **Wide range and number of agreements –** When running an HPC cloud infrastructure in order to make the environment easy to use, the provider should have as many licensing agreements in place as possible. The customers should not have to manage the licensing process. With a wide range of offerings in terms of ISV software, the experience will be more fluid.
- **Libraries, dependencies in place, cloud provider owns –** In most environments, the application requires a number of libraries to use, and depends on a number of factors for productive work. Dependencies (including specific versions) involving libraries and run-time software must be set up and integrated before actual runs take place.

4 Transparent cost model

Customers using a cloud provider for HPC applications acknowledge that a cost is involved. It is important to make transparent to the users the entire cost of running their applications in the cloud, in order to make valid and detailed comparisons to a comparison of on-premises versus the cloud.

- **CPU granularity billing –** Since a physical server contains a number of sockets, each of which contains multiple cores and an HPC application that will run for some time, the actual cost to the user is important. Billing based on actual cores used rather than the entire server will be more reflective of the actual use of the computer system.
- **Prices are public –** A number of charges by the provider to the customer would comprise the final bill. It is important that when using a cloud provider for HPC applications that all costs are made public so that the customer can make their own decision as to the benefits of using the cloud versus in-house options.

- **No transport costs –** Since the data upload and download for an HPC application may be significant, some providers may charge for this service. The value of using the cloud for running HPC applications is the application itself, so transport costs and billing should be avoided.
- **Storage costs –** As with any significant computing application, the amount of data produced can be large. It is reasonable to expect that there will be some sort of storage cost associated with HPC applications. Even if the transport costs are low to zero such that data could be moved back and forth to the customer site from the cloud provider, this would delay the startup of future runs on the same data as that data would have to then transferred back up to the provider. Understanding the storage costs in terms of capacity and time frame is very important.

5 Simple GUI

Access to a cloud provider's resources requires that the user interact in some way with those resources. The simpler the interface, the more time a customer can dedicate to using the cloud provider's resources without frustration over having to learn a new and unfamiliar access model. Obviously, a simple GUI run within a web browser will allow for the maximum level of ease-of-use for new customers.

- **Made for users, not system admins –** Users want to get their HPC applications run and not have to deal with complex access methods. Simple GUIs tailored to what the user must do to become productive are always more valuable than having to learn a new system.
- **Command line available –** For power users, in addition to a simple GUI, a command line interface may be desirable, especially when having to create scripts for various actions.
- **ISV licenses easy to use –** Accessing an ISV license should not have to involve setting path names and linking to remote systems. When an application starts, that cloud provider should be able to access the necessary licenses seamlessly.

6 Remote visualization

HPC applications will most likely produce significant amounts of data that must be post-processed in some way. Data from the application will be stored on a storage device and then can be either downloaded to the customer site or remotely visualized from the cloud site.

- **Post-processing after the HPC application completes –** Visualization of the data that resides on the cloud provider site requires specialized hardware (high-end graphics cards) to produce accurate visualization at acceptable frame rates. The challenge then becomes sending the images over the network, back to the user's desktop or other client device. The simpler that this process is for the user, the better and the more productive an engineer or analyst can perform. Various techniques and compression technologies can be used. It is important to understand the requirements for the client and the limitations on performing remote visualization. By creating the visualization on the cloud-based server, significantly less amounts of data must be transported back to the user. As the simulation post-processing increases in size, the use of cloud-based visualization can dramatically reduce network requirements.
- **Fast –** Experiments have shown that the maximum latency when performing interactive 3D graphics must be in the range of 100 to 120 milliseconds. If the delay from mouse movement to graphics screen update is more than this value, the user becomes frustrated and loses productivity. Thus, assuming an adequate network, the combination of hardware graphics on the remote end and display at the client end must be fast to enable a productive experience.
- **No added software –** Various techniques exist to move the image data from the cloud provider to the client. In some cases, software must be installed on the client to display the images that are being streamed across the network. However, a simpler solution exists that enables the display of images in a browser tab, using an HTML5-compatible browser. No added software is necessary.

7 API available

In some cases, the ability to script the entire session, without using a GUI for access to the cloud provider, can be valuable. Data uploading, job submission and notifications of job completion could all be automated and invoked automatically or with a simple action by the user on their local (client) system. However, this should be in addition to an easy-to-use GUI for a majority of the users. APIs also offer software vendors the ability to seamlessly plug a cloud resource into their software products and interfaces.

8 Queueing systems

When a user uses a cloud provider for HPC applications, maximum performance from the hardware, network and storage is the goal. However, a public cloud provider is designed for a multi-tenant use case. Thus, it is important that an HPC cloud provider implement the tools and controls to make sure that each application has complete and exclusive use of the cores and sockets in a specific server. Modern resource management systems can assure this, but the cloud provider must create the policies that ensure this.

Industries/Uses of Cloud Computing for HPC Applications

Various industries have adopted or are in the planning and evaluation phase for using the cloud for HPC applications. Within the realm of technical computing, certain workloads are suited to a cloud-based HPC environment. Workloads could be considered either loosely coupled or tightly coupled. In each of the industries discussed below, multiple jobs submitted with different input parameters would be loosely coupled and not require a low-latency, high-speed interconnect, while a job that requires the use of multiple systems working in concert would be tightly coupled and need IB. ***Figure 1*** shows a loosely coupled system where there is no need for node-to-node communication at high speed.

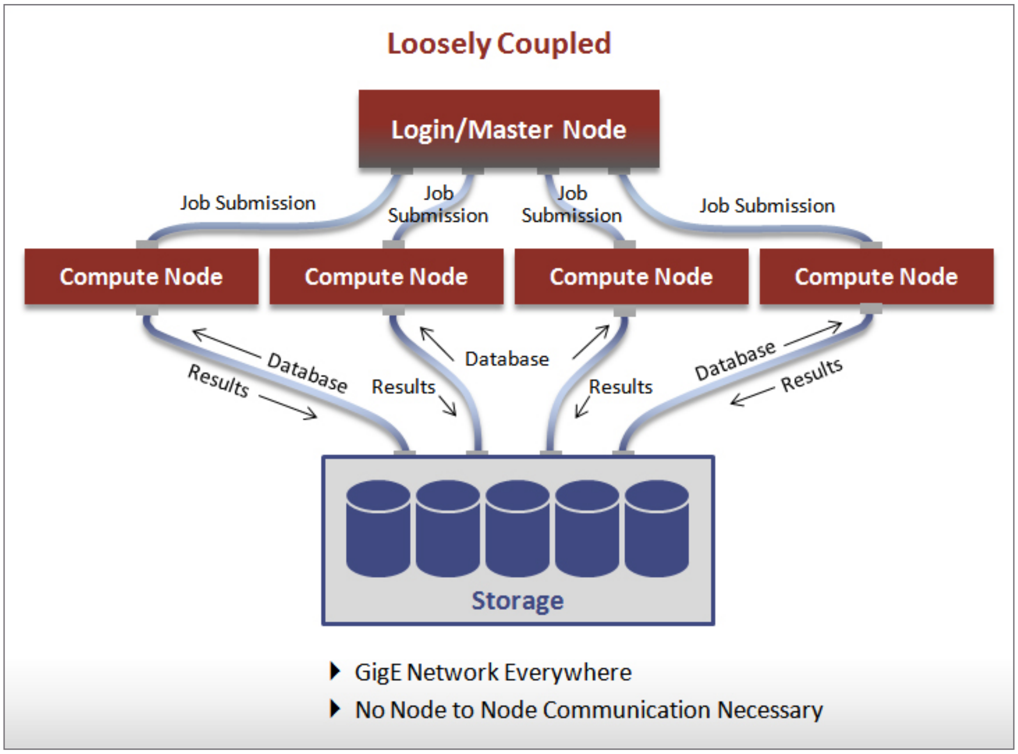

Figure 1 – Loosely Coupled System

Figure 2 diagrams a tightly coupled system where an IB network between the nodes is used for message passing between systems.

Computer-Aided Engineering (CAE)

As organizations design and develop products, the virtual testing for strength, reliability and operational capacity is well-suited for cloud computing. An organization's CAE environment may include finite element analysis, computational fluid dynamics, large-scale deformation, electronic design automation and dynamic simulation. All of these workflows can be offloaded from internal systems to a cloud infrastructure, as these types of applications are self-contained.

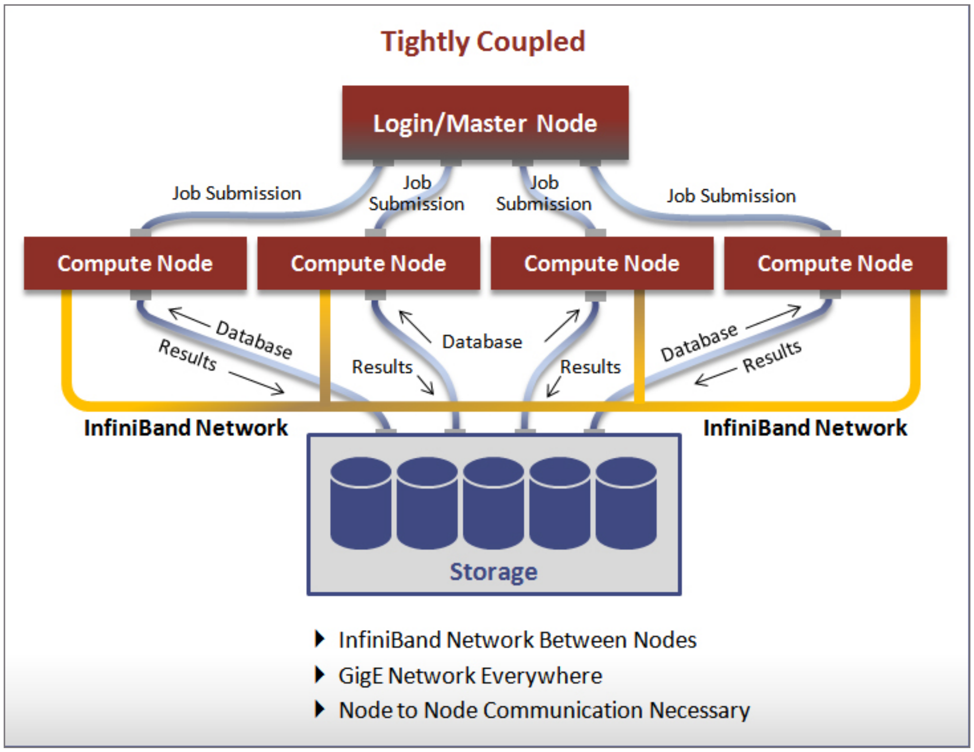

Figure 2 – A Tightly Coupled System with InfiniBand

Life sciences

Genomics and computational chemistry are two examples of applications that can be offloaded to a cloud infrastructure. Each requires significant computing resources with a large number of cores for maximum application performance. In addition, many of these applications and algorithms lead to tremendous performance when GPU accelerators are available.

Research

General research and simulations of physical phenomena can be moved to the computing resources available from cloud providers. In many cases in this category, applications will be developed by the users and can be easily packaged to use the cloud resources. However, applica-tions that would normally be run on a Top100 (www.top500.org) type system would not be applicable to a public cloud setup.

Energy exploration

Searching for the next oil or gas field requires significant computational power to move from the surveys to drilling locations. In addition, once the reservoir is discovered, simulating oil extraction is data and computationally demanding. A public cloud infrastructure would be ideal for this application. The only drawback would be the amount of data that must be uploaded for the application from the original acoustic sampling information.

Weather modeling and forecasting

Predicting the weather for the next hour, day, week or month has become extremely important for a wide range of organizations. Breaking up the water, land and atmosphere into massive amounts of 3D cells, scientists can model the chemistry and physics within that cell. The data (temperature, wind direction, etc.) then needs to be communicated to its neighboring cells and the simulation can continue. These applications require a very fast and low latency interconnect and are written using the MPI application programming interface.

Penguin Computing On Demand (POD)

POD is a cloud service designed and implemented for HPC applications. POD offers bare-metal access, IB interconnects and support by HPC experts. The servers and networking environment are what an organization might place in their internal data center, and do not require extra layers of software (hypervisor) to be present. Penguin Computing has been delivering HPC servers, storage and services to organizations across the globe for years, and is considered a leading company in the HPC market. Penguin Computing has taken its expertise in serving the HPC market and created a cloud service aimed at giving HPC users a cloud offering, tuned for the unique needs of technical computing professionals. ***Figure 3*** shows the differences between a virtualized server environment and an optimized, bare-metal environment.

"We appreciate the consistent quality of the cores that are being made available to us by Penguin Computing, enabling us flexibility in terms of volume load. The interactive framework is convenient to implement and allows us to use POD not only for solving, but also for pre- and post-processing," said Len Imas, senior CFD engineer for ORACLE TEAM USA.

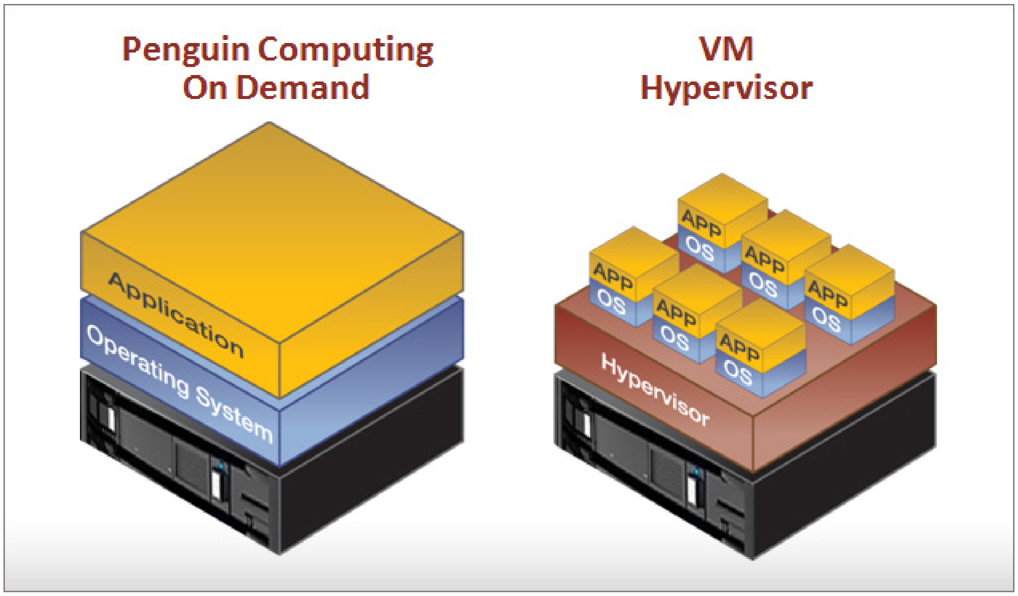

Figure 3 – POD Compared with Virtualized Environment with VM Hypervisor

POD is designed to integrate with and be familiar for technical computing professionals. POD's design allows users to create a workflow, whether using one application or multiple applications, that is not different from using an in-house HPC environment.

POD's user-facing side implements the on-demand self-service, access and measured service portion of the NIST definition of a cloud service mentioned previously, while leveraging established HPC practices for resource pooling and elasticity.

POD additionally implements a variety of user workflows, from the traditional HPC workflow to portal-based workflows or even API-driven workflows.

Job Submission Interface

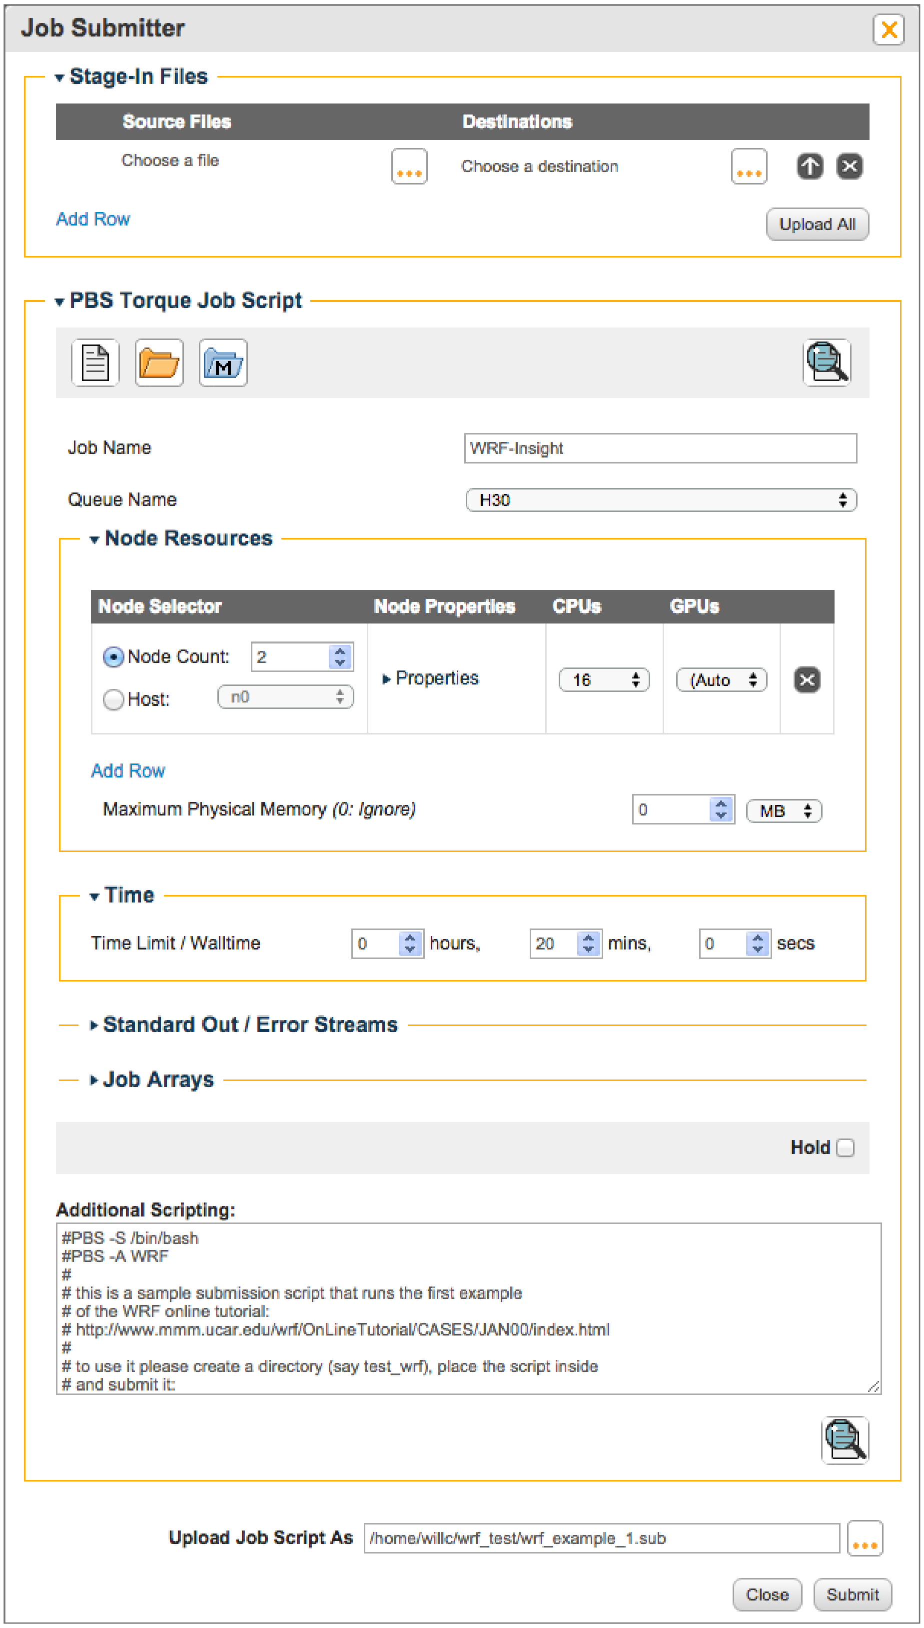

Figure 4 – Job Submission Interface

A typical portal-based workflow would involve the following steps (abbreviated here for illustration).

After logging on to the portal, a user can set up their job and determine the number of nodes and cores and whether a GPU is needed. Through simple scripting mechanisms (examples provided), a user or administrator can use the range of software provide by POD. ***Figure 4*** is a screen shot of the job submission task from the portal.

The next step would be to check if the job submission has no errors and is in the queue and scheduled to be run. ***Figure 5*** shows a summary of the job submission status.

Job Submission Summary

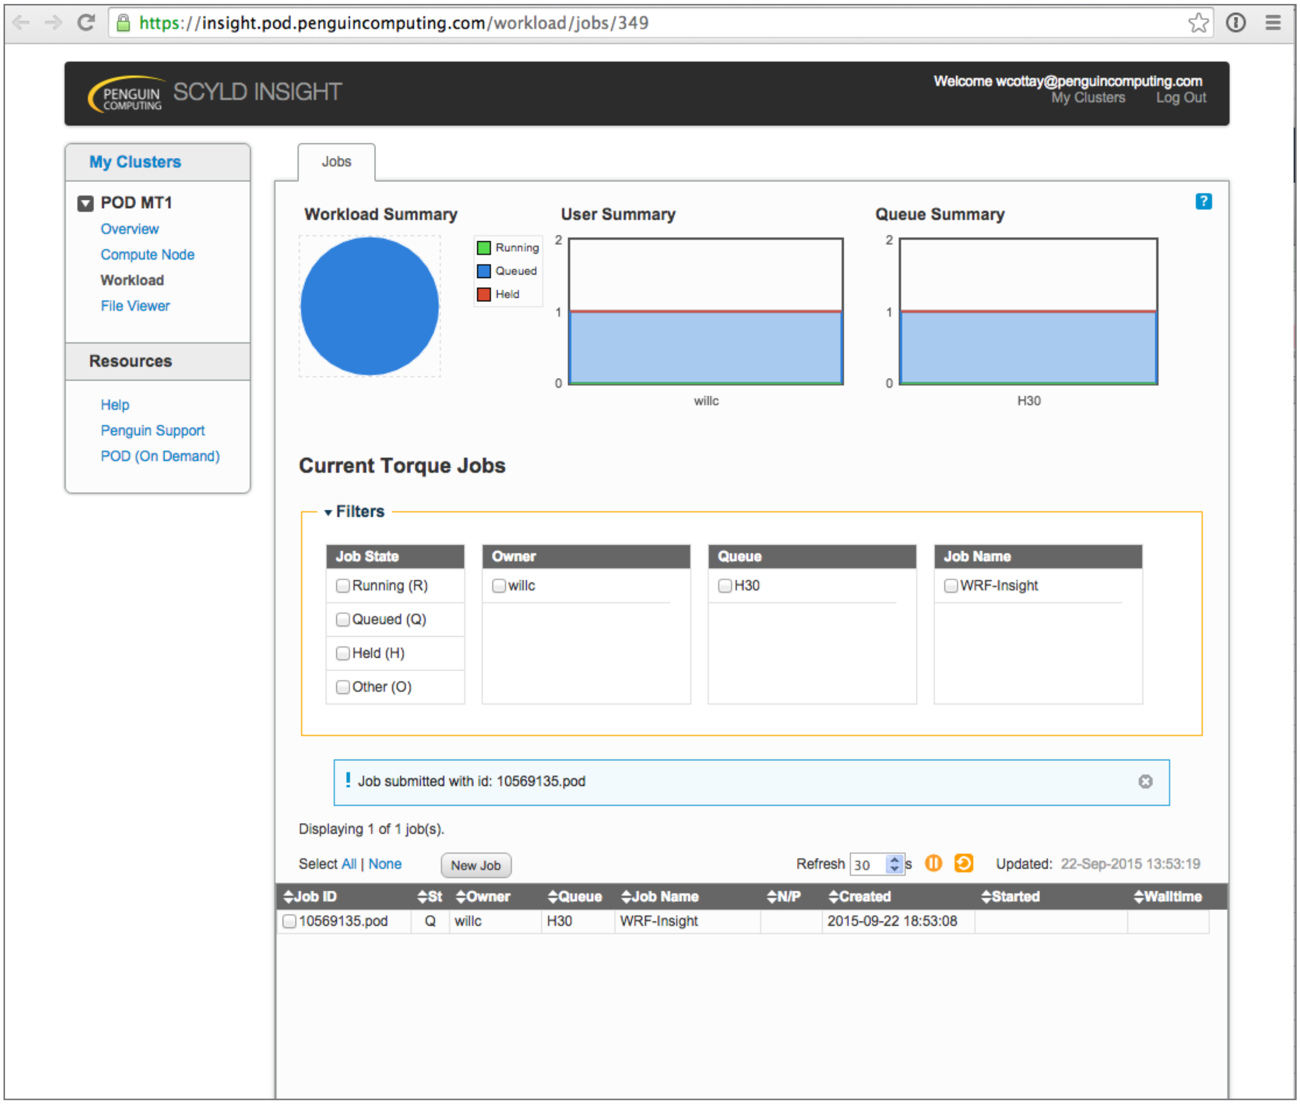

Figure 5 – Job Submission Summary

Once the application is running, it is important to be able to monitor the resources used and the state of those resources. To get a better understanding of the node that a job is running on, a user may wish to get a history of the node workload and other important statistics. ***Figure 6*** shows a number of performance measurements of a single node and the applications running on that node.

Single Node Performance

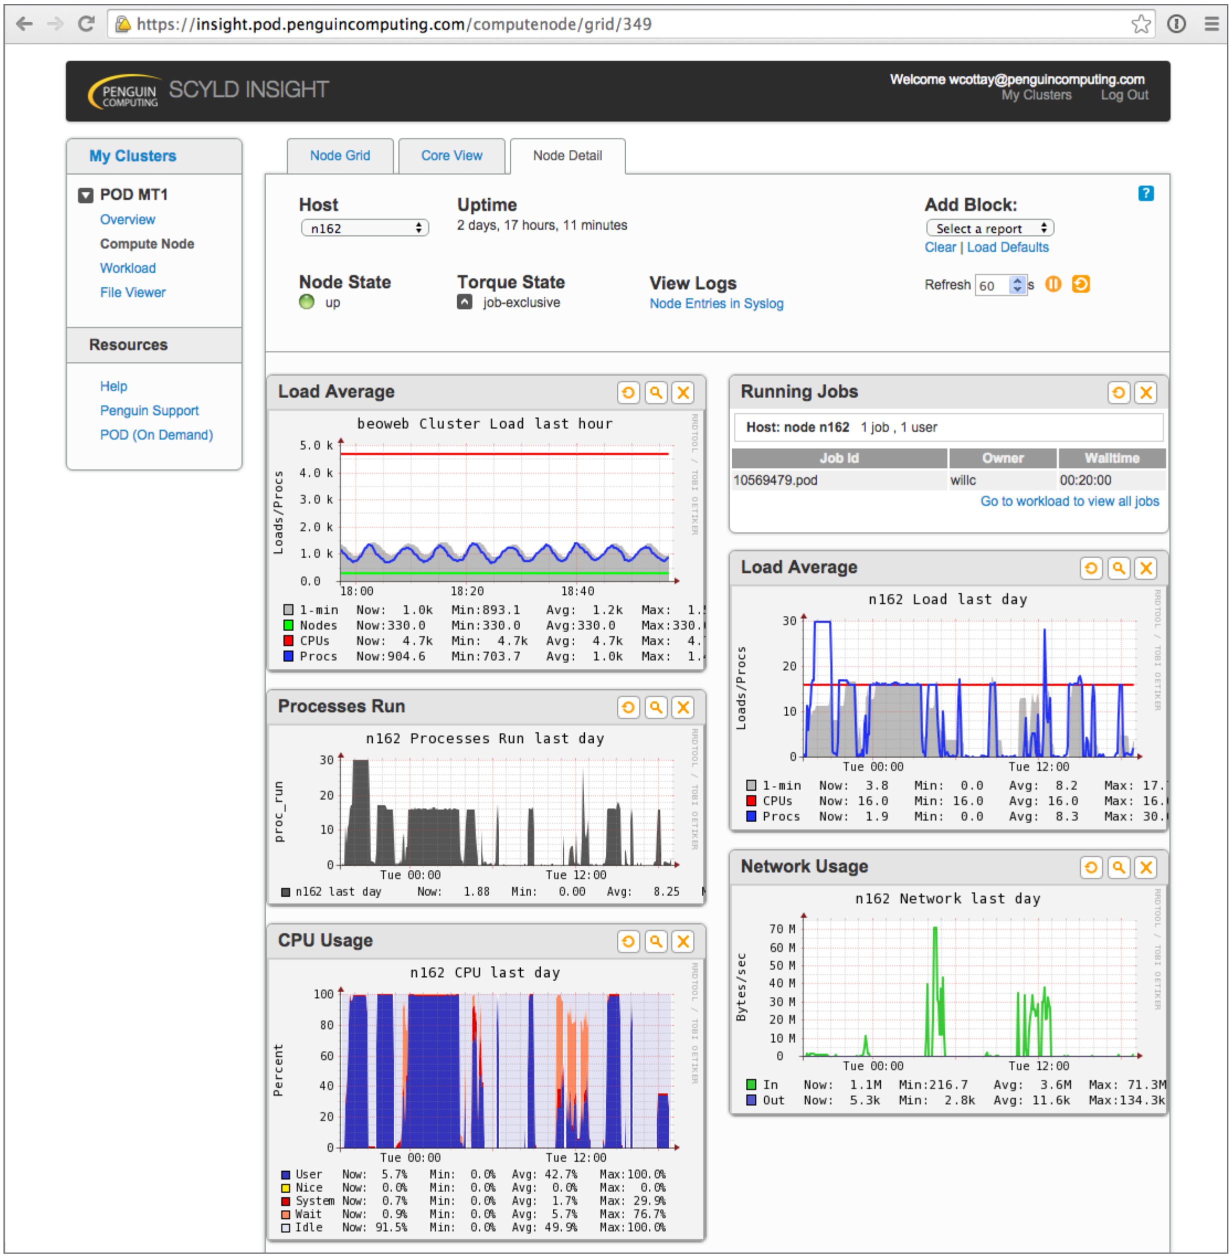

Figure 6 – Single Node Performance

While an application is running performance is key, especially in an HPC environment. Determining how a set of nodes in a cluster is performing — critical when running an application that requires multiple servers — gives users the ability to see if the performance of an application is what is expected. For example, POD allows users to monitor the memory used and the CPU utilization of every core the job is running on. By monitoring this in real time, issues can be addressed sooner rather than later, resulting in more productivity. ***Figure 7*** shows the performance of each core in a two-node setup and the amount of memory being used on each system.

Real-time Monitoring of Nodes Being Used

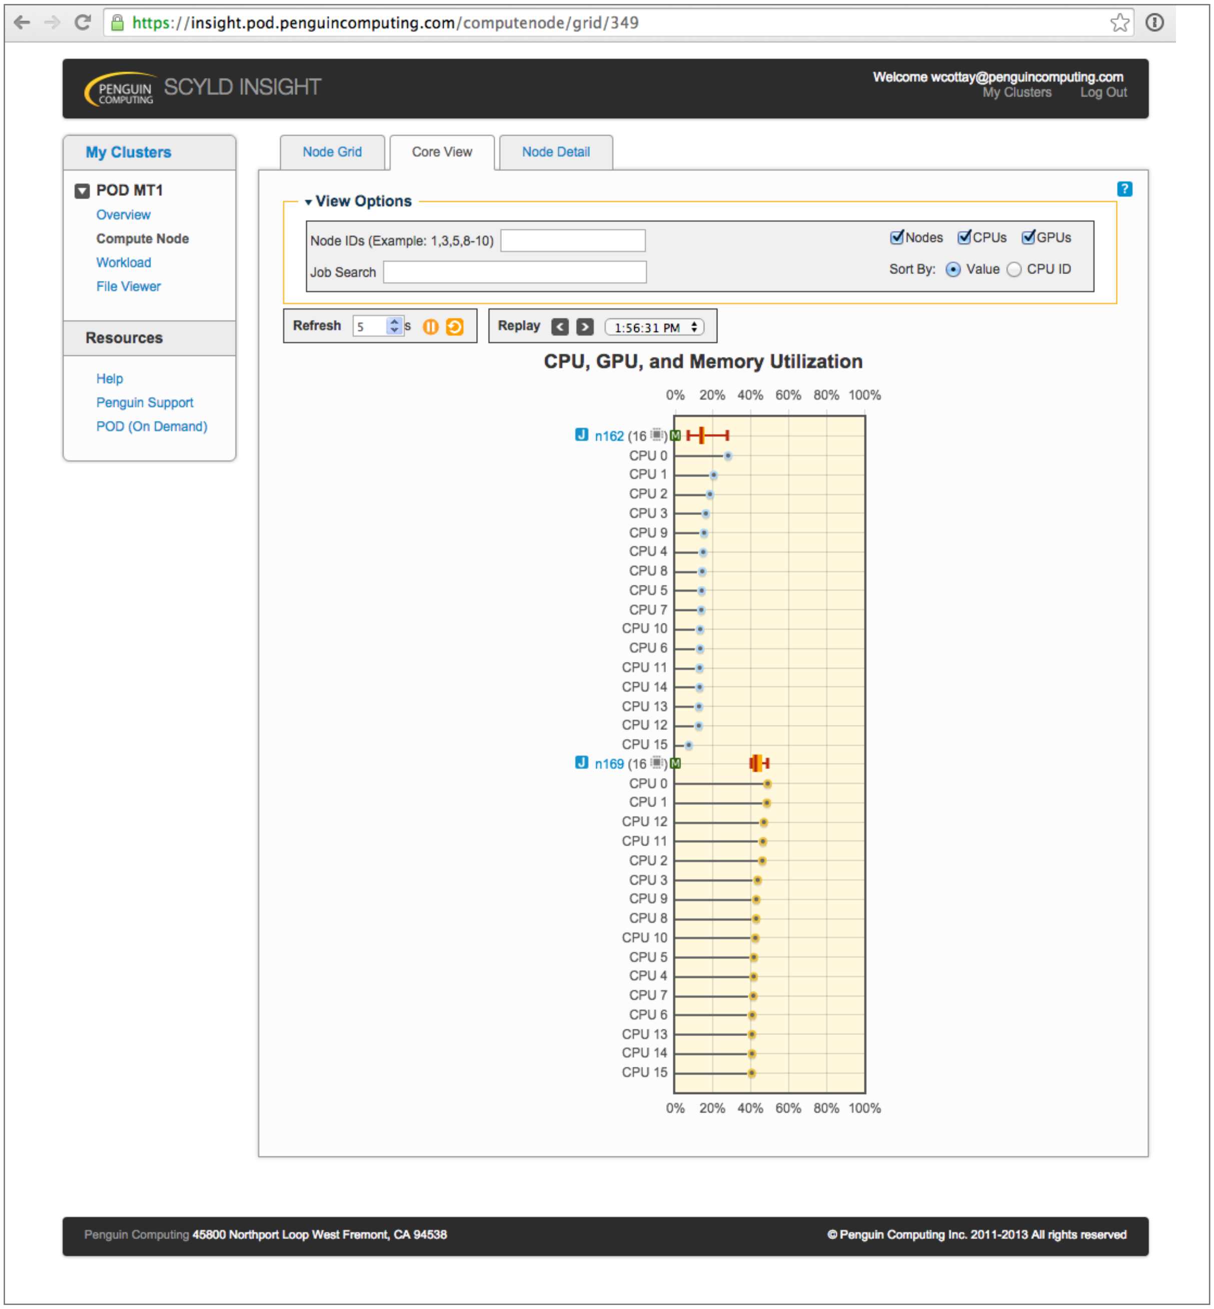

Figure 7 – Real-time View of Account Usage Summary

The POD facility has redundant high-speed data links with speeds up to 1Gbps. All access to POD is over secure, industry-standard, encrypted protocols: HTTPS+SSL, SSH through AES-128 or AES-256, or customized private VPN connections through AES-256.

POD pricing is transparent and is based on the number of cores used for computation and is measured to the nearest three seconds of use. Thus, users are only charged what they actually use. Real-time monitoring of the CPU and storage usage for their applications is easily available, so that users can determine what their monthly budgets should be for POD. ***Figure 8*** shows a sample of an account overview, so that team members can monitor usage of resources and accounting information.

Real-time View of Account Usage Summary

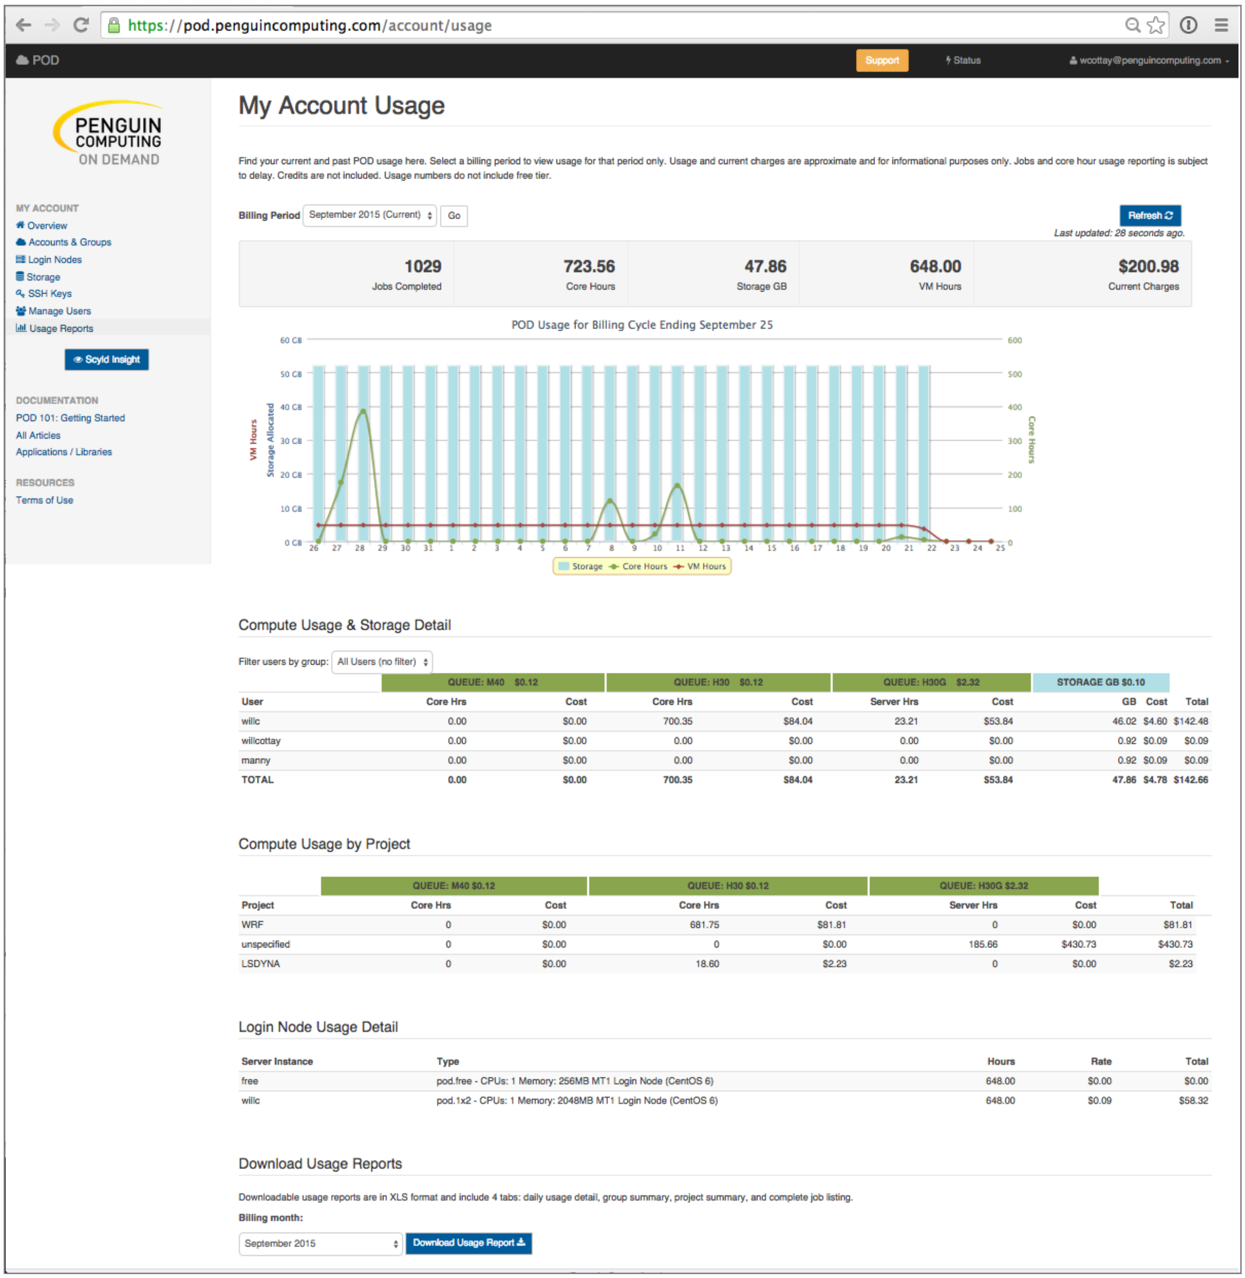

Figure 8 – Real-time View of Account Usage Summary

"We are pleased to work with Penguin Computing to enable users to run their applications faster and solve larger, more intricate problems without increasing the IT cost associated with maintaining high-performance capacity," said Silvina Grad-Freilich, senior manager of Parallel Computing Marketing at The MathWorks.

After the application has completed, 3D postprocessing can be performed. Using the Scyld Cloud Workstation software, which is integrated with the computation elements in POD, users can interactively view the results using popular post-processing packages on their client. No software is required to be installed on the client and is accessed with an HTML5-compatible browser using H.264 compression technology. ***Figure 9*** shows the general process for viewing on a thin client, such as a laptop or tablet. All of the major browsers are supported.

Performance of applications is critical in the HPC environment and results are expected quickly. Due to the fact that POD has been designed for HPC workloads, the application performance will be similar to running in-house on the same types of computer systems.

POD offers bare-metal access to servers. No virtualization is used, unlike some larger cloud providers, which can decrease the performance of the application. The application runs directly on the Linux operating system with no virtualization layers to slow performance.

Remote Visualization Architecture

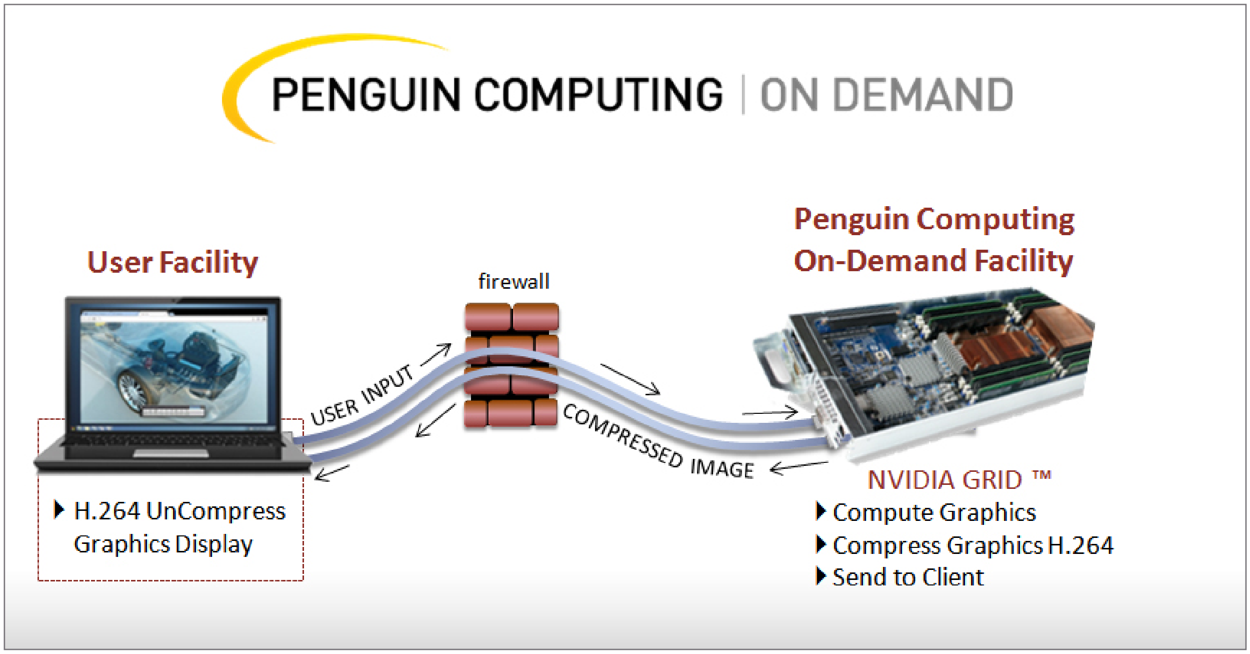

Figure 9 – Remote Visualization Architecture

The POD systems utilize the latest generation of Intel® Xeon® processors. Since Penguin Computing is an Intel® partner, the POD servers are internally designed and sold to end-user customers. A typical node will contain two sockets, and have at least 4GB to 6GB per core of RAM. GPUs from NVIDIA are also available for those applications that can take advantage of the GPU capabilities. NVIDIA is also a partner of Penguin Computing and has access to the latest generation of NVIDIA products. On the networking side, all POD servers are connected through a QDR IB fabric, which is non-blocking with switches that can switch faster than wire speed. This gives the user's application the fastest available interconnect for applications that run on multiple nodes at once. There is no extra charge for using the IB fabric, as all compute nodes contain this interface. 1 and 10 GigE interconnects are also standard on all servers. ***Figure 10*** shows a typical POD server.

An alternative and highly useful capability is the ability to use the resources from a cloud provider without the need for a full-featured GUI. POD makes available an API as part of the Scyld Cloud Management System, which allows for scripting and automation of authenticating and provisioning of an entire HPC cluster. The following steps can be automated at the customer site, so that, for example, after a user has completed their computer-aided design, the CAE analysis can be automatically sent to POD, with the following steps automatically performed:

- Authentication
- Generate a list of available PODs (Scyld Cloud Controller instances and their API endpoints)
- Create a system account on a POD cluster
- Provision a storage volume on a POD cluster
- Create a login node instance on POD
- Install SSH keys to access your POD environment

Figure 10 – Typical POD Server Consisting of Two Sockets Intel® Xeon Haswell CPUs

Penguin Computing has a long history of supplying and supporting HPC customers. With this history comes significant expertise in the HPC market, which is the foundation of the POD offering. HPC experts are part of the POD service team and have experience in a wide range of HPC domains. The architects of the POD infrastructure are experienced in designing systems for customers and thus all POD servers are Penguin Computing-designed.

Summary

The POD cloud offering has been designed for HPC professionals by HPC professionals using years of experience. As cloud computing has gained a strong foothold in companies of all sizes for developer and enterprise applications, the uptake of the cloud for HPC use has been lacking. By using years of experience and focusing on the HPC market, the POD offering is what HPC users have come to expect in terms of application performance, while reducing both CAPEX and OPEX. Users can easily upload data, perform calculations and visualize the results in an integrated, cost-effective manner.

Author: Michael A. Schulman

Michael is an experienced marketing professional in High Performance Computing. His interests lie in the areas of how HPC technologies can be used to produce new insights in various technical domains, as well as new HPC technologies that make access easier. His experience includes working at Silicon Graphics, Inc., and Sun Microsystems in the computer graphics and HPC organizations. He has a B.S. and an M.S. from Cornell University.

References

[1] NISTSpecial Publication 800-145 – "The NIST Definition of Cloud Computing" http://csrc.nist.gov/publications/nistpubs/800-145/SP800-145.pdf

[2] http://www.slideshare.net/insideHPC/hpc-market-update-from-idc-51061896, slide 11 and 12

[3] http://www.businesscloudnews.com/2014/06/19/hot-metal-how-virtualisation-is-no-longer-synonymous-with-cloud/

[4] http://www.prnewswire.com/news-releases/bare-metal-cloud-a-non-virtualized-cloud-option-for-performance-sensitive-workloads-241634951.html