An Introduction to High Performance Computing on AWS

Scalable, Cost-Effective Solutions for Engineering, Business, and Science

August 2015

Notices

This document is provided for informational purposes only. It represents AWS's current product offerings and practices as of the date of issue of this document, which are subject to change without notice. Customers are responsible for making their own independent assessment of the information in this document and any use of AWS's products or services, each of which is provided "as is" without warranty of any kind, whether express or implied. This document does not create any warranties, representations, contractual commitments, conditions or assurances from AWS, its affiliates, suppliers or licensors. The responsibilities and liabilities of AWS to its customers are controlled by AWS agreements, and this document is not part of, nor does it modify, any agreement between AWS and its customers.

Contents

Abstract

This paper describes a range of high performance computing (HPC) applications that are running today on Amazon Web Services (AWS). You will learn best practices for cloud deployment, for cluster and job management, and for the management of third-party software. This whitepaper covers HPC use cases that include highly distributed, highly parallel grid computing applications, as well as more traditional cluster computing applications that require a high level of node-to-node communications. We also discuss HPC applications that require access to various types of high performance data storage.

This whitepaper covers cost optimization. In particular, we describe how you can leverage Amazon Elastic Compute Cloud (EC2) Spot Instances[1] and storage options such as Amazon Simple Storage Service (S3), Amazon Elastic Block Store (EBS), and Amazon Glacier for increased performance and significant cost savings when managing large, scalable HPC workloads.

Introduction

Amazon Web Services (AWS) provides on-demand scalability and elasticity for a wide variety of computational and data-intensive workloads, including workloads that represent many of the world's most challenging computing problems: engineering simulations, financial risk analyses, molecular dynamics, weather prediction, and many more. Using the AWS Cloud for high performance computing enables public and private organizations to make new discoveries, create more reliable and efficient products, and gain new insights in an increasingly data-intensive world.

Organizations of all sizes use AWS. Global enterprises use AWS to help manage and scale their product development and manufacturing efforts, to evaluate financial risks, and to develop new business insights. Research and academic institutions use AWS to run calculations and simulations at scales that were previously impractical, accelerating new discoveries. Innovative startups use AWS to deploy traditional HPC applications in new and innovative ways, especially those applications found in science and engineering. AWS also provides unique benefits for entirely new categories of applications that take advantage of the virtually limitless scalability that cloud has to offer.

Using AWS, you can focus on design, simulation, and discovery, instead of spending time building and maintaining complex IT infrastructures. AWS provides a range of services: from virtual servers and storage that you can access on-demand, to higher level computing and data services such as managed databases and software development and deployment tools. AWS also provides services that enable cluster automation, monitoring, and governance.

What Is HPC?

One way to think of HPC is to compare HPC requirements to requirements for a typical server. HPC applications require more processor cores–perhaps vastly more–than the cores available in a typical single server, and HPC applications also require larger amounts of memory or higher storage I/O than is found in a typical server. Most HPC applications today need parallel processing—either by deploying grids or clusters of standard servers and central processing units (CPUs) in a scale-out manner, or by creating specialized servers and systems with unusually high numbers of cores, large amounts of total memory, or high throughput network connectivity between the servers, and from servers to high-performance storage. These systems might also include non-traditional compute processing, for example using graphical processing units (GPUs) or other accelerators attached to the servers. These specialized HPC systems, when deployed at large scale, are sometimes referred to as supercomputers.

HPC and supercomputers are often associated with large, government-funded agencies or with academic institutions. However, most HPC today is in the commercial sector, in fields such as aerospace, automotive, semiconductor design, large equipment design and manufacturing, energy exploration, and financial computing.

HPC is used in other domains in which very large computations—such as fluid dynamics, electromagnetic simulations, and complex materials analysis—must be performed to ensure a high level of accuracy and predictability, resulting in higher quality, and safer, more efficient products. For example, HPC is used to model the aerodynamics, thermal characteristics, and mechanical properties of an automotive subassembly or components to find exactly the right design that balances efficiency, reliability, cost, and safety, before spending millions of dollars prototyping a real product.

HPC is also found in domains such as 2D and 3D rendering for media and entertainment, genomics and proteomics analysis for life sciences and healthcare, oil and gas reservoir simulation for energy exploration, and design verification for the semiconductor industry. In the financial sector, HPC is used to perform institutional liquidity simulations and to predict the future values and risks of complex investments. In architectural design, HPC is used to model everything from the structural properties of a building, to the efficiency of its cooling systems under thousands of different input parameters, resulting in millions of different simulation scenarios.

HPC platforms have evolved along with the applications they support. In the early days of HPC, computing and data storage platforms were often purpose-built and optimized for specific types of applications. For example, in computational fluid dynamics (CFD) and molecular dynamics (MD), two dimensional engineering applications are widely used that have very different needs for CPU densities, amounts and configurations of memory, and node-to-node interconnects.

Over time, the growing use of HPC in research and in the commercial sector, particularly in manufacturing, finance, and energy exploration, coupled with a growing catalog of HPC applications, created a trend toward HPC platforms built to handle a wider variety of workloads, and these platforms are constructed using more widely available components. This use of commodity hardware components characterizes the *cluster and grid* era of HPC. Clusters and grids continue to be the dominant methods of deploying HPC in both the commercial and research/academic sectors. Economies of scale, and the need to centrally manage HPC resources across large organizations with diverse requirements, have resulted in the practical reality that widely divergent applications are often run on the same, shared HPC infrastructure.

Grids and Clusters

Grid computing and cluster computing are two distinct methods of supporting HPC parallelism, which enables applications that require more than a single server. Grid computing and cluster computing using widely available servers and workstations has been common in HPC for at least two decades, and today they represent the overwhelming majority of HPC workloads.

When two or more computers are connected and used together to support a single application, or a workflow consisting of related applications, the connected system is called a *cluster*. Cluster management software may be used to monitor and manage the cluster (for example, to provide shared access to the cluster by multiple users in different departments) or to manage a shared pool of software licenses across that same set of users, in compliance with software vendor license terms.

Clusters are most commonly assembled using the same type of computers and CPUs, for example a rack of commodity dual or quad socket servers connected using high-performance network interconnects. An HPC cluster assembled in this way might be used and optimized for a single persistent application, or it might be operated as a managed and scheduled resource, in support of a wide range of HPC applications. A common characteristic of HPC clusters is that they benefit from locality: HPC clusters are normally constructed to increase the throughput and minimize the latency of data movement between computing nodes, to data storage devices, or both.

Grid computing, which is sometimes called high throughput computing (HTC), differs from cluster computing in at least two ways: locality is not a primary requirement, and the size of the cluster can grow and shrink dynamically in response to the cost and availability of resources. Grids can be assembled over a wide area, perhaps using a heterogeneous collection of server and CPU types, or by "borrowing" spare computing cycles from otherwise idle machines in an office environment, or across the Internet.

An extreme example of grid computing is the UC Berkeley SETI@home[2] experiment, which uses many thousands of Internet-connected computers in the search for extraterrestrial intelligence (SETI). SETI@home volunteers participate by running a free program that downloads and analyzes radio telescope data as a background process without interrupting the normal use of the volunteer's computer. A similar example of web-scale grid computing is the Stanford Folding@home[3] project, which also uses many thousands of volunteers' computers to perform molecular-level proteomics simulations useful in cancer research.

Similar grid computing methods can be used to distribute a computer-aided design (CAD) 3D rendering job across underutilized computers in an architectural office environment, thus reducing or eliminating the need to purchase and deploy a dedicated CAD cluster.

Due to the distributed nature of grid computing, applications deployed in this manner must be designed for resilience. The unexpected loss of one or more nodes in the grid must not result in the failure of the entire computing job. Grid computing applications should also be horizontally scalable, so they can take advantage of an arbitrary number of connected computers with near-linear application acceleration.

A Wide Spectrum of HPC Applications in the Cloud

Demand for HPC continues to grow, driven in large part by ever-increasing demands for more accurate and faster simulations, for greater insights into ever-larger datasets, and to meet new regulatory requirements, whether for increased safety or for reduced financial risk.

The growing demand for HPC, and the time and expense required to deploy and manage physical HPC infrastructures, has led many HPC users to consider using AWS, either to augment their existing HPC infrastructure, or to entirely replace it. There is growing awareness among HPC support organizations—public and private—that cloud provides near-instant access to computing resources for a new and broader community of HPC users, and for entirely new types of grid and cluster applications.

HPC has existed on the cloud since the early days of AWS. Among the first users of Amazon Elastic Compute Cloud (EC2) were researchers looking for scalable and cost-effective solutions to problems ranging from genome analysis in life sciences, to simulations in high-energy physics, and other computing problems requiring large numbers of CPU cores for short periods of time. Researchers quickly discovered that the features and capabilities of AWS were well suited for creating massively parallel grids of virtual CPUs, on demand. Stochastic simulations and other "pleasingly parallel" applications, from molecular modeling to Monte Carlo financial risk analysis, were particularly well suited to using Amazon EC2 Spot Instances, which allow users to bid on unused EC2 instance capacity at cost savings of up to 90 percent off the normal hourly on-demand price.

As the capabilities and performance of AWS have continued to advance, the types of HPC applications that are running on AWS have also evolved, with open source and commercial software applications being successfully deployed on AWS across industries, and across application categories.

In addition to the many public sector users of cloud for scalable HPC, commercial enterprises have also been increasing their use of cloud for HPC, augmenting or in some cases replacing, their legacy HPC infrastructures.

Pharmaceutical companies, for example, are taking advantage of scalability in the cloud to accelerate drug discovery by running large-scale computational chemistry applications. In the manufacturing domain, firms around the world are successfully deploying third-party and in-house developed applications for computer aided design (CAD), electronic design automation (EDA), 3D rendering, and parallel materials simulations. These firms routinely launch simulation clusters consisting of many thousands of CPU cores, for example to run thousands or even millions of parallel parametric sweeps.

In the financial services sector, organizations ranging from hedge funds, to global banks, to independent auditing agencies such as FINRA are using AWS to run complex financial simulations, to predict future outcomes, and to back-test proprietary trading algorithms.

Mapping HPC Applications to AWS Features

Amazon EC2 provides a wide selection of instance types optimized to fit different use cases. Instance types comprise varying combinations of CPU, memory, storage, and networking capacity and give you the flexibility to choose the appropriate mix of resources for specific HPC applications. AWS also offers a wide variety of data storage options, and higher-level capabilities for deployment, cluster automation, and workflow management. To better understand how these capabilities are used for HPC, we'll first discuss the broad categories of HPC applications.

Loosely Coupled Grid Computing

This category of HPC applications is sometimes characterized as high throughput computing (HTC). Examples include Monte Carlo simulations for financial risk analysis, materials science for proteomics, and a wide range of applications that can be distributed across very large numbers of CPU cores or nodes in a grid, with little dependence on high performance node-to-node interconnect, or on high performance storage.

These applications are often designed for fault-tolerance, meaning the application is tolerant of individual nodes being added or removed during the course of a run. Such applications are ideally suited to Amazon EC2 Spot Instances, and benefit as well from automation using Auto Scaling[4]. Customers with highly scalable applications can choose from many EC2 instance types[5]. They can optimize the choice of instance types for the specific compute tasks they plan to execute or for controlling total costs of completing a large set of batch tasks over time. Many applications in this category are able to take advantage of GPU acceleration, using Amazon EC2 G2 instances in combination with programming methods such as NVIDIA's CUDA parallel computing platform, or with OpenCL.

Tightly Coupled HPC

Applications in this category include many of the largest, most established HPC workloads: example workloads include weather modeling, electromagnetic simulations, and computational fluid dynamics. These applications are often written using the messaging passing interface (MPI) or shared memory programming models, using libraries such as MPITCH, OpenMP, or other methods for managing high levels of inter-node communications.

Tightly coupled applications can be deployed effectively on the cloud at small to medium scale, with a maximum number of cores per job being dependent on the application and its unique set of requirements, for example to meet the constraints of packet size, frequency, and latency sensitivity of node-to-node communications. A significant benefit of running such workloads on AWS is the ability to scale out to achieve a higher quality of results. For example, an engineer running electromagnetic simulations could run larger numbers of parametric sweeps than would otherwise be practical, by using very large numbers of Amazon EC2 On-Demand or Spot Instances, and using automation to launch independent and parallel simulation jobs. A further benefit for such an engineer is using Amazon Simple Storage Service (S3), Amazon DynamoDB, and other AWS capabilities to aggregate, analyze, and visualize the results.

Amazon EC2 capabilities that help with applications in this category include EC2 placement groups and enhanced networking[6], for reduced node-to-node latencies and consistent network performance, and the availability of GPU instance types, which can reduce the need to add more computing nodes by offloading highly parallel computations to the GPU.

Data-Intensive Computing

When grid and cluster HPC workloads such as those described earlier are combined with large amounts of data, the resulting applications require fast, reliable access to various types of data storage. Representative HPC applications in this category include genomics, high-resolution image processing, 3D animation rendering, mix-signal circuit simulation, seismic processing, and machine learning, among others.

Note that HPC in this category has similarities to “big data” but has different goals: big data is used to answer questions you didn’t know to ask, or it is used to discover correlations and patterns in large and diverse datasets. Examples of big data include website log analysis, financial fraud detection, consumer sentiment analysis, and ad placements.

HPC may also generate or consume very large amounts of data, but HPC applications most often operate on well-structured data models, for example a 3D mesh representing a complex physical shape, or the individual frames of an animated feature film. HPC applications use computing to calculate an answer to a known question, or to simulate a scenario based on a predefined model, using predefined sets of inputs. In the domain of semiconductor design, for example, digital and mixed-signal simulations are often run on large computing clusters, with many thousands of individual simulation tasks that all require access to high-performance shared storage. This pattern is also found in life sciences, in particular genomics workflows such as DNA and RNA sequence assembly and alignment.

AWS services and features that help HPC users optimize for data-intensive computing include Amazon S3, Amazon Elastic Block Store (EBS), and AmazonEC2 instance types such as the I2 instance type, which includes locally attached solid-state drive (SSD) storage. Solutions also exist for creating high performance virtual network attached storage (NAS) and network file systems (NFS) in the cloud, allowing applications running in Amazon EC2 to access high performance, scalable, cloud-based shared storage resources.

Factors that Make AWS Compelling for HPC

Scalability and Agility

AWS allows HPC users to scale applications horizontally and vertically to meet computing demands, eliminating the need for job queues and decreasing the time to results. Horizontal scalability is provided by the elasticity of Amazon EC2—additional compute nodes can be added as needed and in an automated manner. Vertical scalability is provided by the wide range of EC2 instance types, and through Amazon EC2 features such as placement groups and enhanced networking.

Automated methods of HPC deployment, including the CfnCluster framework[7] developed at AWS help customers get started quickly and benefit from scalability in the cloud.

Global Collaboration and Remote Visualization

HPC users deploying on AWS quickly find that that running workloads on the cloud is not simply a means to doing the same kinds of work as before, at lower cost. Instead, these customers are seeing that cloud enables a new way for globally distributed teams to securely collaborate on data, and to manage their non-HPC needs even more efficiently, including desktop technical applications.

Such collaboration in manufacturing, for example, can include using the cloud as a secure, globally accessible big data platform for production yield analysis, or enabling design collaboration using remote 3D graphics. The use of the cloud for collaboration and visualization allows a subcontractor or remote design team to view and interact with a simulation model in near real time, without the need to duplicate and proliferate sensitive design data.

Reducing or Eliminating Reliance on Job Queues

HPC users today are accustomed to using open source or commercial cluster and job management tools, including job schedulers. In a typical HPC environment, individual HPC users—researchers, engineers, and analysts who rely on HPC applications—will submit their jobs to a shared resource using a job queue submission system, using either the command line or an internal job submission portal. The submitted job typically includes a script that specifies the applications to be run and includes other information, such as whether and where data need to be pre-staged, the number of cores or threads to be allocated to the job, and possibly the maximum allowable runtime for the job. At this point, the cluster management software takes over, and it schedules the various incoming jobs, which may have different priorities, to the cluster resources.

Depending on the mix of jobs being submitted, their inter-dependencies and priorities, and whether they are optimized for the shared resource, the HPC grid or cluster may operate at very high or very low levels of effective utilization.

When workloads are highly variable (such as when there is a simultaneous high demand for simulations from many different groups, or when there are unexpected high-priority jobs being submitted), the queue wait times for a centrally managed physical cluster can grow dramatically, resulting in job completion times that are far in excess of the actual time needed to compete each job. Errors in the input scripts, mistakes in setting job parameters, or unanticipated runtimes can result in additional scheduling complexities and longer queue wait times due to queue contention.

When running HPC in the AWS Cloud, the problem of queue contention is eliminated, because every job or every set of related, interdependent jobs can be provided with its own purpose-built, on-demand cluster. In addition, the on-demand cluster can be customized for the unique set of applications for which it is being built. For example, you can configure a cluster with the right ratios of CPU cores, memory, and local storage. Using the AWS Cloud for HPC applications means there is less waste of resources and a more efficient use of HPC spending.

Faster Procurement and Provisioning

Rapid deployment of cloud-based, scalable computing and data storage is compelling for many organizations, in particular those seeking greater ability to innovate. HPC in the cloud removes the burden of IT procurement and setup from computational scientists and from commercial HPC users. The AWS Cloud allows these HPC users to select and deploy an optimal set of services for their unique applications, and to pay only for what they use.

The AWS Cloud can be deployed and managed by an individual HPC user, such as a geophysicist needing to validate a new seismic algorithm at scale using on-demand resources. Or the AWS Cloud can be deployed and managed by a corporate IT department, using procedures similar to those used for managing physical infrastructure. In both cases, a major benefit of using the AWS Cloud is the speed at which new infrastructure can be brought up and be ready for use, and the speed at which that same infrastructure can be reduced or eliminated to save costs. In both cases—scale-up and scale-down—you can commission and decommission HPC clusters in just minutes, rather than in days or weeks.

Sample Architectures

Grid Computing in the Cloud

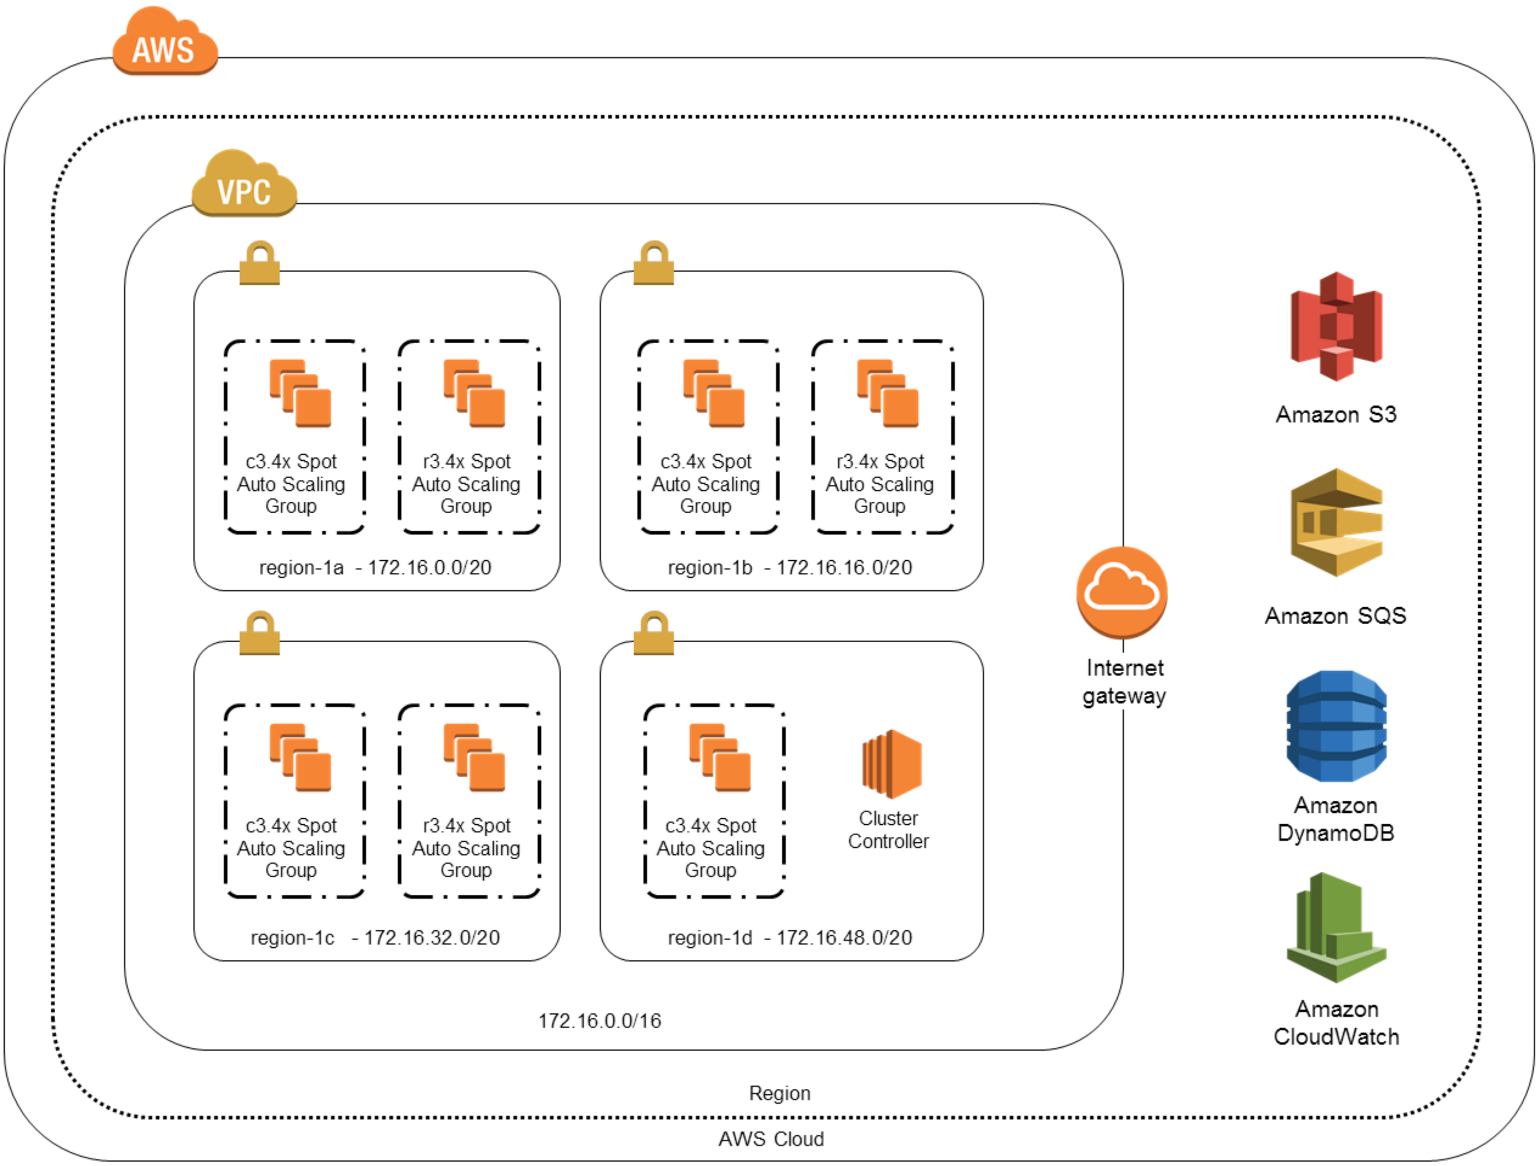

Figure 1: "Loosely-coupled" grid

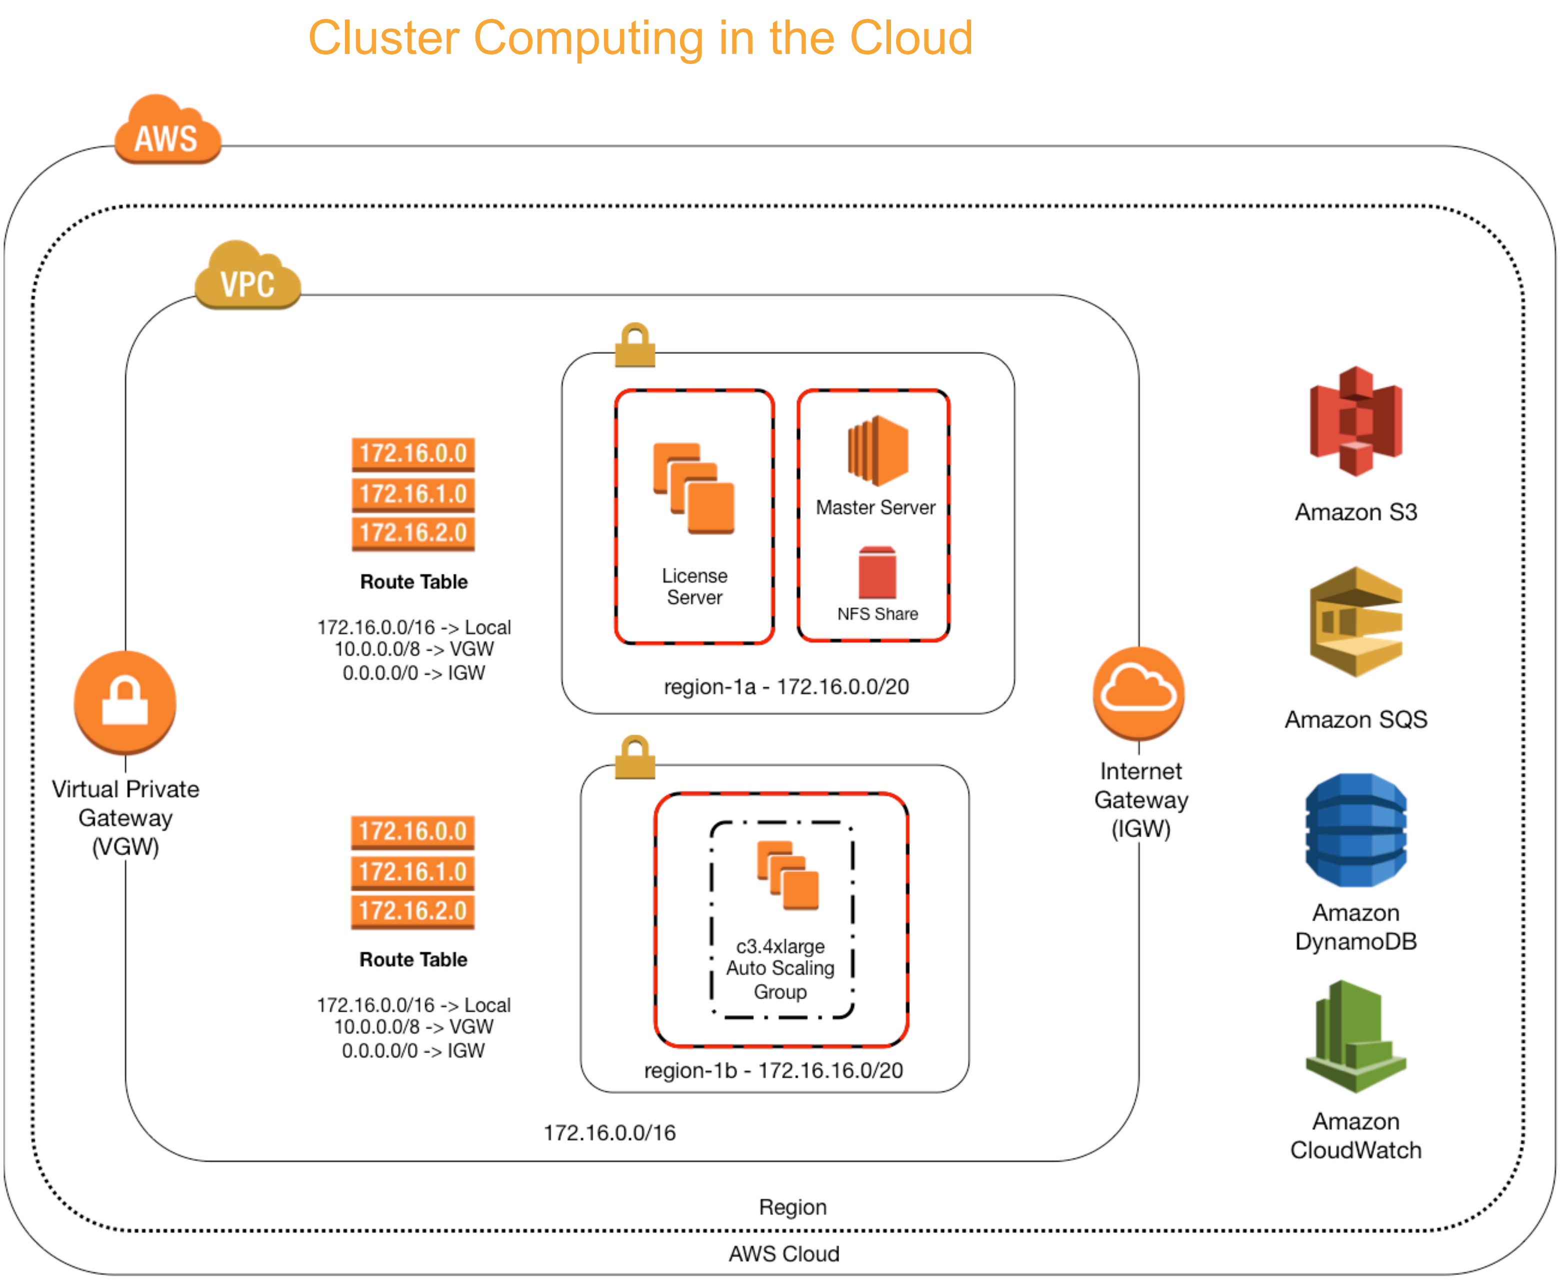

Figure 2: "Tightly coupled" cluster

Running Commercial HPC Applications on AWS

There are many independent software vendors (ISVs) providing innovative solutions for HPC applications. These ISVs include providers of computer-aided design (CAD), computer-aided engineering (CAE), electronic design automation (EDA), and other compute-intensive applications, as well as providers of HPC middleware, such as cluster management and job scheduling solutions. Providers of HPC-oriented remote visualization and remote desktop tools are also part of the HPC software ecosystem, as are providers of libraries and development software for parallel computing.

In most cases, these third-party software products can run on AWS with little or no change. By using the features of Amazon Virtual Private Cloud (VPC), HPC users and HPC support teams can ensure that licensed ISV software is being operated in a secure and auditable manner, including the use of license servers and associated logs.

In some cases, it will be necessary to discuss your proposed use of technical software with the ISV, to ensure compliance with license terms. AWS is available to help with such discussions, including providing ISVs with deployment assistance via the AWS Partner Network (APN).

In other cases, the ISV may have alternative distributions of software that are optimized for use on AWS, or can provide a more fully managed software-as-a-service (SaaS) alternative to customer-managed cloud deployments.

Security and Governance for HPC

The AWS Cloud infrastructure has been architected to be one of the most flexible and secured cloud computing environments available today. For HPC applications, AWS provides an extremely scalable, highly reliable, and secured platform for the most sensitive applications and data.

World-Class Protection

With the AWS Cloud, not only are infrastructure headaches removed, but so are many of the allied security issues.

The AWS virtual infrastructure has been designed to provide optimum availability while designed for customer privacy and segregation.

For a complete list of all the security measures built into the core AWS Cloud infrastructure, platforms, and services, please read our "Overview of Security Processes" whitepaper[8].

Built-In Security Features

Not only are your applications and data protected by highly secured facilities and infrastructure, they're also protected by extensive network and security monitoring systems. These systems provide basic but important security measures such as distributed denial of service (DDoS) protection and password brute-force detection on AWS Accounts. A discussion of additional security measures follows.

Secure Access

Customer access points, also called API endpoints, allow secure HTTP access (HTTPS) so that you can establish secure communication sessions with your AWS services using Secure Sockets Layer (SSL).

Built-In Firewalls

You can control how accessible your instances are by configuring built-in firewall rules—from totally public to completely private, or somewhere in between. When your instances reside within an Amazon Virtual Private Cloud (VPC) subnet, you can control egress as well as ingress.

Unique Users

AWS Identity and Access Management (IAM) allows you to control the level of access your own users have to your AWS infrastructure services. With IAM, each user can have unique security credentials, eliminating the need for shared passwords, or keys, and allowing the security best practices of role separation and least privilege.

Multi-Factor Authentication

AWS provides built-in support for multi-factor authentication (MFA) for use with AWS accounts as well as individual IAM user accounts.

Private Subnets

Amazon VPC allows you to add another layer of network security to your instances by creating private subnets and even adding an IPsec VPN tunnel between your home network and your VPC.

Encrypted Data Storage

Customers can have the data and objects they store in Amazon S3, Amazon Glacier, Amazon Redshift, and Amazon Relational Database Service (RDS) for Oracle encrypted automatically using Advanced Encryption Standard (AES) 256, a secure symmetric-key encryption standard using 256-bit encryption keys.

Direct Connection Option

The AWS Direct Connect service allows you to establish a dedicated network connection from your premises to AWS. Using industry standard 802.1q VLANs, this dedicated connection can be partitioned into multiple logical connections to enable you to access both public and private IP environments within your AWS Cloud.

Security Logs

AWS CloudTrail provides logs of all user activity within your AWS account. You can see who performed what actions on each of your AWS resources.

Isolated GovCloud

For customers who require additional measures in order to comply with US ITAR regulations, AWS provides an entirely separate region called AWS GovCloud (US) that provides an environment where customers can run ITAR-compliant applications, and provides special endpoints that utilize only FIPS 140-2 encryption.

AWS Cloud HSM

For customers who must use Hardware Security Module (HSM) appliances for cryptographic key storage, AWS CloudHSM provides a highly secure and convenient way to store and manage keys.

Trusted Advisor

Provided automatically when you sign up for AWS Premium Support, the AWS Trusted Advisor service is a convenient way for you to see where you could use a little more security. It monitors AWS resources and alerts you to security configuration gaps, such as overly permissive access to certain EC2 instance ports and Amazon S3 storage buckets, minimal use of role segregation using IAM, and weak password policies.

Because the AWS Cloud infrastructure provides so many built-in security features, you can simply focus on the security of your guest operating system (OS) and applications. AWS security engineers and solutions architects have developed whitepapers and operational checklists to help you select the best options for your needs, and they recommend security best practices, such as storing secret keys and passwords in a secure manner and rotating them frequently.

Conclusion

Cloud computing helps research and academic organizations, government agencies, and commercial HPC users gain fast access to grid and cluster computing resources, to achieve results faster and with higher quality, at a reduced cost relative to traditional HPC infrastructure. The AWS Cloud transforms previously complex and static HPC infrastructures into highly flexible and adaptable resources for on-demand or long-term use.

Contributors

The following individuals contributed to this document:

- David Pellerin, principal BDM (HPC), AWS Business Development
- Dougal Ballantyne, solutions architect, Amazon Web Services
- Adam Boeglin, solutions architect, Amazon Web Services

Further Reading

Get started with HPC in the cloud today by going to aws.amazon.com/hpc.

Notes

[1] http://aws.amazon.com/ec2/purchasing-options/spot-instances/

[2] http://setiathome.ssl.berkeley.edu/

[3] http://folding.stanford.edu/

[4] http://aws.amazon.com/documentation/autoscaling/

[5] http://aws.amazon.com/ec2/instance-types/

[6] http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/enhanced-networking.html

[7] http://aws.amazon.com/hpc/cfncluster

[8] http://d0.awsstatic.com/whitepapers/Security/AWS Security Whitepaper.pdf