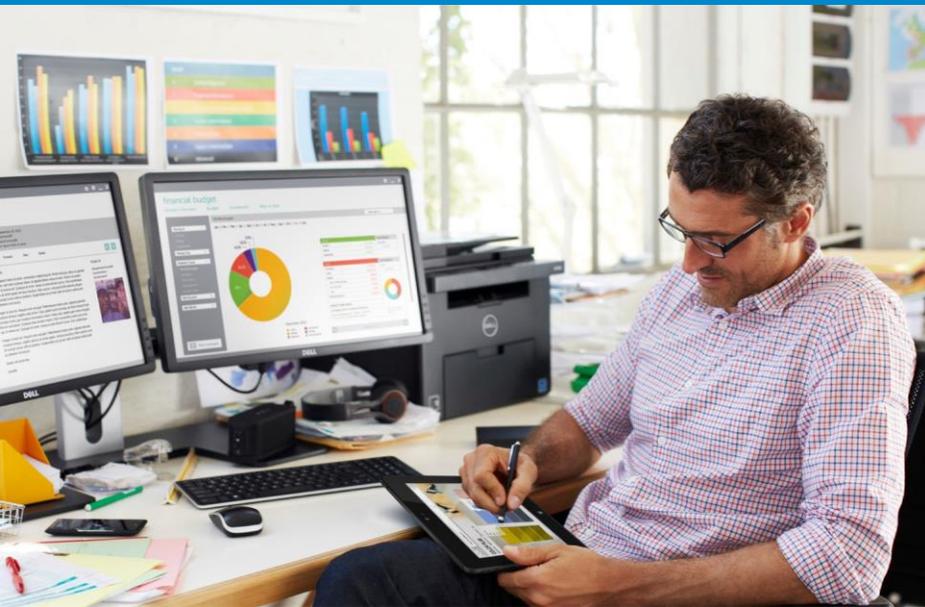
A man wearing glasses and a dark shirt is standing in a server room. He is holding a tablet computer in his left hand and pointing with his right hand towards a server rack. The server racks are filled with various hardware components, and the room is dimly lit with blue ambient lighting.

# A guide to high-performance computing (HPC) cloud-based solutions and big data analysis



“The HPC community has had a long-standing interest in creating scale-out environments for running throughput-oriented and parallel distributed workloads. Both large-scale environments (for example, cloud computing facilities) and scale-out workloads (such as big data) are becoming more important in the enterprise. In fact, with the rise of big data, the advent of affordable, powerful clusters and strategies that take advantage of commodity systems for scale-out applications, these days the enterprise computing environment is looking a lot like HPC.”

— Andrew Nelson, VMware at VMworld 2014\*



\* VMworld 2014: Virtualized HPC as a Service, Andrew Nelson, speaker. <http://www.youtube.com/watch?v=ssA0dZehULU>.

# Dell Cloud Solutions harness big data so many more questions are made answerable for businesses and researchers

HPC is helping to design, test and fly planes before they're physically built. We're helping to simulate dangerous weather patterns and save lives by creating early warning systems. We're even helping make virtual crash tests for automobiles literally a matter of a wall of monitors and computer algorithms.

Today, high-performance computing (HPC) has become the backbone of science and product development. This breakthrough computing power is an absolutely irreplaceable tool for scientists, academics and cutting-edge researchers. All over the world, IT departments are using Dell Services cloud-based HPC solutions to provide maximum computational power, while employing standards-based technologies that have been tested and validated.

All of which helps today's most visionary, scientific researchers do what they do best: unearth the discoveries that fundamentally change the world.

## Making HPC available to all who need it

HPC used to be the domain of specialists using expensive, proprietary equipment. Now commodity-based infrastructure and cloud-based solutions offer affordable

HPC tools and strategic advantages to any research organization or line of business.

## Dell Services approach to high-performance computing

Dell recognizes that no two customers have exactly the same needs and requirements. Our customized approach to HPC focuses on three areas to achieve faster results by enabling discovery and driving the creation of intellectual property with bare metal. Bare metal cloud-based configurations, virtualized cloud-based solutions using AWS or the use of Google Cloud Platform. We add value with highly flexible terms, support and Dell Cloud Manager for Google and Amazon Web Services (AWS) for hybrid cloud management.

## Dell Cloud Manager

Dell Cloud Manager is cloud infrastructure management software that helps you manage the consumption of your cloud services, while deploying and managing enterprise applications in the cloud. Cloud Manager connects to your existing cloud accounts and provides a unified central management console from which you can control all actions taken within the cloud.

Cloud Manager provides secure management of your cloud resources, including granular role-based access controls and budget controls for all resources, with the ability to set quotas to ensure budgets are capped. It's easy to extend existing governance policies to the cloud —including integration with Lightweight Directory Access Protocol (LDAP), Active Directory and Security Assertion Markup Language (SAML) — via our application programming interface (API), key management and encryption. And centralized management of resources across all clouds or cloud accounts makes monitoring and reporting much easier than tracking through disparate systems.

## Security

Dell Cloud-based HPC and big data analysts are secure. We work with each client to address these issues to maintain the security of their data.

## Secure Data Migration

R Systems can load data from disc or tape. Peering with Google provides an excellent way to move large amounts of data globally over Google's dark fiber backbone. Data migration to AWS can be accomplished with different methodologies.

# Dell Cloud-Based HPC questions answered

Issues	Perception	Dell Cloud-based HPC
Seasonal business overtaxes compute resources	<ul style="list-style-type: none"> <li>Poor customer experience due to slow website response times</li> <li>Loss of business to competitors</li> <li>System downtime stresses IT personnel due to overburdened server workloads</li> </ul>	<ul style="list-style-type: none"> <li>Rapidly scale-up or down capacity</li> <li>Cut application related operational costs</li> <li>Provide bursting to an alternate data center for additional capacity, when needed</li> </ul>
Lengthy time to provision IT resources	<ul style="list-style-type: none"> <li>Slow time to market for applications reduces competitive advantage</li> <li>Perceived inefficiencies of IT personnel by the business</li> <li>High infrastructure costs charged back to the business or directly impacting CIO budget</li> </ul>	<ul style="list-style-type: none"> <li>Automate key deployment tasks</li> <li>Enable self-service deployment of servers by the business end users</li> <li>Cut application provisioning times</li> <li>Shift cost/benefit decisions to the business</li> </ul>
Poor utilization of virtualized servers	<ul style="list-style-type: none"> <li>Sub-optimal price/performance of compute</li> <li>Inflexible infrastructure required over-images to maintain consistent</li> <li>Feature limitations prohibit dynamic reallocation</li> </ul>	<ul style="list-style-type: none"> <li>Minimize upfront investment requirements and risk resources</li> <li>Provide management and monitoring tools that architecture automatically reassign server provisioning performance and utilization</li> </ul>
IT CapEx budget reductions or freezes	<ul style="list-style-type: none"> <li>Inability to procure new servers and storage to meet computing demand</li> <li>Lack of funds to investigate emerging technologies which could lend a competitive advantage</li> </ul>	<ul style="list-style-type: none"> <li>Help transition your IT costs from CapEx to an OpEx growing model</li> <li>Help you leverage public cloud scale to release data center capacity and/or facilitate disaster recovery</li> <li>Enable a shift of resources from lights-on to strategic</li> <li>Drive your green agenda for energy efficiency in the data center</li> </ul>
Intense compute resources	<ul style="list-style-type: none"> <li>Powerful compute resources expensive, require large upfront capital high maintenance</li> </ul>	<ul style="list-style-type: none"> <li>Enjoy the benefits of powerful are fast computing with no expenses</li> <li>Focus on running your business, not maintaining hardware</li> </ul>

# Dell Services bare metal high-performance computing with R Systems



Dell HPC Cloud with R Systems Service is designed to provide a secure private cloud environment, hosted and managed by Dell and R Systems, an HPC Solutions Partner. HPC infrastructure consists of Dell processing, power, storage and memory capacity. R Systems provides white-glove HPC services with custom solutions in your choice of locations, including their company-owned data centers in Champaign, Illinois located at the University of Illinois' Research Park. R Systems offers complete Dell hardware-based systems, as well as custom engagements/configurations based on specific business needs. Dell bare metal service is only available in North America.

## The scope of this service

If you're a research or commercial organization that wants to host sensitive data or high-performance applications with a cloud provider, you need architectural options that are not available from standard public or multi-tenant cloud services. Our service can offer greater flexibility than public cloud by enhancing your ability to operate production workloads that might otherwise be limited to custom, on-site operations in a hosted private cloud.

Services include:

- Customizable burst or short-term HPC capacity served as an operating expense in 24 hour increments

- Large-scale core count clusters with the option of Infiniband interconnectivity
- High memory configurations available for memory-intensive workloads
- User support for all knowledge levels (from novice to experienced)
- Customizable security for public or private cloud
- Easy to set up access
- Available for Linux and Windows
- In-depth experience with a wide variety of independent software vendors (ISVs) such as ANSYS, CD-adapco, Milliman and Simulia
- Pool resources
- Share unused cycles on separate clusters
- Have a single interface for multiple clusters

The benefits:

- Lower technical barrier to HPC via burst models
- Empower researchers and increase collaboration
- Increase institutional utilization of cloud resources
- Pay only for the resources you use through on-demand cloud computing
- Reduce capital outlay and total cost of ownership (TCO)
- Take advantage of the agility and efficiencies of cloud
- Experience superior user support with R Systems' highly trained and experienced technicians.

The services

Build a private cloud from existing individual clusters if you want to:

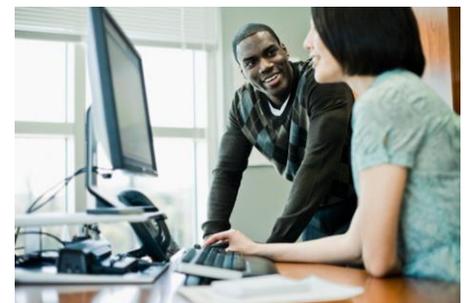
The bursting ability

Public cloud resources can help you:

- Service intermittent processing spikes that surpass local capacity
- Use multiple operating systems and applications

The hosting and management

- Domain specific (for example, a genomics cloud) for researchers who:
- Want their data, their collaborators' data and reference data in one location for processing
- Want access to high-performance computing



To learn more about R Systems and the "white glove" approach to HPC they offer, visit [Dell bare metal hpc solutions](#) on [youtube.com](#)

# Google Cloud Platform

Google Cloud Platform harnesses the 14 years of experience Google has building out the world's fastest, most powerful, highest quality cloud infrastructure on the planet. Google powers Google — and that means some of the largest, highest availability software products ever created. The Cloud Platform brings that power to you.

When you are running on Cloud Platform, you're using the same infrastructure Google uses to run their critical tasks such as searches, Docs or Gmail.

For those unfamiliar with the offering, it's very straightforward. It is offered a variety of ways to perform compute functions at scale, store data and build software better. It has the capability of going from data to insight in milliseconds.

Google Cloud Platform is composed of a family of products, each including a web interface, a command-line tool and a representational state transfer (REST) application programming interface (API).

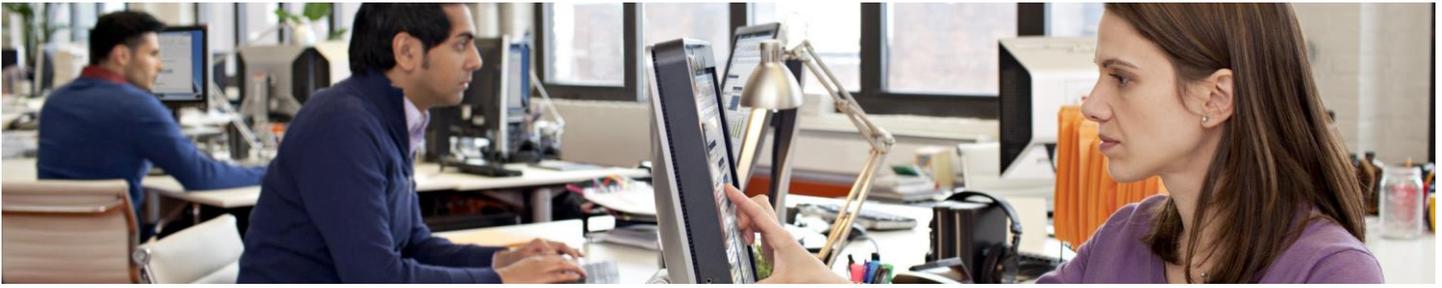
- **Google App Engine** is a **platform as a service** for sandboxed web applications. App Engine offers automatic scaling with resources increased automatically to handle server load.
- **Google Compute Engine** is the **infrastructure-as-a-service** component of Google Cloud Platform that enables users to launch virtual machines (VMs) on demand.
- **Google Cloud Storage** is an online storage service for files.
- **Google Cloud Datastore** is a fully managed, highly available NoSQL data storage for non-relational data that includes a REST API.
- **Google Cloud SQL** is a fully managed MySQL database that lives in the Google Cloud infrastructure.
- **Google BigQuery** is a data analysis tool that uses SQL-like queries to process big datasets in seconds.
- **Google Cloud Endpoints** is a tool to create services inside App Engine that can be easily connected from iOS, Android and JavaScript clients.
- **Google Cloud DNS** is a domain name system (DNS) service hosted in the Google Cloud infrastructure.

For instance, Google Genomics provides an API to store, process, explore and share DNA sequence reads, reference-based alignments and variant calls using Google's cloud infrastructure.

Google Genomics helps researchers:

- Store alignments and variant calls for one genome or a million.
- Process genomic data in batch by running principal component analysis or Hardy-Weinberg equilibrium, in minutes or hours, by using parallel computing frameworks like MapReduce.
- Explore data by slicing alignments and variants by genomic range across one or multiple samples — for your own algorithms or for visualization. You can also interactively process entire cohorts to find transition/transversion ratios, allelic frequency, genome-wide association and more using BigQuery.
- Share genomic data with your research group, collaborators, the broader community or the public. You decide.





Google Genomics is implementing the API defined by the [Global Alliance for Genomics and Health](#) for visualization, analysis and more. Compliant software can access Google Genomics, local servers or any other implementation. From a business perspective, it attacks big data making it easy to find meaning from a data warehouse or wherever the data resides.

Google Cloud Platform is a pay-as-you-go service. BigQuery is billed by project making it easy to charge back lines of business.

Google provides a [pricing calculator](#) that will help you get a better idea of the cost of using Google Cloud Platform. The use of Dell Cloud Manager is required with Google Cloud Manager. An in-depth presentation on Google Cloud Manager and Dell is available. C3 instances provide you with the

highest performance processors and the lowest price/compute performance compared to all other Amazon EC2 instances. I2 instances are engineered for extremely I/O intensive workloads. The largest I2 instance type can deliver over 365,000 random reads per second and over 315,000 random writes per second. AWS recommends using C3 instances for compute-intensive applications and the I2 instances for transactional systems.

## Google Cloud Platform



Compute solutions

- Compute Engine: Run large-scale workloads on virtual machines, using infrastructure as a service.
- App Engine: Develop your applications easily using built-in services with our platform-as-a-service tool.



Storage solutions

- Cloud SQL: Store and manage data using a fully managed, relational MySQL database.
- Cloud Storage: Use a durable and highly available object storage service.
- Cloud Datastore: Provides a managed, NoSQL, schema-less database for storing non-relational data.



Big data solutions

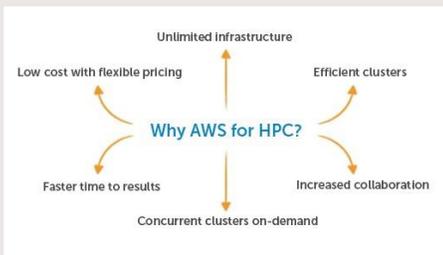
- BigQuery: Analyze big data in the cloud. Run fast, SQL queries against multi-terabyte datasets in seconds.



Developer tools

- Cloud SDK: Contains tools and libraries, allowing you to create and manage resources on our cloud platform.
- Push to Deploy: Use the Git version control system to automatically deploy your application to our App Engine.
- Cloud Playground: Run services like App Engine, Cloud Storage and Cloud SQL from your browser.
- Android Studio: Add our cloud platform as a back end to your application—right from the Android Studio IDE.
- Plugin for Eclipse: Provides tooling, API support and easy deployment for App Engine Java developers.

# Amazon Web Services High Performance Computing



Amazon Web Services offers a solution to solve complex science, engineering and business problems using applications that require high bandwidth, low latency networking and very high compute capabilities.

**TOP500: 76th fastest supercomputer on-demand**

Hun 2014 Top 500 list  
484.2 Flop/s  
26,496 cores in a cluster of EC2 C3 instances  
Intel Xeon E5-2680v2  
10C  
2.800 GHz processors  
LinPack Benchmark

AWS increases the speed of research by running high-performance computing in the cloud and reduces costs by providing cluster compute or cluster graphics processor unit (GPU) servers on demand and without large capital investments.

Access to a full-bisection, high-bandwidth network for tightly-coupled, I/O-intensive workloads enables you to scale out across thousands of cores for throughput-oriented applications.

**Compute Services**

**Elastic Compute Cloud (EC2)**

Basic unit of compute capacity, virtual machines ranges of CPU, memory & local disk options choice of instance types, from micro to cluster compute

c3.xlarge  
g2.medium  
m3.large

High Performance Computing (HPC) on Amazon Elastic Compute Cloud (EC2) is enabled by the cluster compute-optimized and GPU instance types (virtual machines). You can use them just like other EC2 instances, but they also have been specifically engineered to provide high-performance networking and you can scale to tens of thousands of instances on-demand.



GPU instances are ideally suited for:

- Graphical high-performance computing applications including 3D modeling and simulation
- Compute workloads including computational chemistry, financial analysis and engineering design
- General purpose GPU computing

AWS GPU instances provide access to NVIDIA GPUs, each with up to 1,536 Compute Unified Device Architecture (CUDA) cores and 4GB of video memory. With the latest driver releases, these GPUs

provide support for OpenGL, DirectX, CUDA, OpenCL and the GRID SDK.

Cluster instances can be launched within a placement group. All instances launched within a placement group have low latency, full bisection and 10Gbps of bandwidth between instances. Like many other Amazon EC2 resources, placement groups are dynamic and are elastically scalable as needed. You can also connect multiple placement groups to create very large high-performance computing clusters for massively parallel processing.

AWS currently supports enhanced networking capabilities using single root I/O virtualization (SR-IOV) for the C3 and I2 instance types. SR-IOV is a method of device virtualization that provides higher I/O performance and lower CPU utilization compared to traditional implementations. For supported Amazon EC2 instances, this feature provides higher packet per second (PPS) performance, lower inter-instance latencies and very low network jitter.

AWS offers many pricing models to support different workloads including free tier, on demand, reserved, spot and dedicated use. Find out more information on AWS Cloud-based HPC. [Click here](#) for a complete introduction to AWS

Many pricing models to support different workloads

#### On-Demand

Pay for compute capacity by the hour with no long-term commitments

For spiky workloads or to define needs



#### Reserved

Make a low, one-time payment and receive a significant discount on the hourly charges

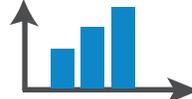
For committed utilization



#### Spot

Bid for unused capacity, charged at a spot price which fluctuates based on supply and demand

For time-insensitive or transient workloads



#### Dedicated

Launch instances within Amazon VPC that run on hardware dedicated to a single customer

For highly sensitive or compliance related workloads



# Dell partner guarantee

Dell recognizes that cloud pricing can be volatile, and if any element of the service provided by a partner is reduced in cost, the cost savings is immediately reflected in your bill.

Regardless of which partner or partners are used, a single invoice for services comes from Dell on a net 30-day basis.

Providing the broadest cloud ecosystem in the industry



Global coverage



Multiple cloud platforms  
Proprietary | Microsoft | VMware  
| OpenStack | CloudStack



Most popular use cases  
Vertical apps | Disaster recovery  
| ERP and CRM | Scale-out web  
infrastructure | Test and  
development | Productivity apps  
| Computational analysis



Dell Cloud Onsite and Dedicated Services



CenturyLink™



HOSTING EVOLVED



Six Degrees Group



For more information about any of our service offerings, please visit [Dell.com/services](http://Dell.com/services) or contact your Dell representative.



Scan or click this code to learn how Dell Services can help your organization.

Product and service availability varies by country. To learn more, customers and Dell Channel Partners should contact their sales representative for more information. Specifications are correct at date of publication but are subject to availability or change without notice at any time. Dell and its affiliates cannot be responsible for errors or omissions in typography or photography. Dell's Terms and Conditions of Sales and Service apply and are available on request. Dell and the Dell logo are trademarks of Dell Inc. Other trademarks and trade names may be used in this document to refer to either the entities claiming the marks and names or their products. Dell disclaims proprietary interest in the marks and names of others.

© 2014 Dell Inc. All rights reserved. October 2014 | D468-A Guide to HPC Cloud based solutions.indd | Rev. 1.0

