

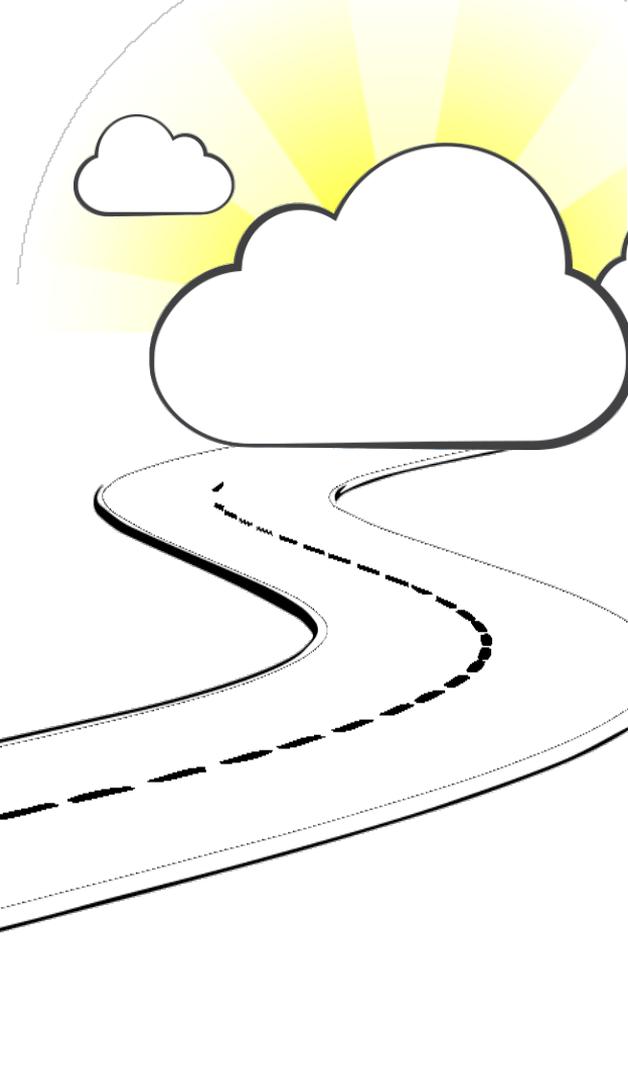


HPC Current Use and Practice: Industry

Cloud Use-Cases, Architectures, and Security

David Pellerin, Principal HPC Business Development

September, 2016

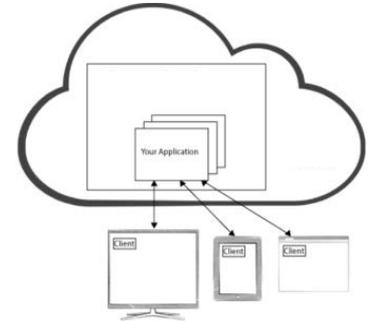


Motivations for Cloud in HPC

Scalability and Agility



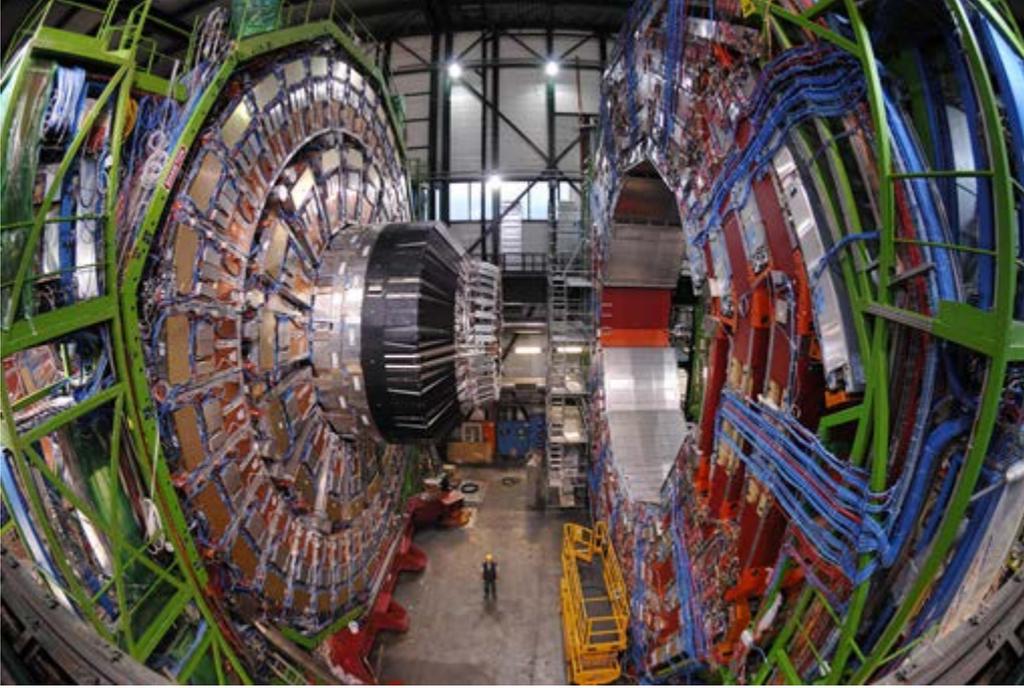
Secure Global Collaboration



Big Data Convergence with HPC



Scale Matters: for Big Compute and Big Data



Finding Patterns in the Data

This is

Big Data

TECH TIMES

PERSONAL TECH BIZ TECH FUTURE TECH

Scientists; Evidence Of New, Unknown Particle?

By [Jim Algar](#), Tech Times | January 12, 9:52 AM

[Like](#) [Follow](#) [Share](#) [Tweet](#) [Reddit](#) [1 Comment](#) [...](#)

[SUBSCRIBE](#)



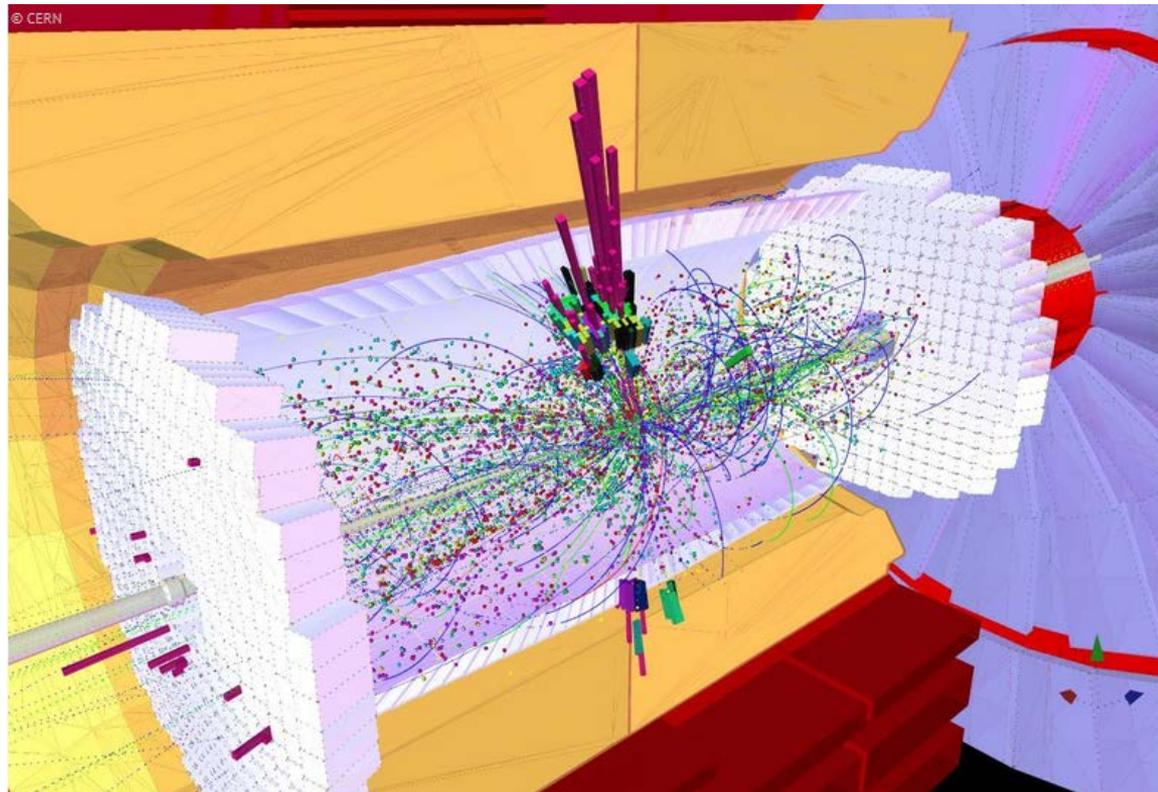
An unexpected "bump" in the data coming out of collision experiments with the Large Hadron Collider in Switzerland has scientists wondering if they've witnessed evidence of previously-unknown subatomic particles.

The collision, which cannot be explained by the Standard Model of physics, may have been evidence of a previously-undiscovered particle, or maybe even two particles, researchers say.

An anomaly in data from particle collisions in the Large Hadron Collider has researchers scratching their heads. Is this evidence of new particles that could turn the Standard Model of Physics on its head?
(Photo : Getty Images)

Constructing Models, Running Simulations, Viewing Results

This is
Big Compute



High Throughput Computing at Scale



HEP Cloud Summary



Zoom Out

Jan 21, 2016 14:49:25 to Feb 10, 2016 00:21:48



AWS VM Status

GCloud VM Status

HEP Cloud HTCondor Status

HEP Cloud Slots

HEP Cloud Slots (old)

Job Status



Slots Summary



Innovation in Energy Management

FUGRO ROAMES



[About Us](#) [Services](#) [Case Studies](#) [Media](#) [Events](#) [Library](#) [Contact Us](#)



About Us

Providing a virtual world environment uniquely tailored for power distribution management, Roames' asset network models facilitate comprehensive vegetation management, infrastructure condition evaluation and enhanced performance monitoring - reducing costs and resources.

[About Us](#)

Services

- 3D Virtual World
- Vegetation Management
- Conductor Clearance
- Asset Condition Assessment

[Services](#)

News

- 15 Oct Fugro Awarded National Grid Framework Contract in the UK
- 8 Jul Fugro Roames Aids Network Providers in Meeting AER Requirements
- 2 Jul Award Winning Roames Service Gains Momentum in UK

[News](#)

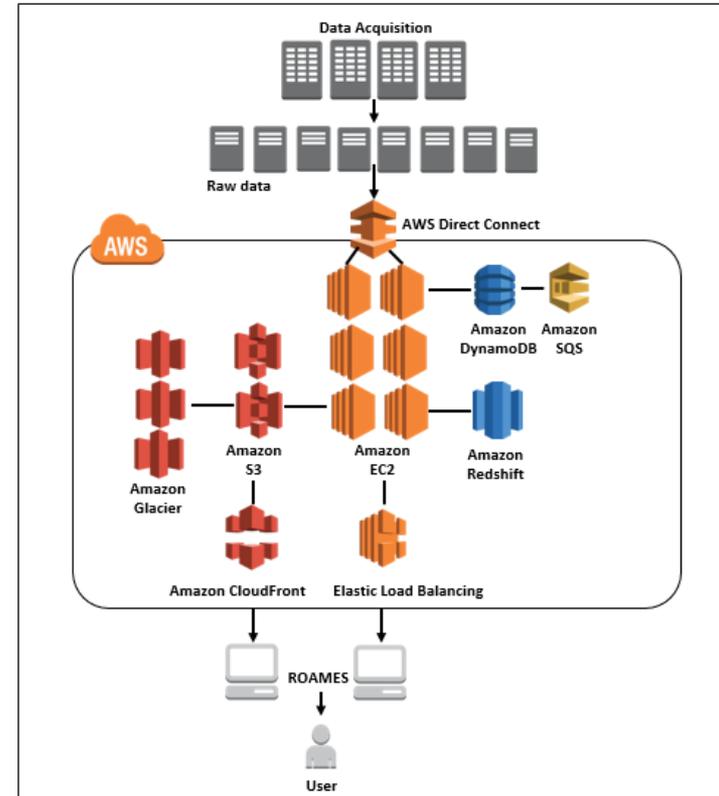


Big Data meets Big Compute

- Aircraft equipped with cameras, laser sensors
- Repeated overflights of power networks
- Captured data is used to render detailed 3D models of the power lines, and the environment
- Analytics and simulations are run to generate actionable reports
- Goal: directing post-disaster repair and prioritizing ongoing maintenance

"Fugro Roames has enabled Ergon Energy to reduce the cost of vegetation management from AU\$100 million to AU\$60 million per year."

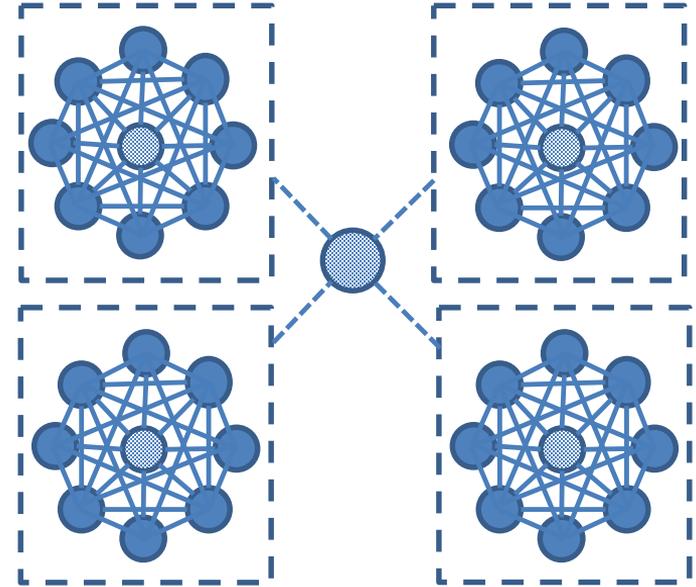
- Josh Passenger, Technical Architect, Fugro Roames



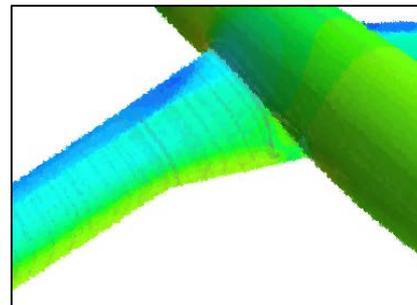
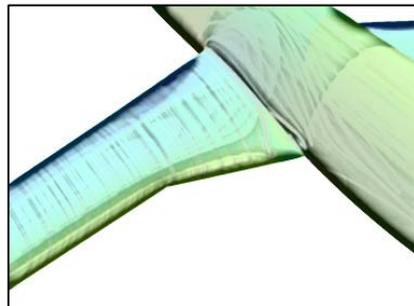
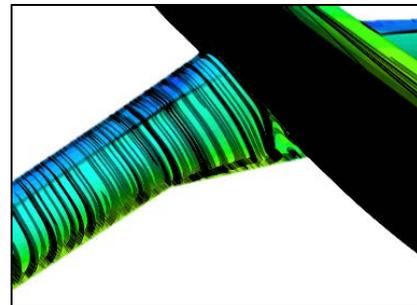
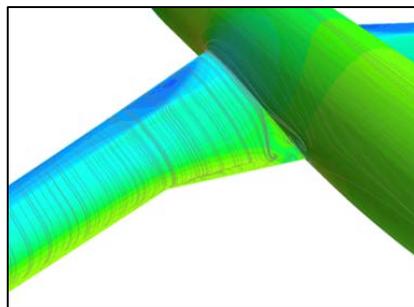
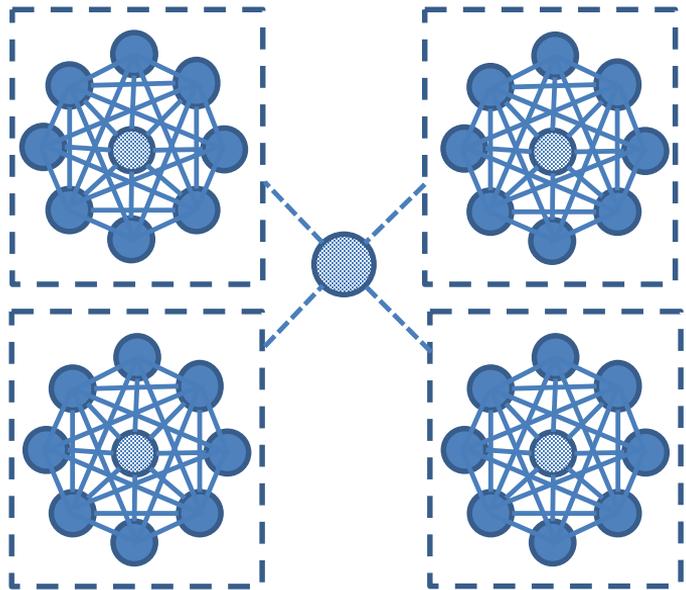
Primary Goal: Improve Throughput



- Run many jobs in parallel for large-scale parameter studies
- Reduce or eliminate HPC resource contention and queue wait times
- Right-size clusters and computing resources for each workload
- Optimize resources for each application, without procuring new hardware



What Does This Mean for Industry?



Go Wide: Expand the Simulation Domain

Run larger numbers of parallel, clustered HPC jobs

The High Cost of Queues for On-Premise HPC



Conflicting goals

- HPC users seek fastest possible time-to-results and have varying needs for scale
- IT support team seeks highest possible utilization of a fixed-size cluster

Result:

- The job queue becomes the capacity buffer
- Job completion times are hard to predict
- Users are frustrated and run fewer jobs
- **Innovation is throttled by IT resources**

HPC Innovation in Product Engineering



HGST applications for engineering:

- ✓ Molecular dynamics, CAD, CFD, EDA
- ✓ Collaboration tools for engineering
- ✓ **Big data for manufacturing yield analysis**



Running drive-head simulations at scale:

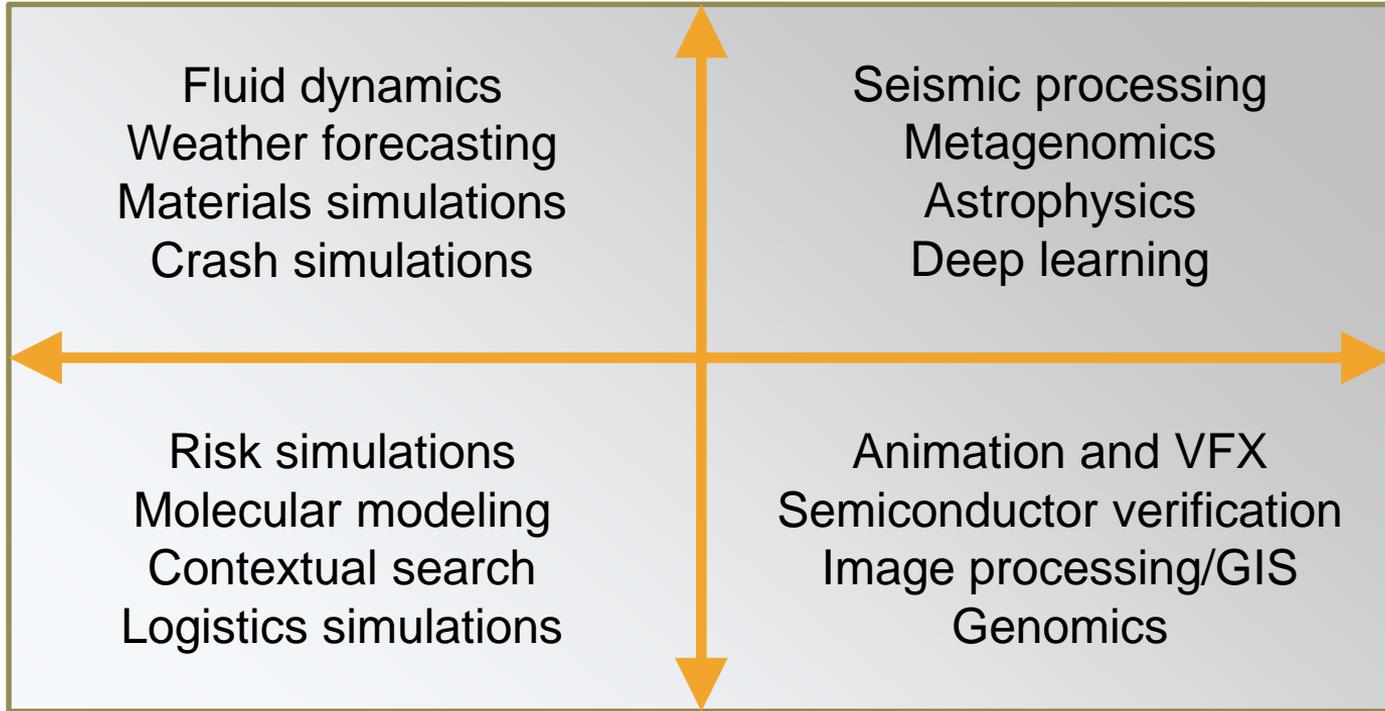
Millions of parallel parameter sweeps, running months of simulations in just hours.

Over 85,000 Intel cores running at peak, using EC2 Spot instances

Mapping HPC Use-Cases: Examples



Clustered (Tightly Coupled)

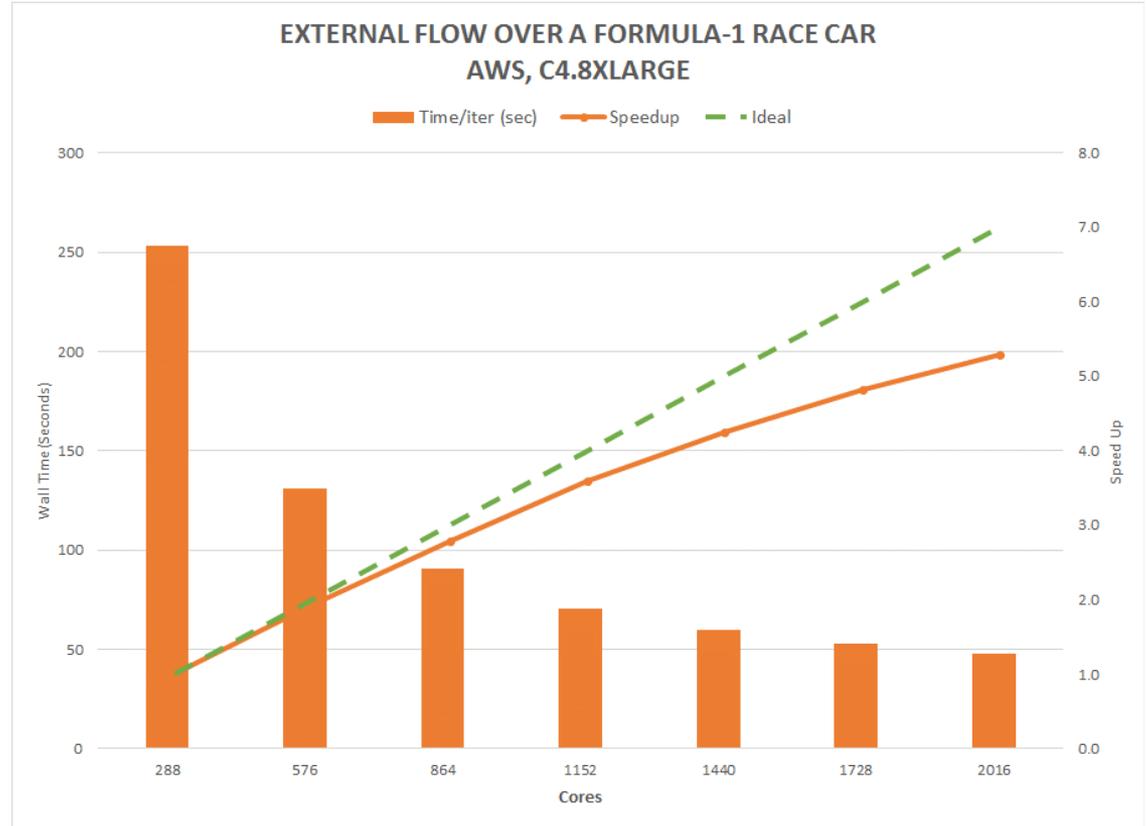
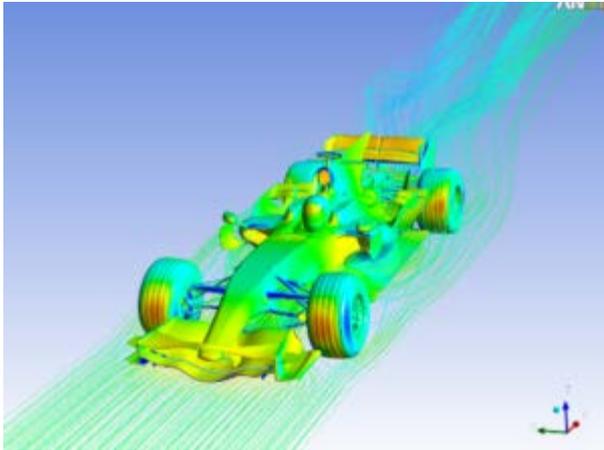


Distributed/Grid (Loosely Coupled)

Cloud Performance Testing: CFD on AWS

ANSYS Fluent

- AWS c4.8xlarge
- 140M cells
- F1 car CFD benchmark



<http://www.ansys-blog.com/simulation-on-the-cloud/>

Fluid dynamics on AWS

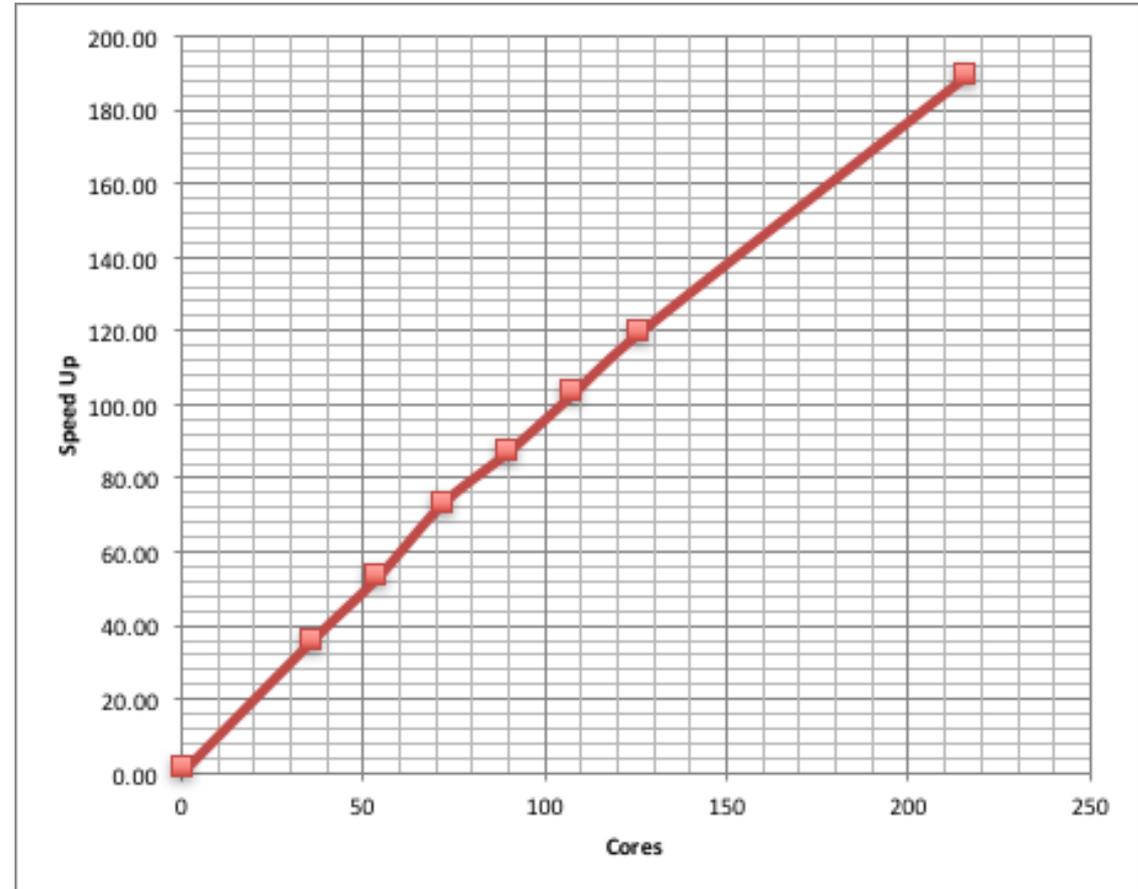


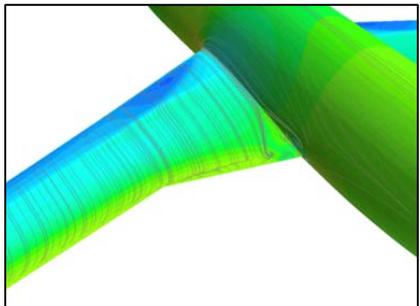
Metacomp CFD++

Fluid Dynamics and other tightly-coupled simulations can scale effectively on AWS

“Metacomp Technologies is very familiar with AWS HPC architecture and is well able to support CFD++ on it with confidence”

Dr. Sukumar Chakravarthy,
President and Founder, Metacomp

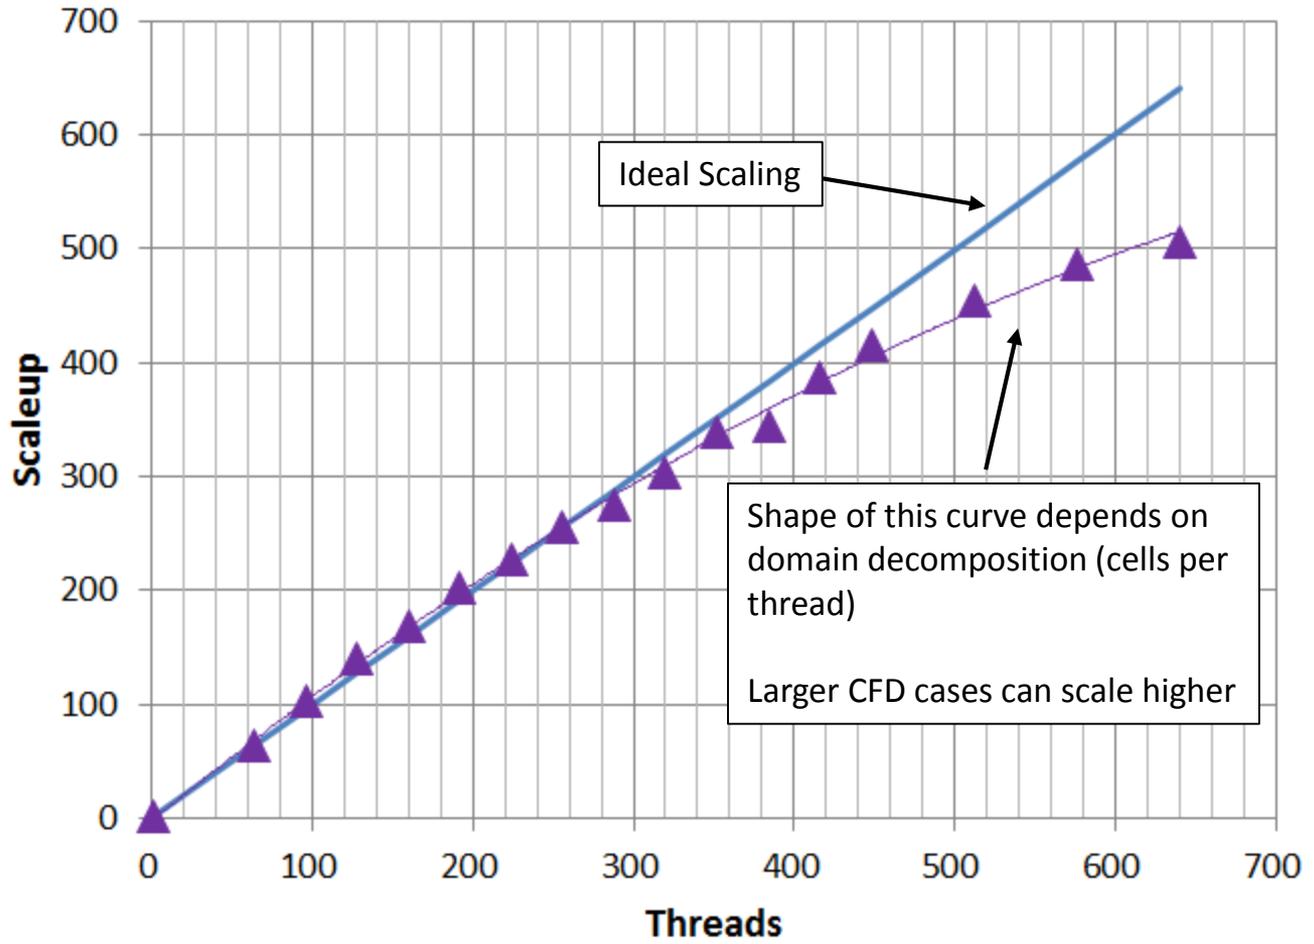




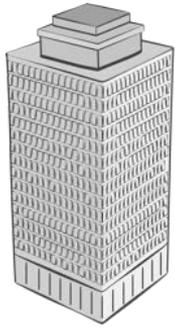
**16M cell, polyhedral,
external aero case,
STAR-CCM+**

**Running on threads,
c4.8xlarge instances**

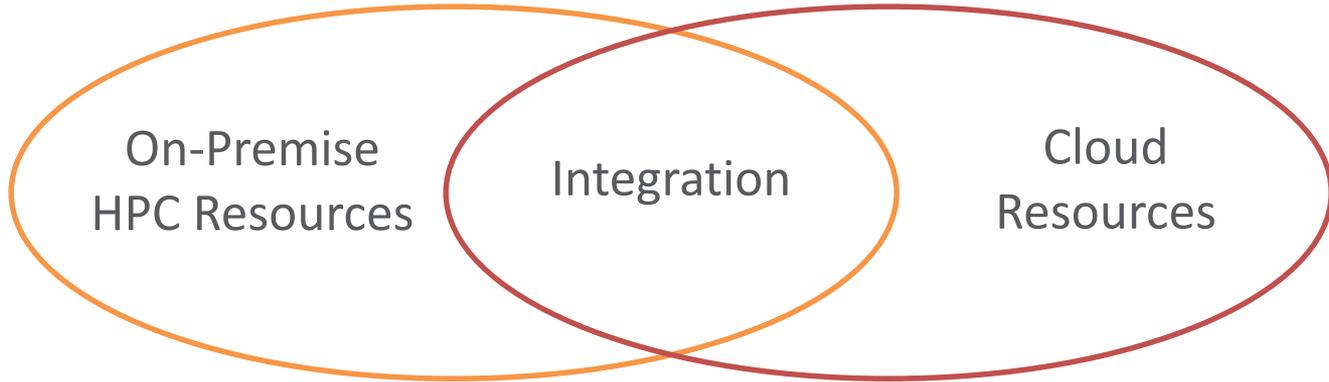
**Demonstrates excellent
scalability for typical
CFD models**



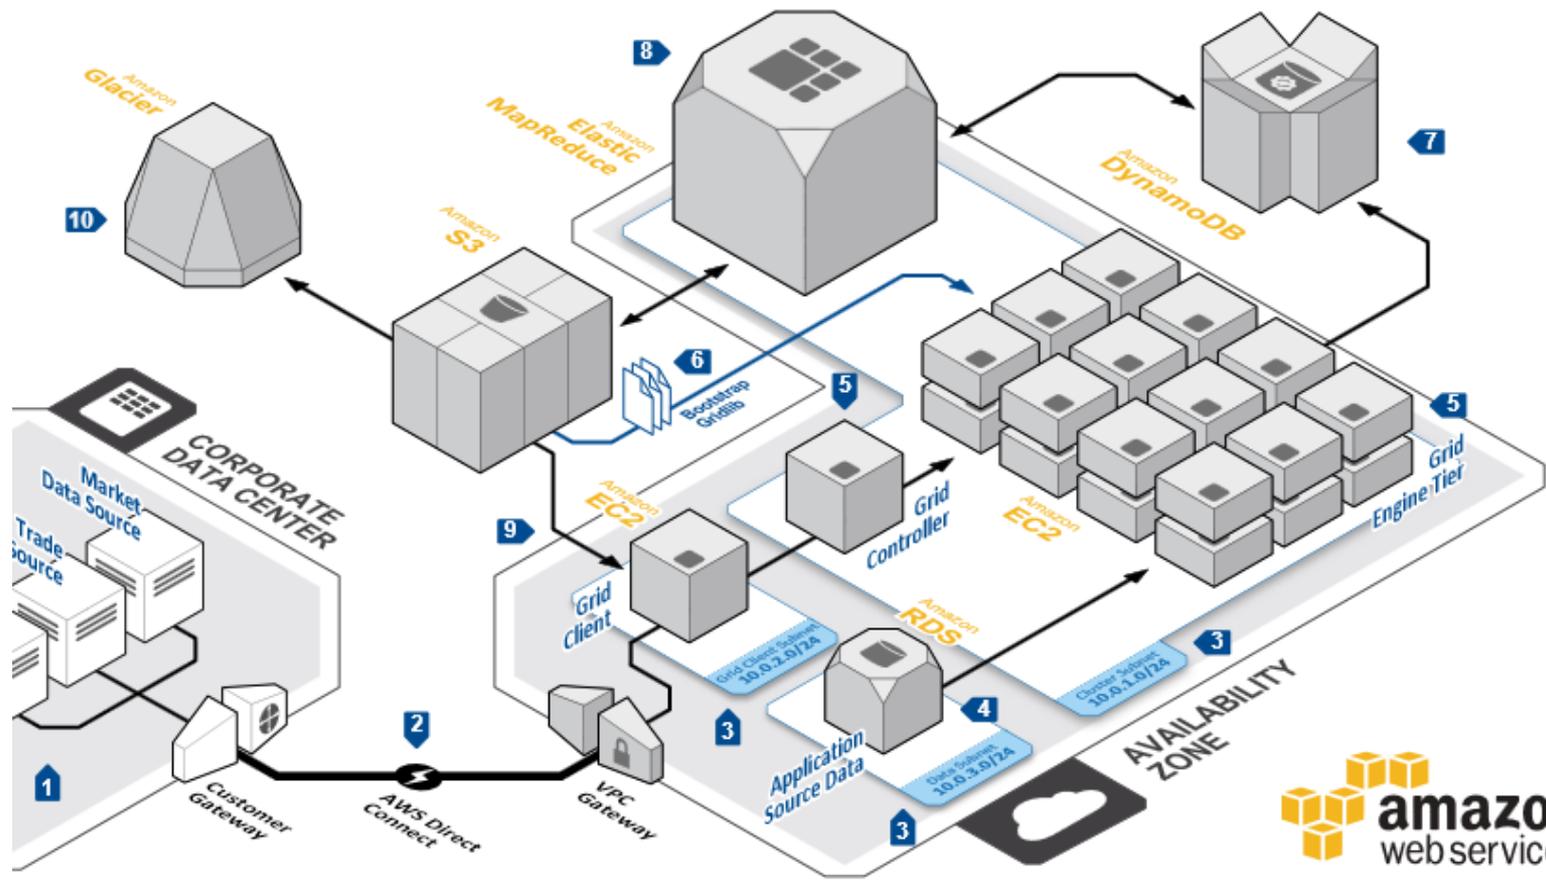
Hybrid HPC is a Common Pattern



**Corporate
Data Centers**

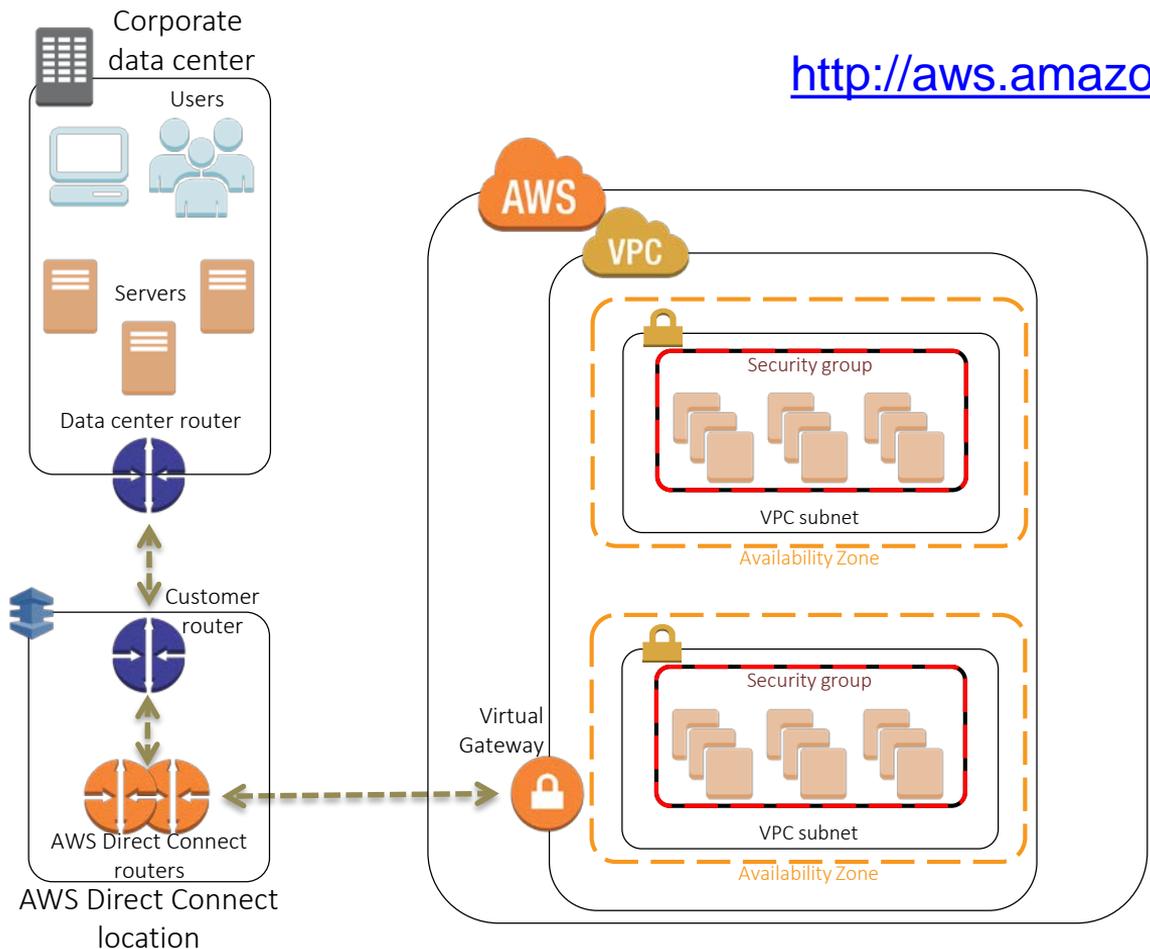


Sample Hybrid HPC Workflow with Auto Scaling



AWS Direct Connect and VPC

<http://aws.amazon.com/directconnect/>

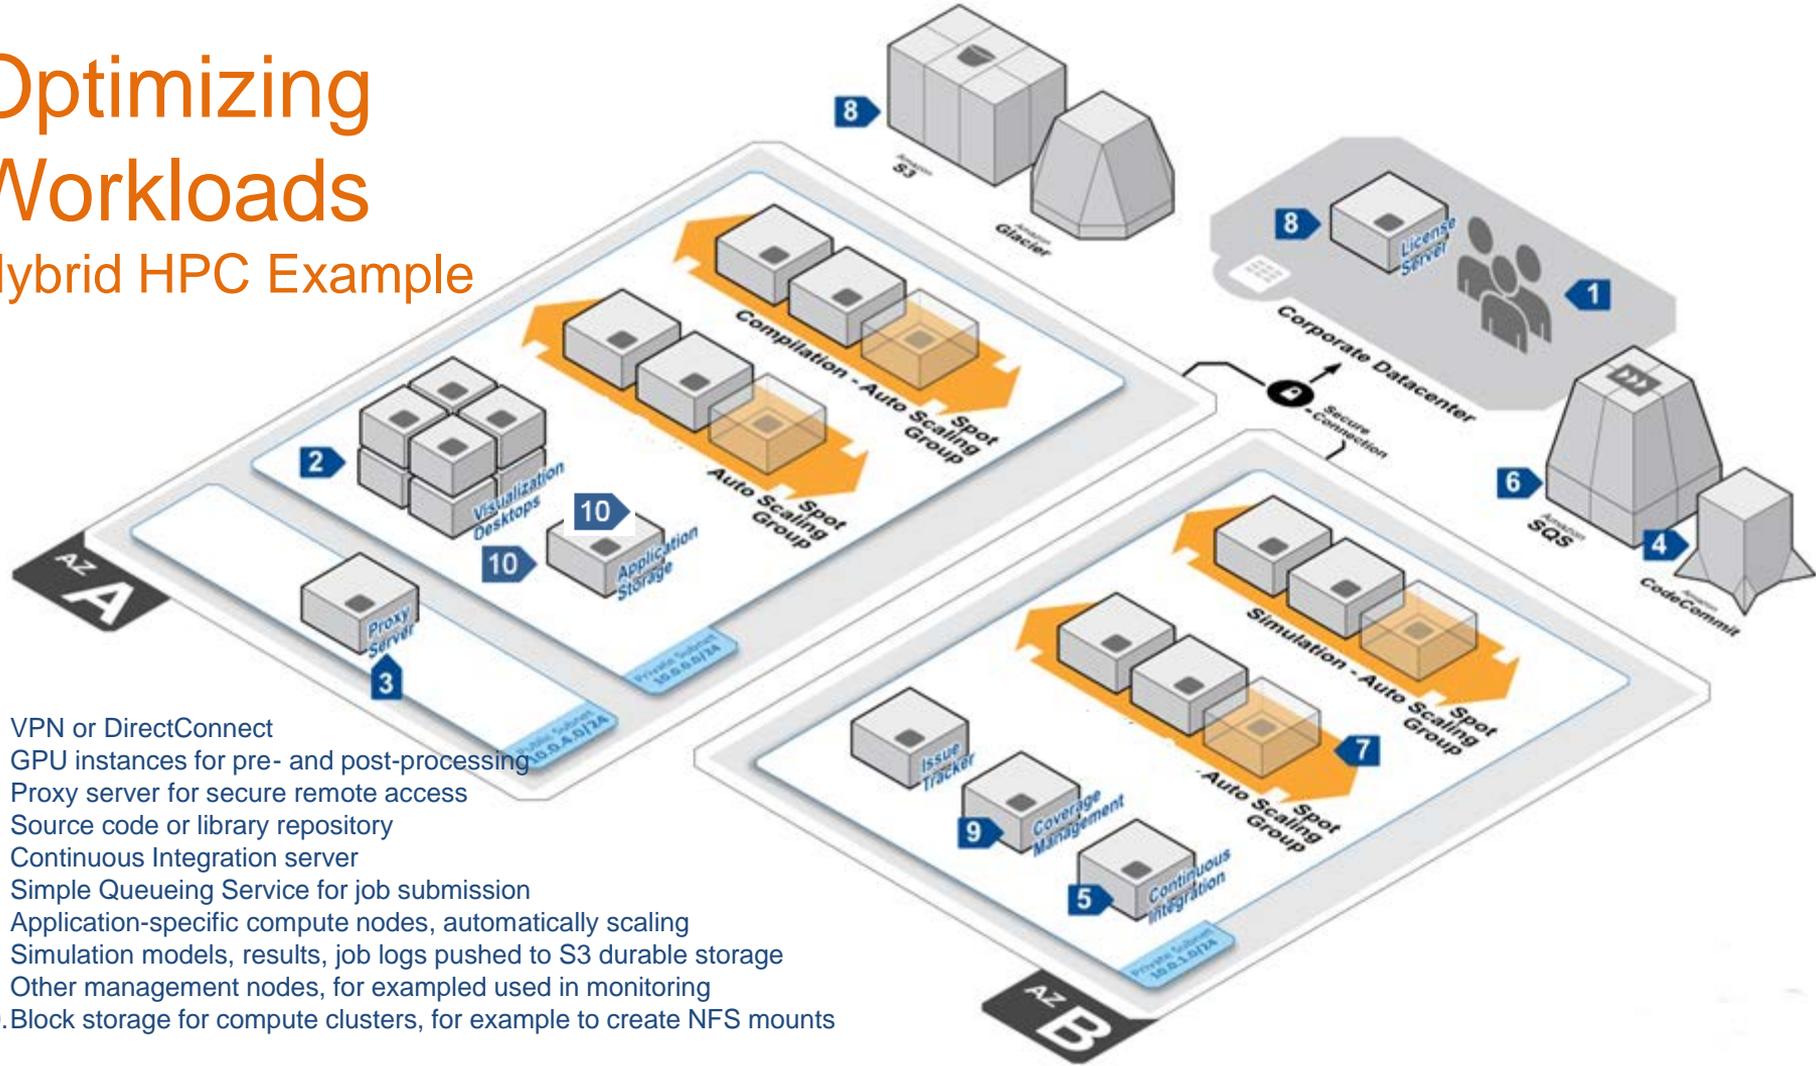


Private connectivity from private data center to an Amazon VPC

Use AWS Direct Connect to establish a private virtual interface from your on-premise network directly to your Amazon VPC

Optimizing Workloads

Hybrid HPC Example



1. VPN or DirectConnect
2. GPU instances for pre- and post-processing
3. Proxy server for secure remote access
4. Source code or library repository
5. Continuous Integration server
6. Simple Queueing Service for job submission
7. Application-specific compute nodes, automatically scaling
8. Simulation models, results, job logs pushed to S3 durable storage
9. Other management nodes, for example used in monitoring
10. Block storage for compute clusters, for example to create NFS mounts

Secure, Petabyte Scale Data Transport

**Ruggedized case
"8.5G Impact"**



E-ink shipping label



**50TB & 80TB
10G network**

**Rain & dust
resistant**



**Tamper-resistant
case & electronics**

**All data encrypted
end-to-end**

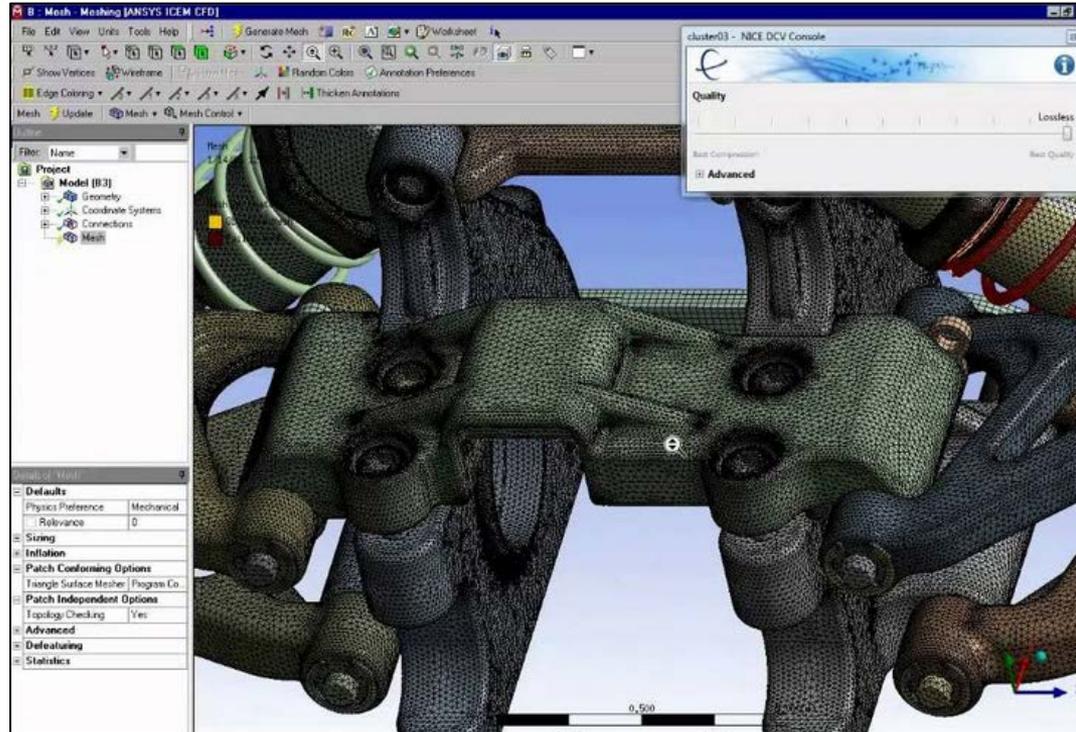
Secure Graphics and Collaboration

Cloud can be used for pre-and post processing as well as HPC

- Use GPUs in the cloud for remote rendering and remote desktops

Cloud is more secure for collaboration

- Encrypt the data in flight and at rest
- Manage your own keys and credentials
- Deliver pixels to your collaborators, not the actual data

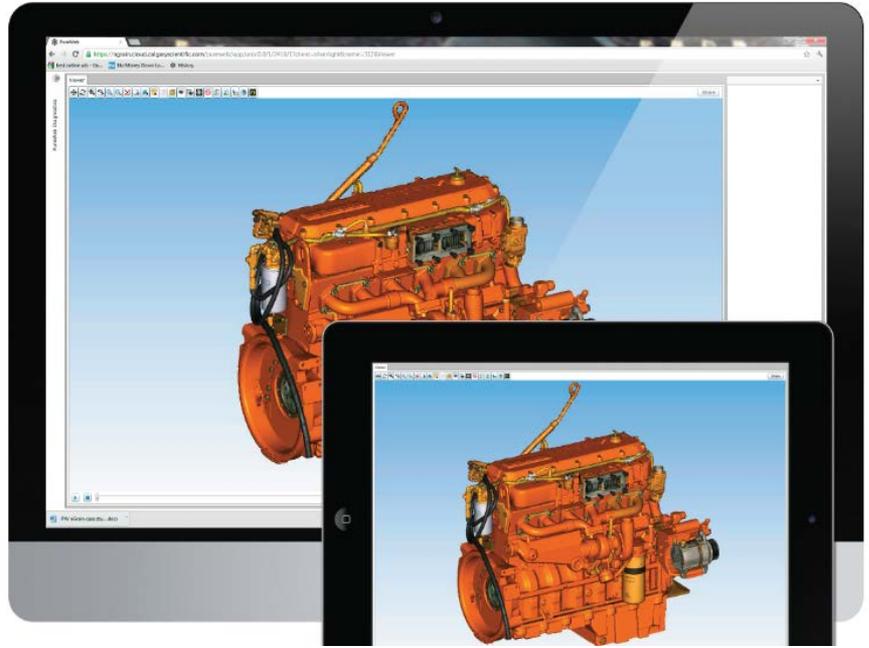


Thin Client Remote Collaboration

Calgary Scientific PureWeb™



calgary  scientific



Enabling Global Collaboration



Bring the users to the data, don't send the data to the users

Global Collaboration for Global Manufacturing

GE on AWS



GE's Global Business Integration Technologies Laboratory wanted to advance traditional manufacturing and create a dynamic network of people and machines that would allow collaboration, rapid prototyping, and product development for complex systems. GE had to adhere to U.S. International Traffic in Arms Regulations (ITAR) regulations and other compliance requirements. By using [AWS GovCloud \(US\)](#), GE developed a revolutionary manufacturing platform, Crowd-driven Ecosystem for Evolutionary Design (CEED), which connects people, materials, models, simulation, and equipment in an ITAR-compliant, secure, and distributed global environment.



JOSEPH J. SALVO
MANAGER, GE

Smart manufacturing Crowd-driven Ecosystem for Evolutionary Design (CEED)

Design



Thermal

Share



Mechanical

Simulate



Electrical

Schedule



Manufacturing

People, materials, models, simulation, equipment,
all connected in one distributed global environment.

Security



SECURITY & COMPLIANCE



Identity Management



Access Control



Key Management & Storage



Monitoring & Logs



Configuration Compliance



Web application firewall



Assessment and reporting



Resource & Usage Auditing

Security, Monitoring, Governance



Infrastructure Security

- **Network firewalls** are built into Amazon VPC
- **Web Application Firewall** lets you create private networks, and control access to your instances and applications
- **Encryption in transit** with TLS across all services
- **Connectivity options** that enable private, or dedicated, connections from your office or on-premises environment

Security, Monitoring, Governance



Inventory and Configuration

- **Amazon Inspector** automatically assesses applications for vulnerabilities or deviations from best practices, including impacted networks, OS, and attached storage
- **Deployment tools** help manage the creation and decommissioning of resources according to organization standards
- **AWS Config** helps identify AWS resources and then track and manage changes to those resources over time
- **AWS CloudFormation** to create standard, preconfigured, secure environments

Security, Monitoring, Governance



Data Encryption

- **Flexible key management** options including AWS Key Management Service
- Choose whether to have AWS manage the encryption keys, or keep complete control over your keys
- **Dedicated, hardware-based cryptographic key storage** is available using AWS CloudHSM
- AWS APIs allow you to integrate encryption and data protection with any of the services you develop or deploy in an AWS environment

Security, Monitoring, Governance



Identity and Access Control

- **Identity and Access Management** lets you define individual user accounts with permissions across AWS resources
- **AWS Multi-Factor Authentication** for privileged accounts, including options for hardware-based authenticators
- **AWS Directory Service** allows you to integrate and federate with corporate directories to reduce administrative overhead
- **Native identity and access management integration** across many AWS services, plus API integration with any of your own applications or services

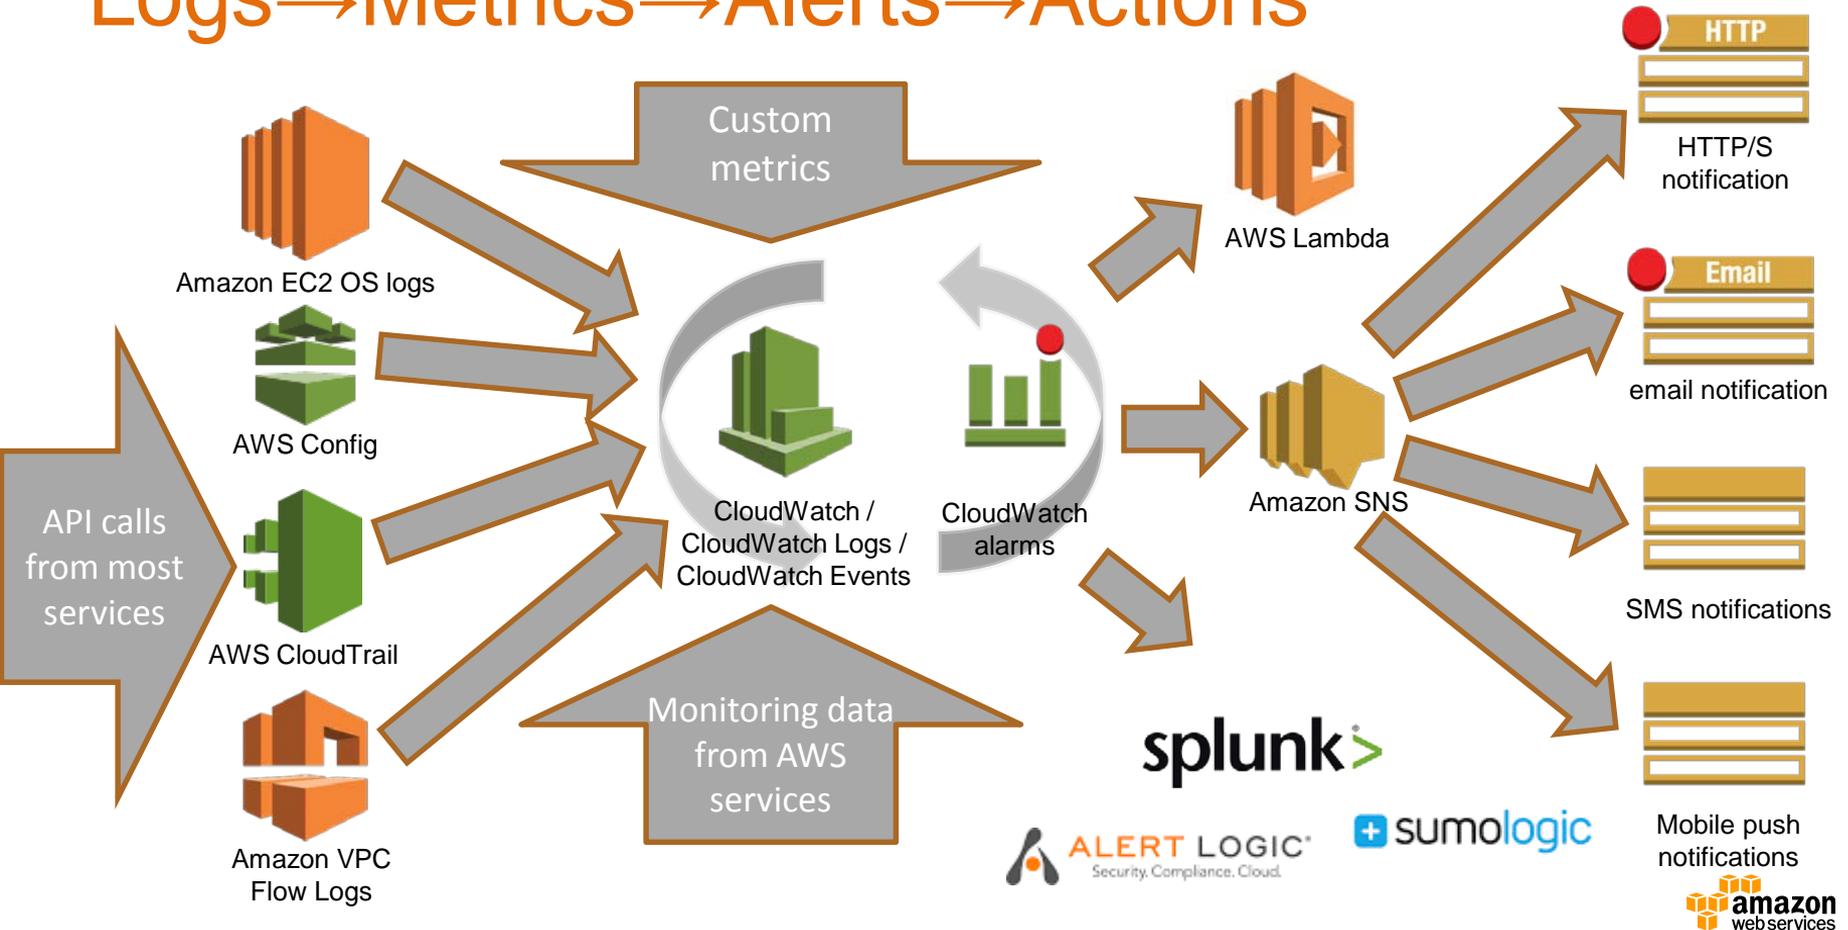
Security, Monitoring, Governance



Monitoring & Logging

- **AWS CloudTrail** provides deep visibility into API calls, including who, what, who, and from where calls were made
- Includes log aggregation options, streamlining investigations and compliance reporting
- **Amazon CloudWatch** to alert on specific events, or if thresholds are exceeded
- **AWS monitoring tools** give you the visibility you need to spot issues before they impact the business

Logs → Metrics → Alerts → Actions



Resources

aws.amazon.com/hpc

aws.amazon.com/big-data/

aws.amazon.com/security

dpelleri@amazon.com