ORACLE

Architecture and Design Principles of the Oracle Cloud Infrastructure

Peter Szegedi

Senior Technology Solution Engineer

Oracle Digital – EMEA Tech North

Safe harbor statement

The following is intended to outline our general product direction. It is intended for information purposes only, and may not be incorporated into any contract. It is not a commitment to deliver any material, code, or functionality, and should not be relied upon in making purchasing decisions.

The development, release, timing, and pricing of any features or functionality described for Oracle's products may change and remains at the sole discretion of Oracle Corporation.

Architected for superior performance

Faster components than other clouds, on par with on-premises

- Bare metal, VMs, GPUs, local storage, block storage, database

More available network bandwidth between components

- No over-subscription, no noisy neighbors, very low latency

First RDMA cluster offering 1.5 microsecond latency in the cloud

- The toughest on-prem workloads can now run in the cloud

OCI Region – HA Building Blocks

- Multiple fault de-correlated, completely independent datacenters: Availability Domain (AD)
- Grouping of hardware and infrastructure within an AD: Fault Domain
- Predictable low latency & high speed, encrypted interconnect between ADs

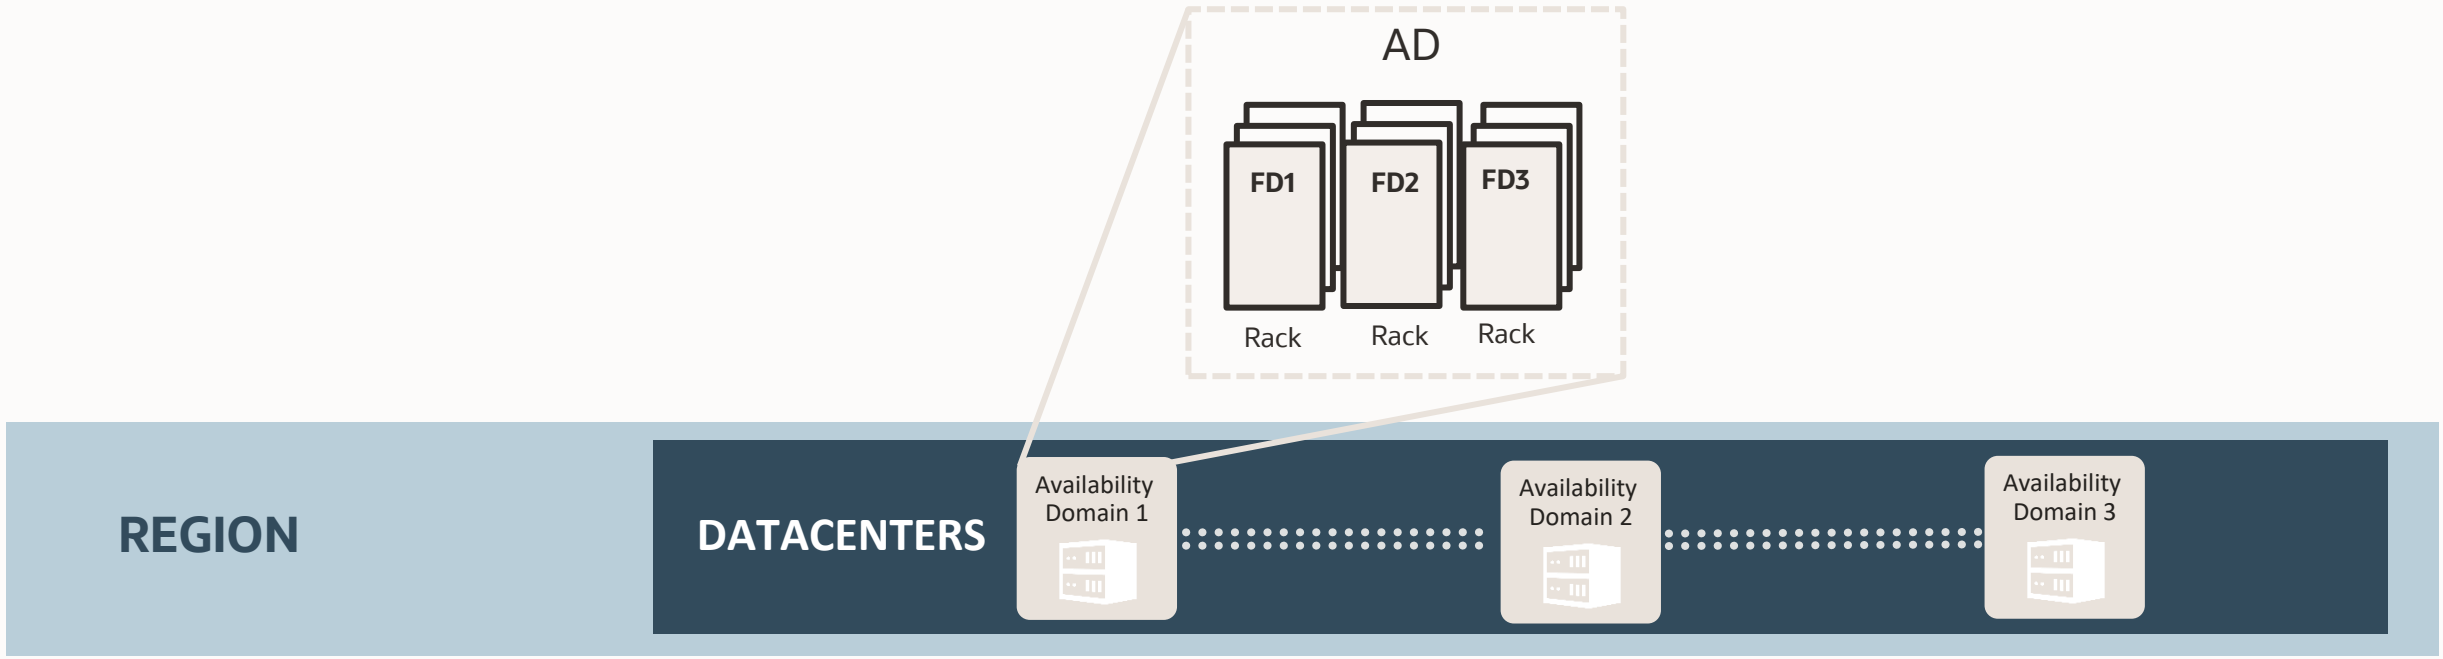

Oracle Cloud Infrastructure Global Footprint

- **As of October 2019**

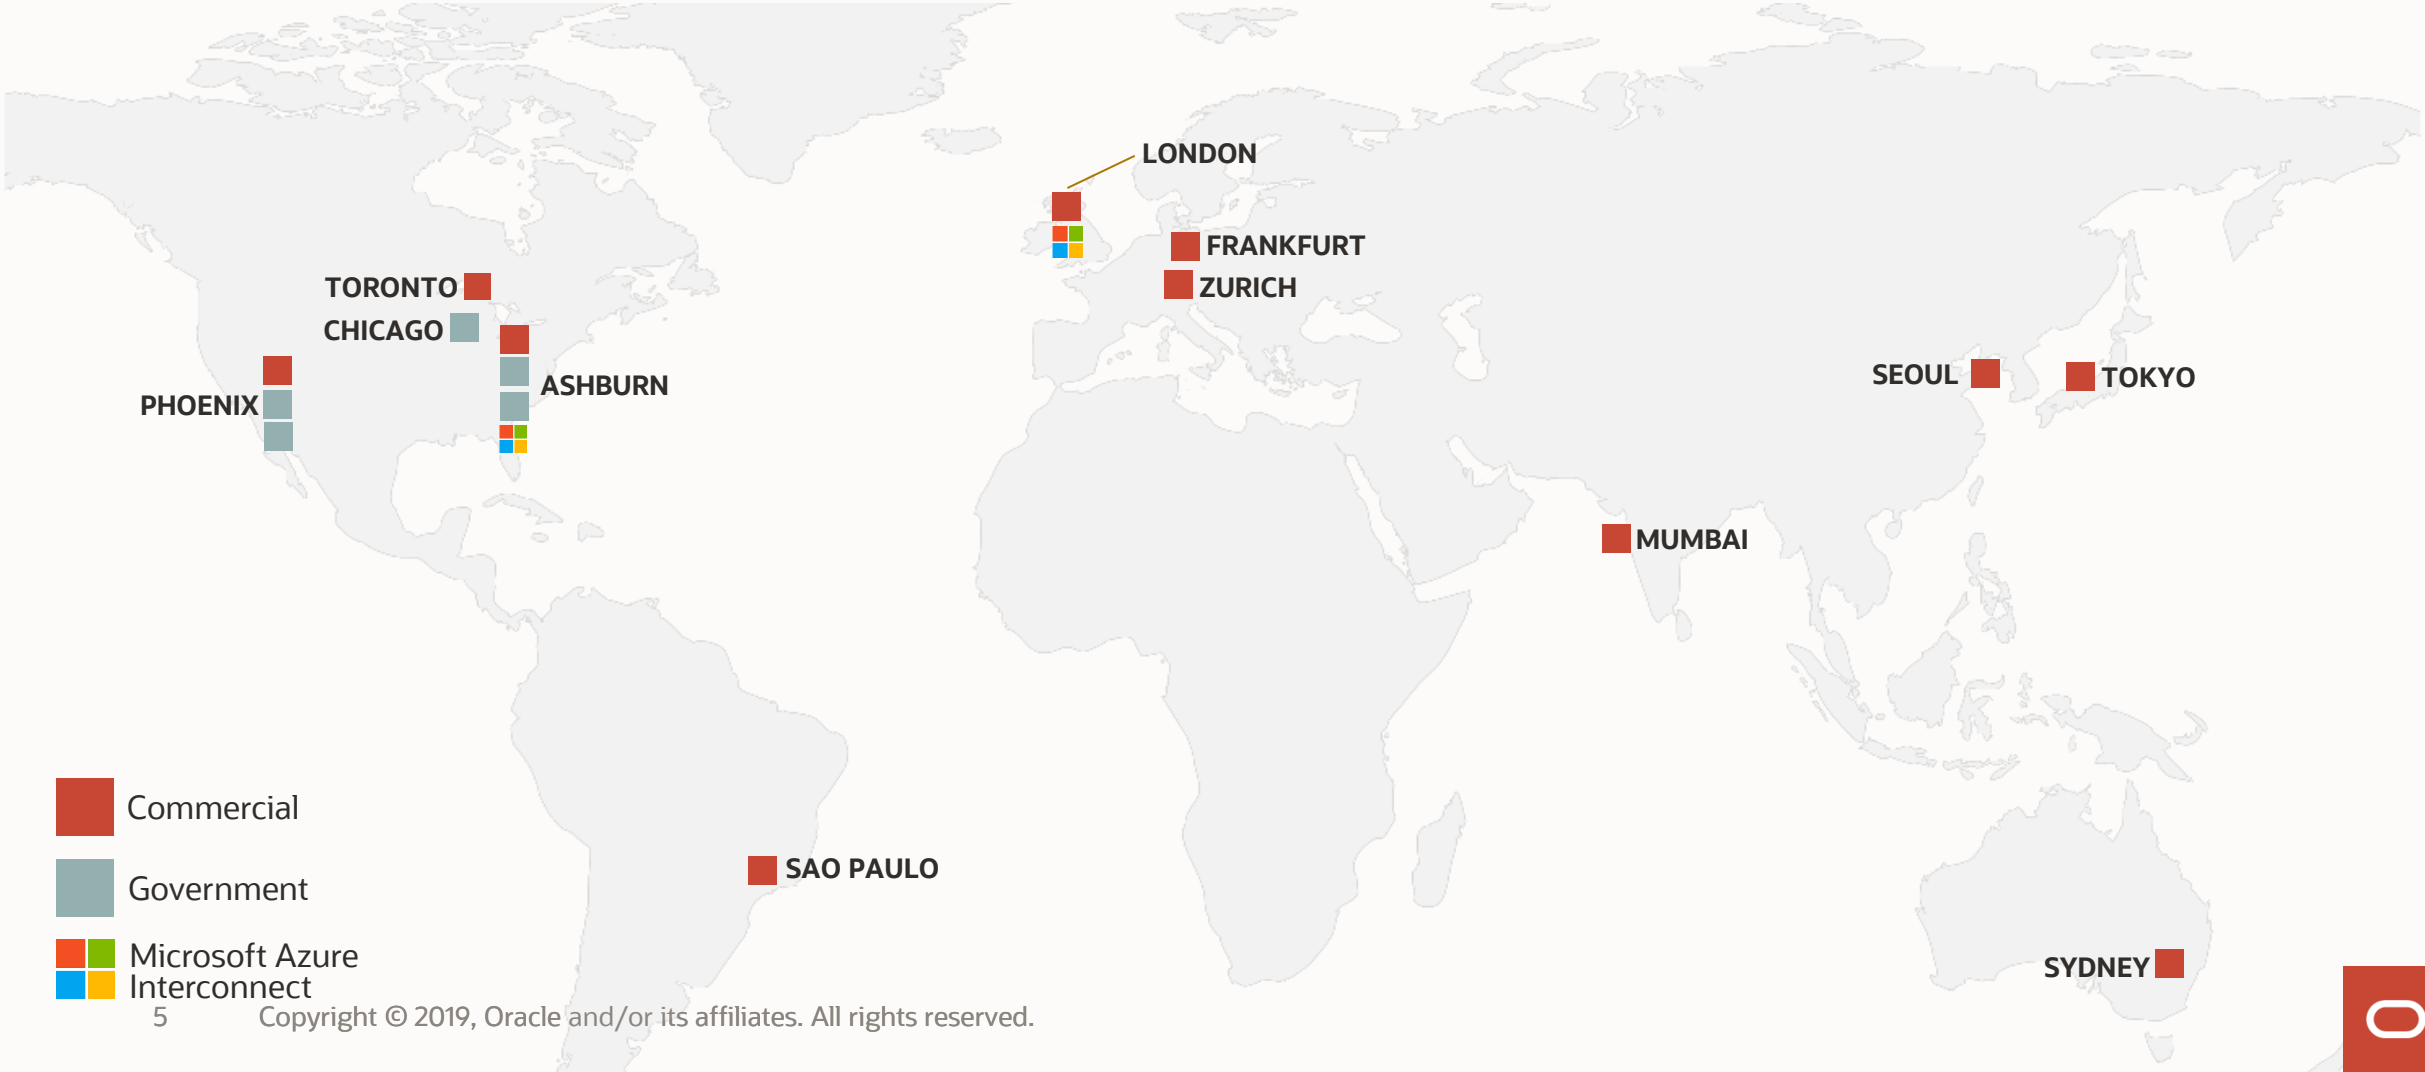

Oracle Cloud Infrastructure Global Footprint

- **End of CY2020: 36 Oracle Regions**

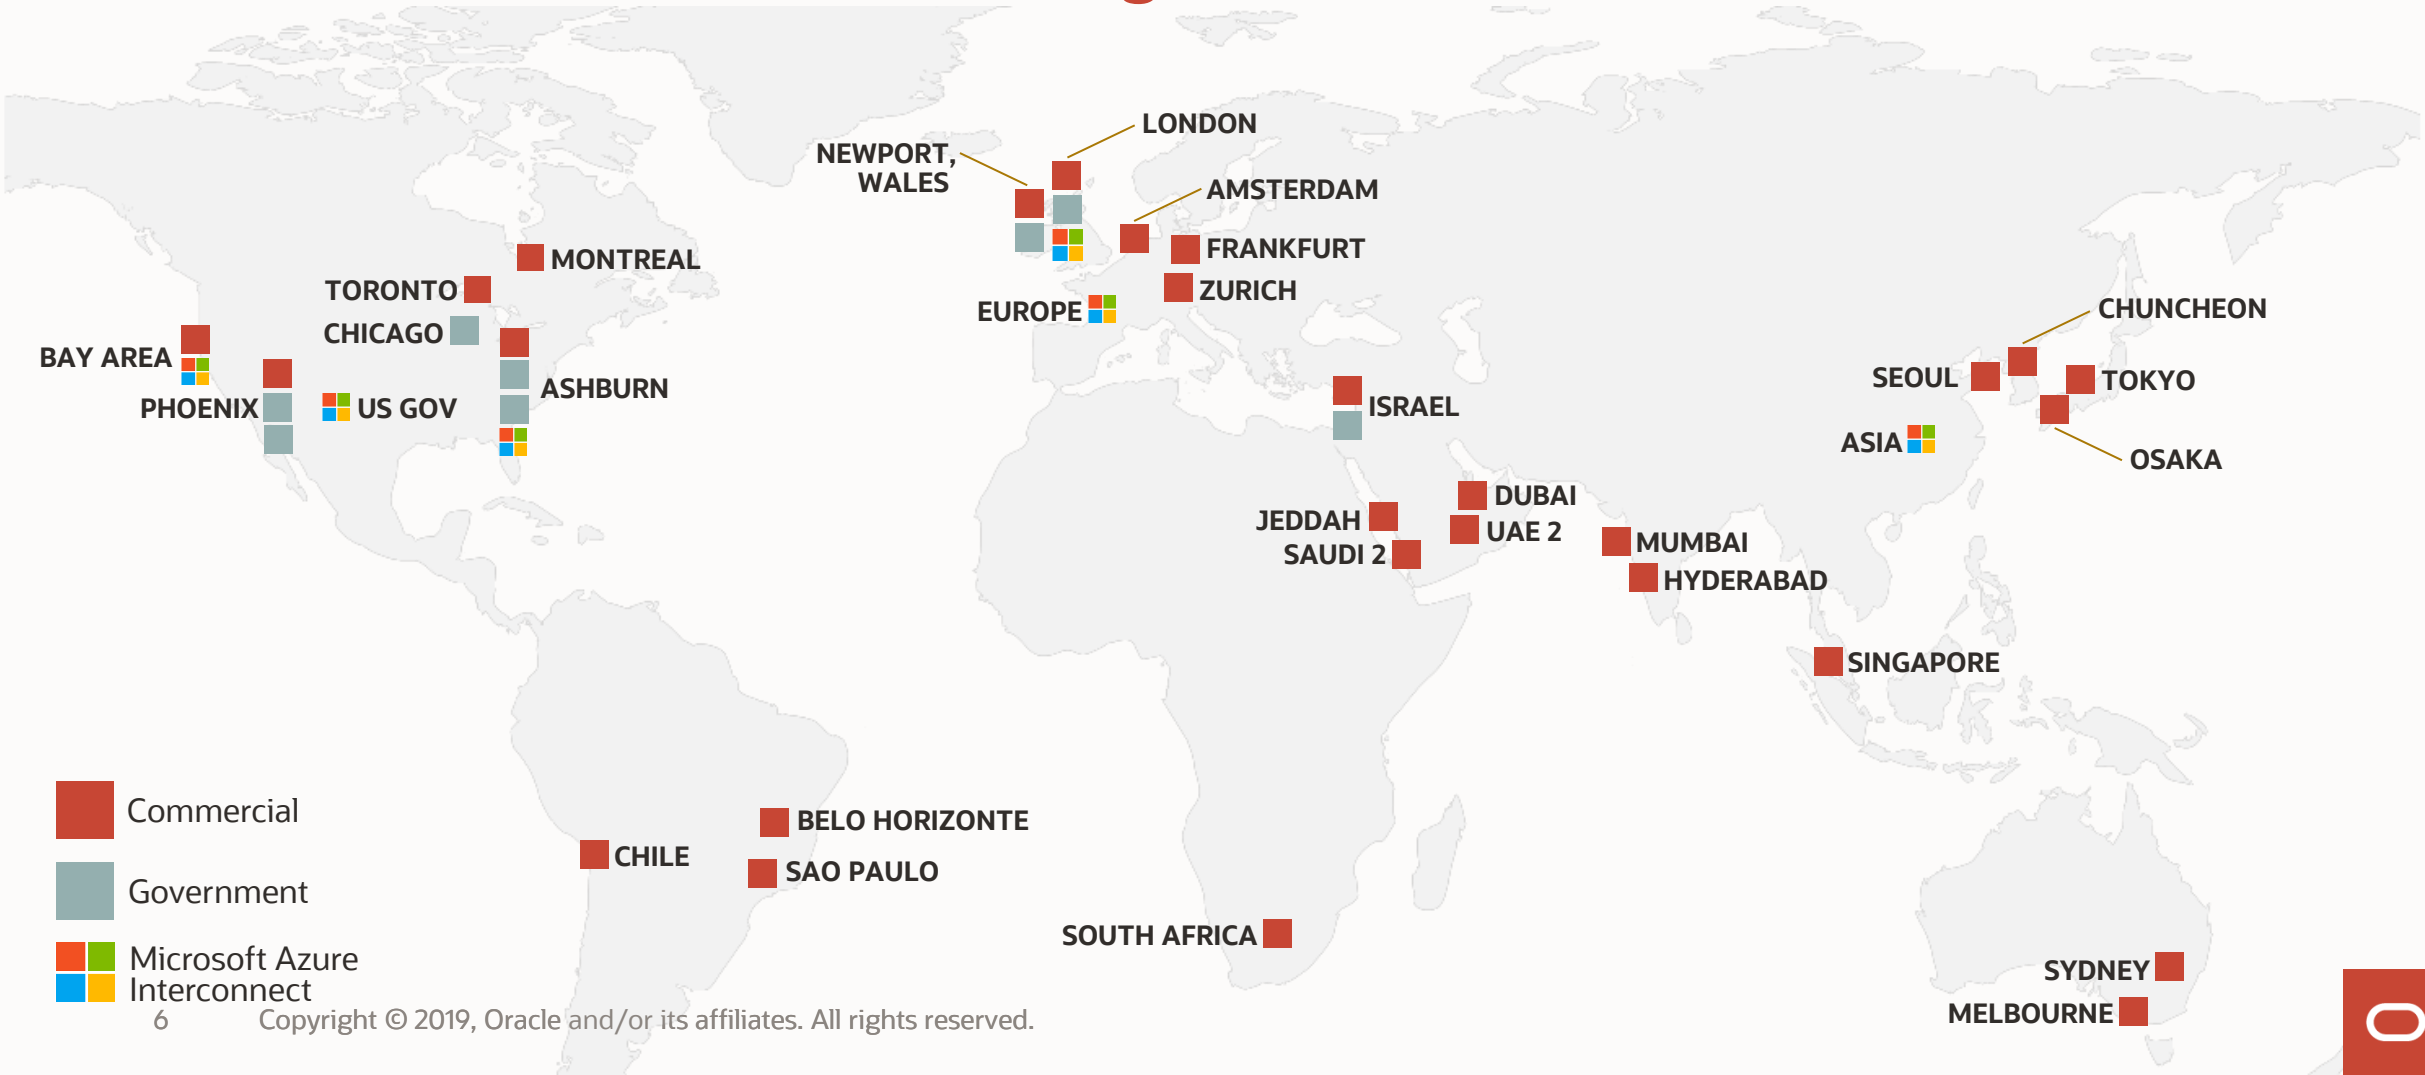

Inside an AD – High Scale, High Performance Network

- Non-oversubscribed network; no noisy-neighbors
- Very high scale – ~1 million network ports in an AD
- Predictable low latency & high speed interconnect between hosts in an AD

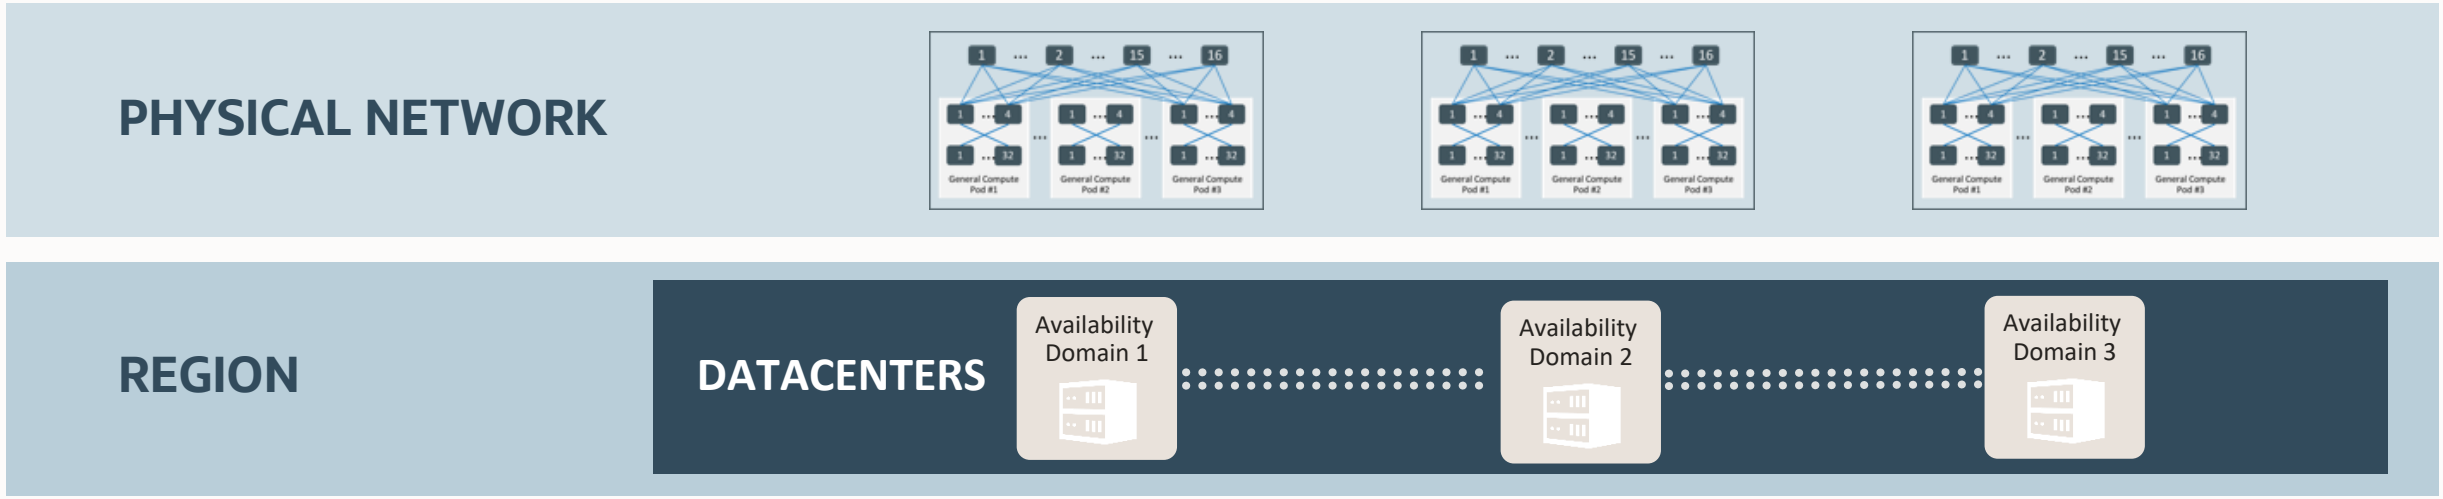

Why are we better?

In other clouds, data has to go through a VM and a hypervisor before it can access the underlying server hardware

Other clouds automatically distribute the VMs across the entire datacenter

This adds significant latency and un-predictability in performance

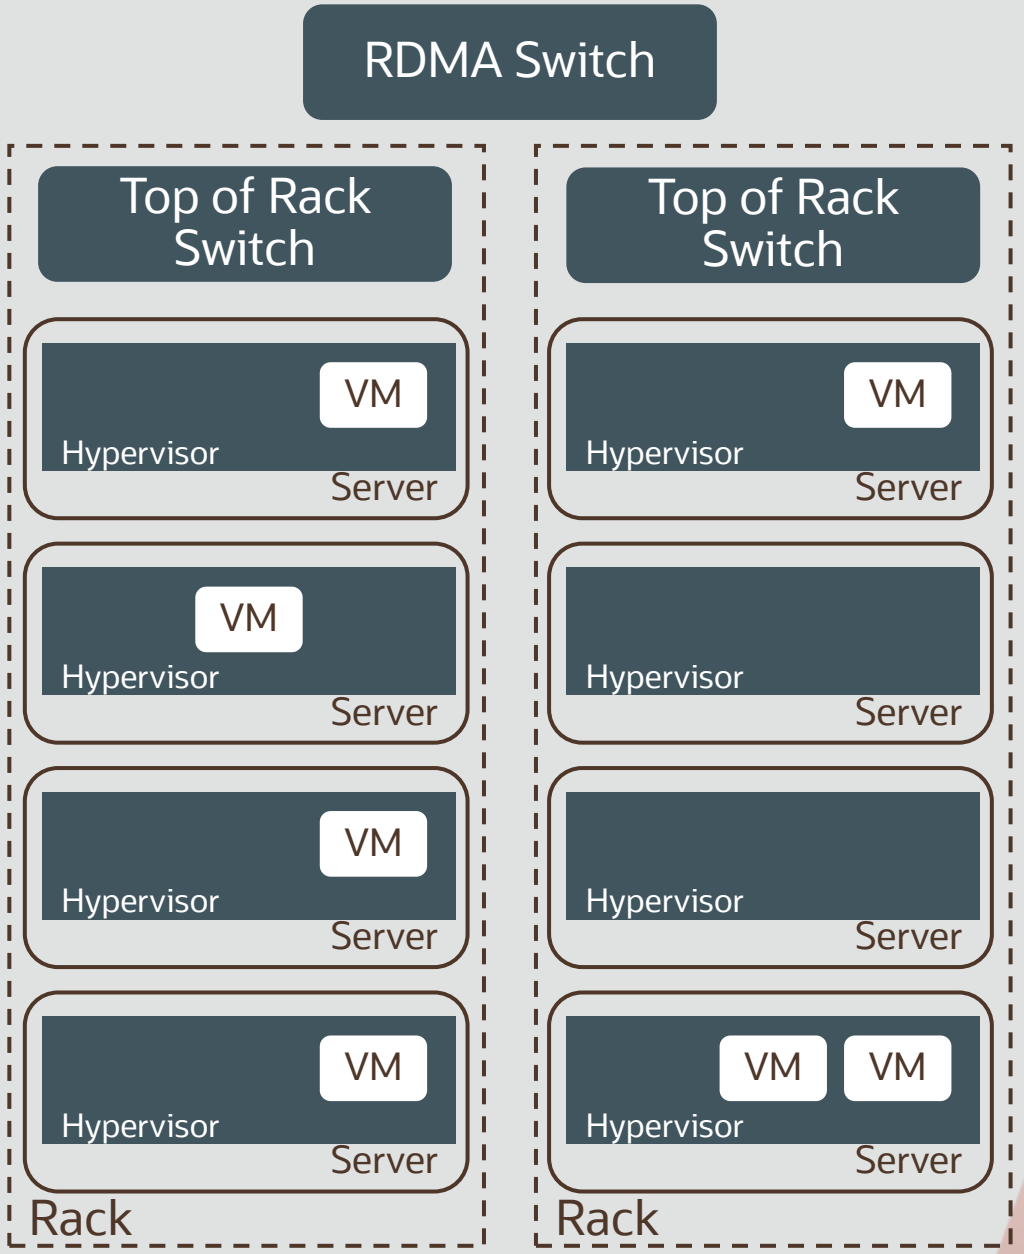

Why are we better?

Oracle connects the servers directly to the RDMA switch

No hypervisor, no virtualization, no jitter

We allow customer to chose their own placement to maximize stability

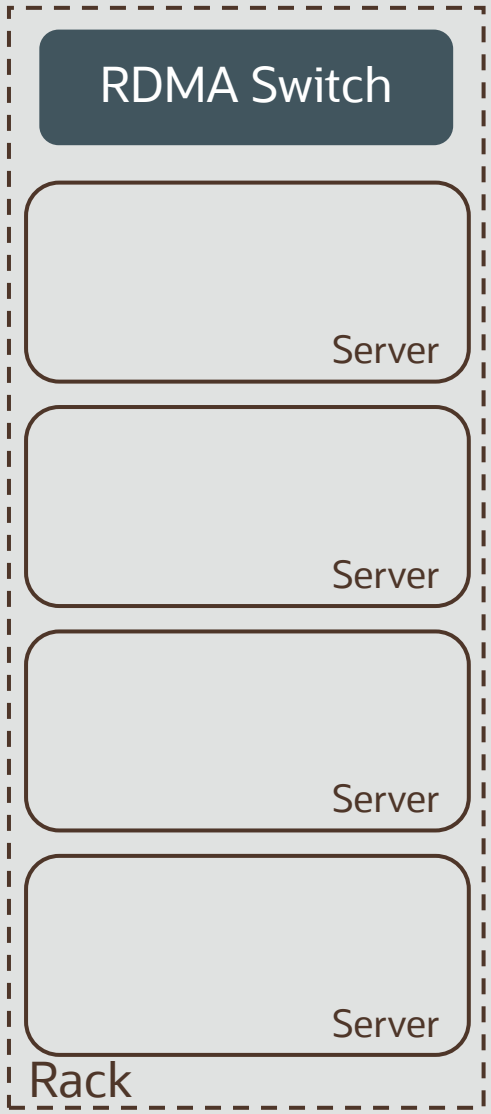

ORACLE

Industry-first cluster networking

Cloud-first network for ultra low latency and high bandwidth

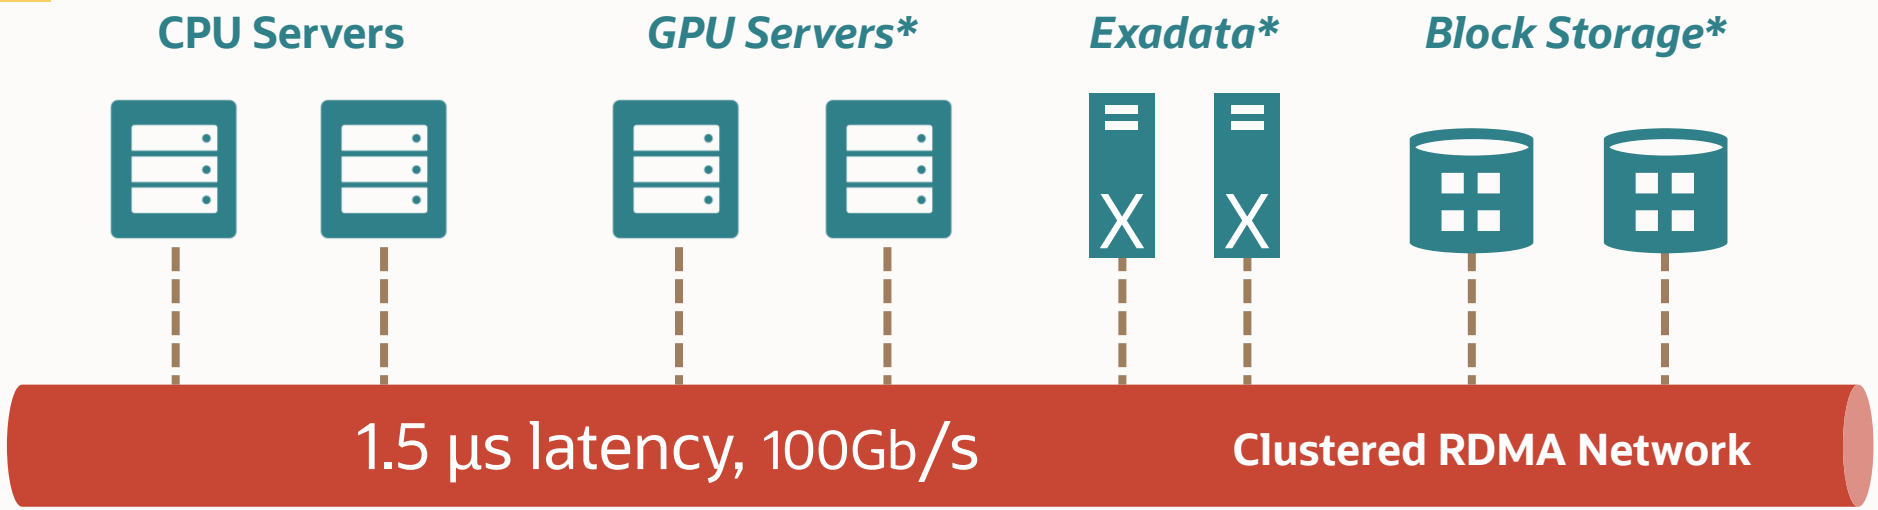

For high performance workloads (HPC, Database, Big Data, AI) including the hardest product development workloads like CFD, Crash Simulations, Reservoir Modelling, DNA Sequencing

Coming Soon: GPU Bare-Metal Instances with 8 GPUs, 2TB RAM, 25TB local NVMe storage, and 8x 100GB Network ports for cluster networking and RDMA

*Planned

Off-box Network Virtualization

Off Box Network Virtualization – moves storage and network IO out of the hypervisor and enables lower overhead and bare metal instances

VIRTUAL NETWORK

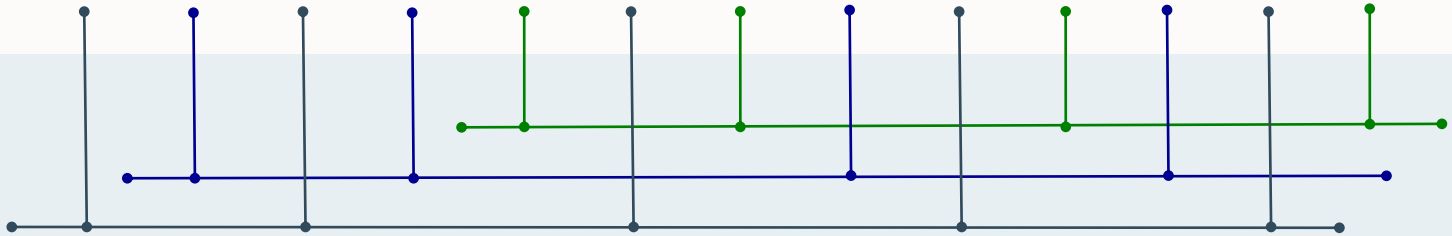

PHYSICAL NETWORK

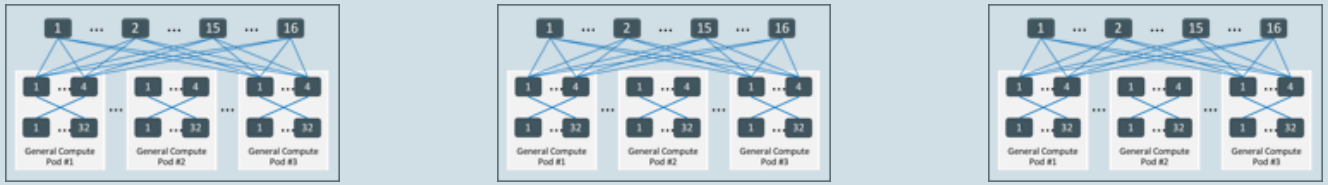

REGION

DATACENTERS

Availability Domain 1

Availability Domain 2

Availability Domain 3

Security

Architected from the ground up for maximum isolation and protection

1st Generation Clouds:
Most Prevalent Today

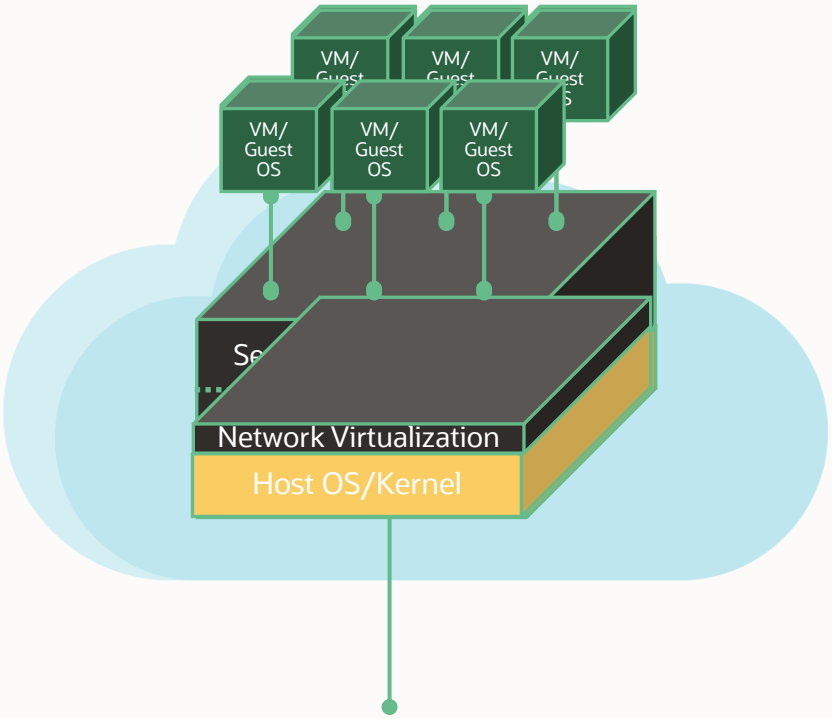

2nd Generation Cloud:
Oracle Cloud Infrastructure-Wide

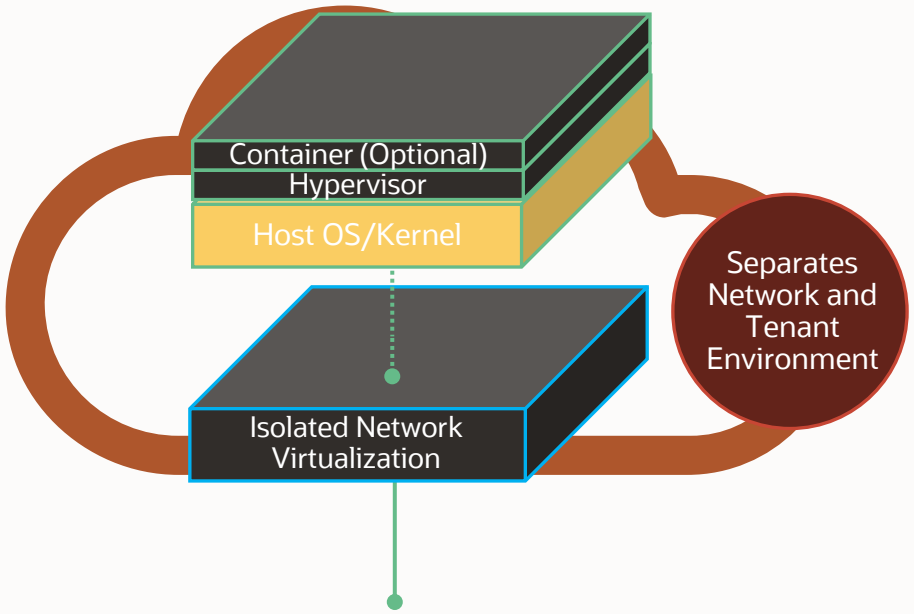

Oracle Cloud Infrastructure Services

COMPUTE, STORAGE, DATABASE, LBs, Security…

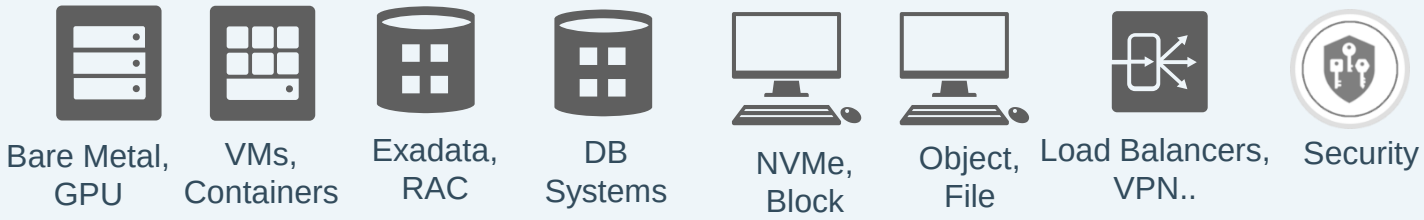

VIRTUAL NETWORK

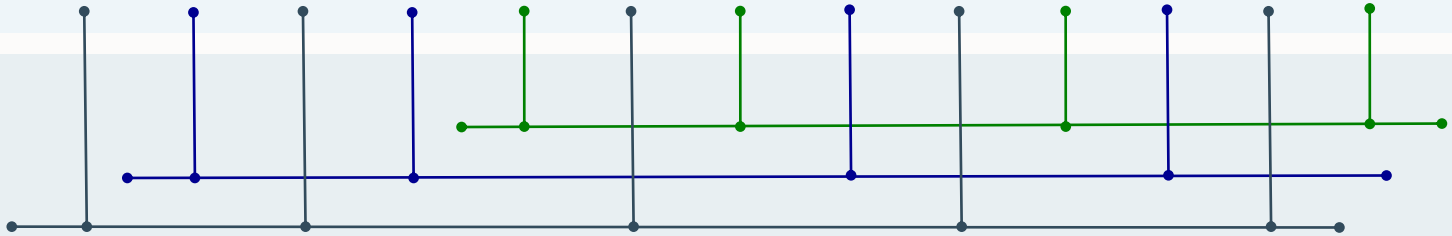

PHYSICAL NETWORK

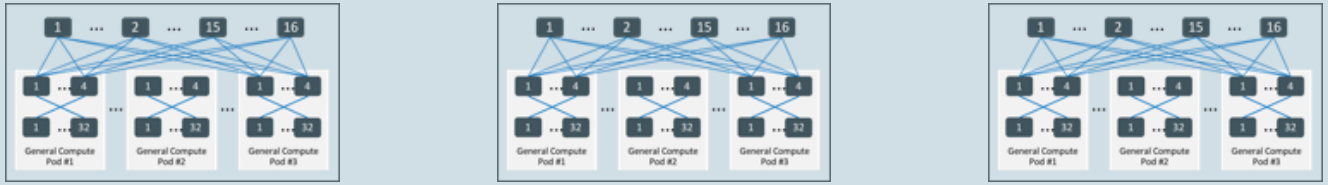

REGION

Networking flexibility and control

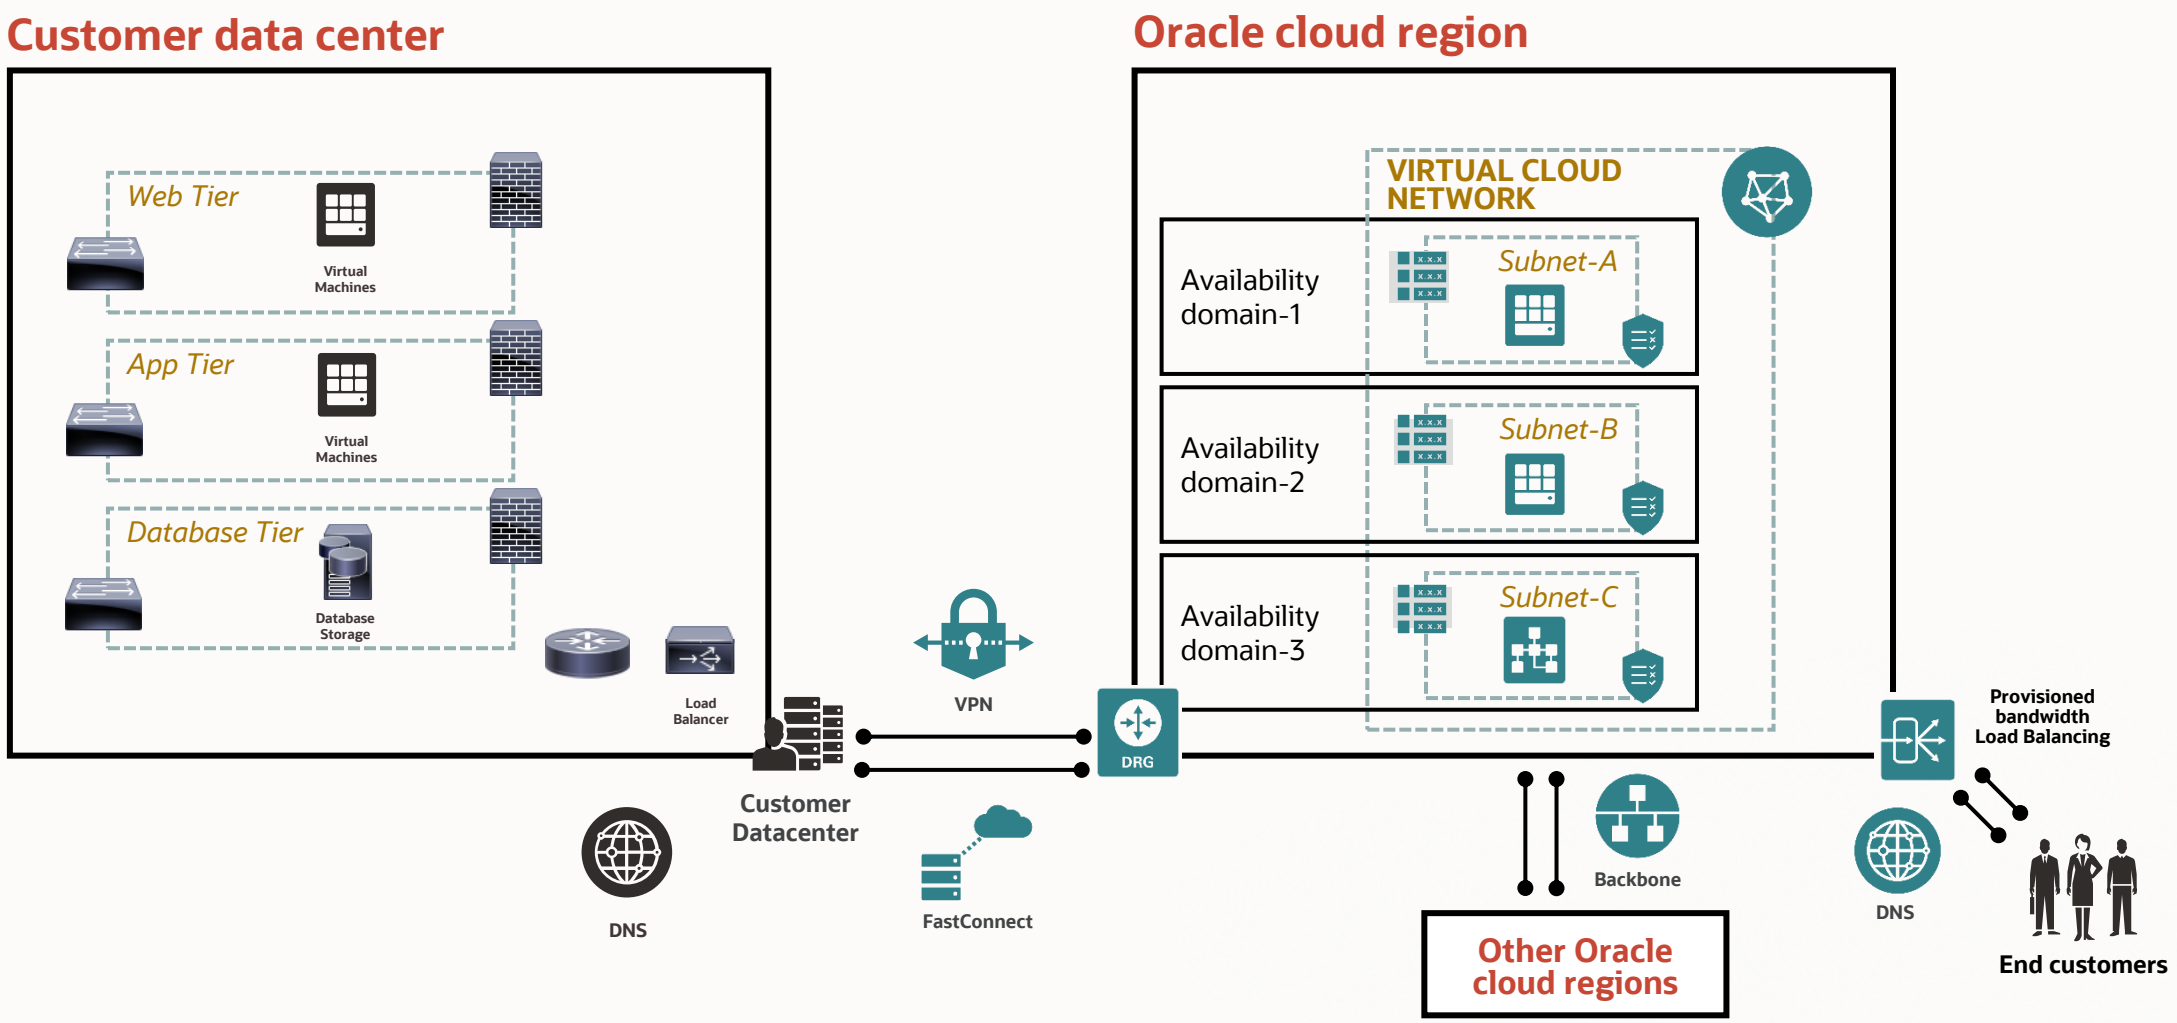

Oracle Gen 2 Cloud: Infrastructure

Oracle Cloud offers the best performance per server

	ORACLE Cloud Infrastructure	AWS	Oracle offers
Compute	2.57 TFLOPs	1.77 TFLOPs	**45% more**
Memory	153.02 GB/s	134.03 GB/s	**14% more**
Block Storage	500,000 IOPS	80,000 IOPS	**525% more**

2-5x faster vs. on-prem and other clouds

Backed by performance SLAs

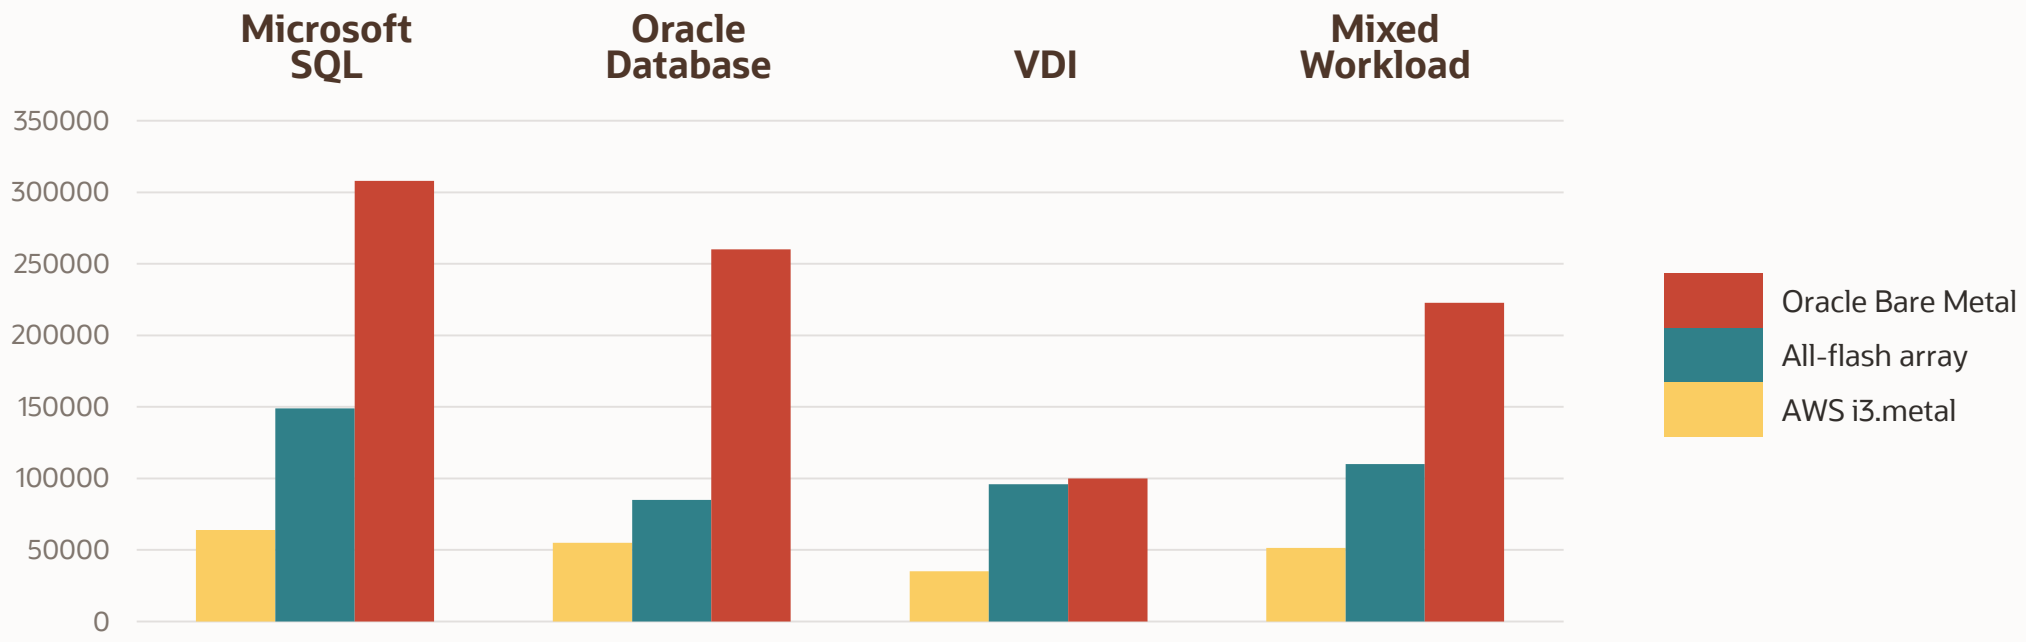

StorageReview

www.storagereview.com/oracle_cloud_infrastructure_compute_bare_metal_instances_review
https://www.storagereview.com/dell_emc_unity_450f_allflash_storage_review

End-to-end cloud infrastructure SLAs

	Oracle	AWS	Azure	GCP
Availability	**Covered**	Covered	Covered	Covered
Performance	**Covered**	No coverage	No coverage	No coverage
Manageability	**Covered**	No coverage	No coverage	No coverage

Stronger isolation and control from core to edge

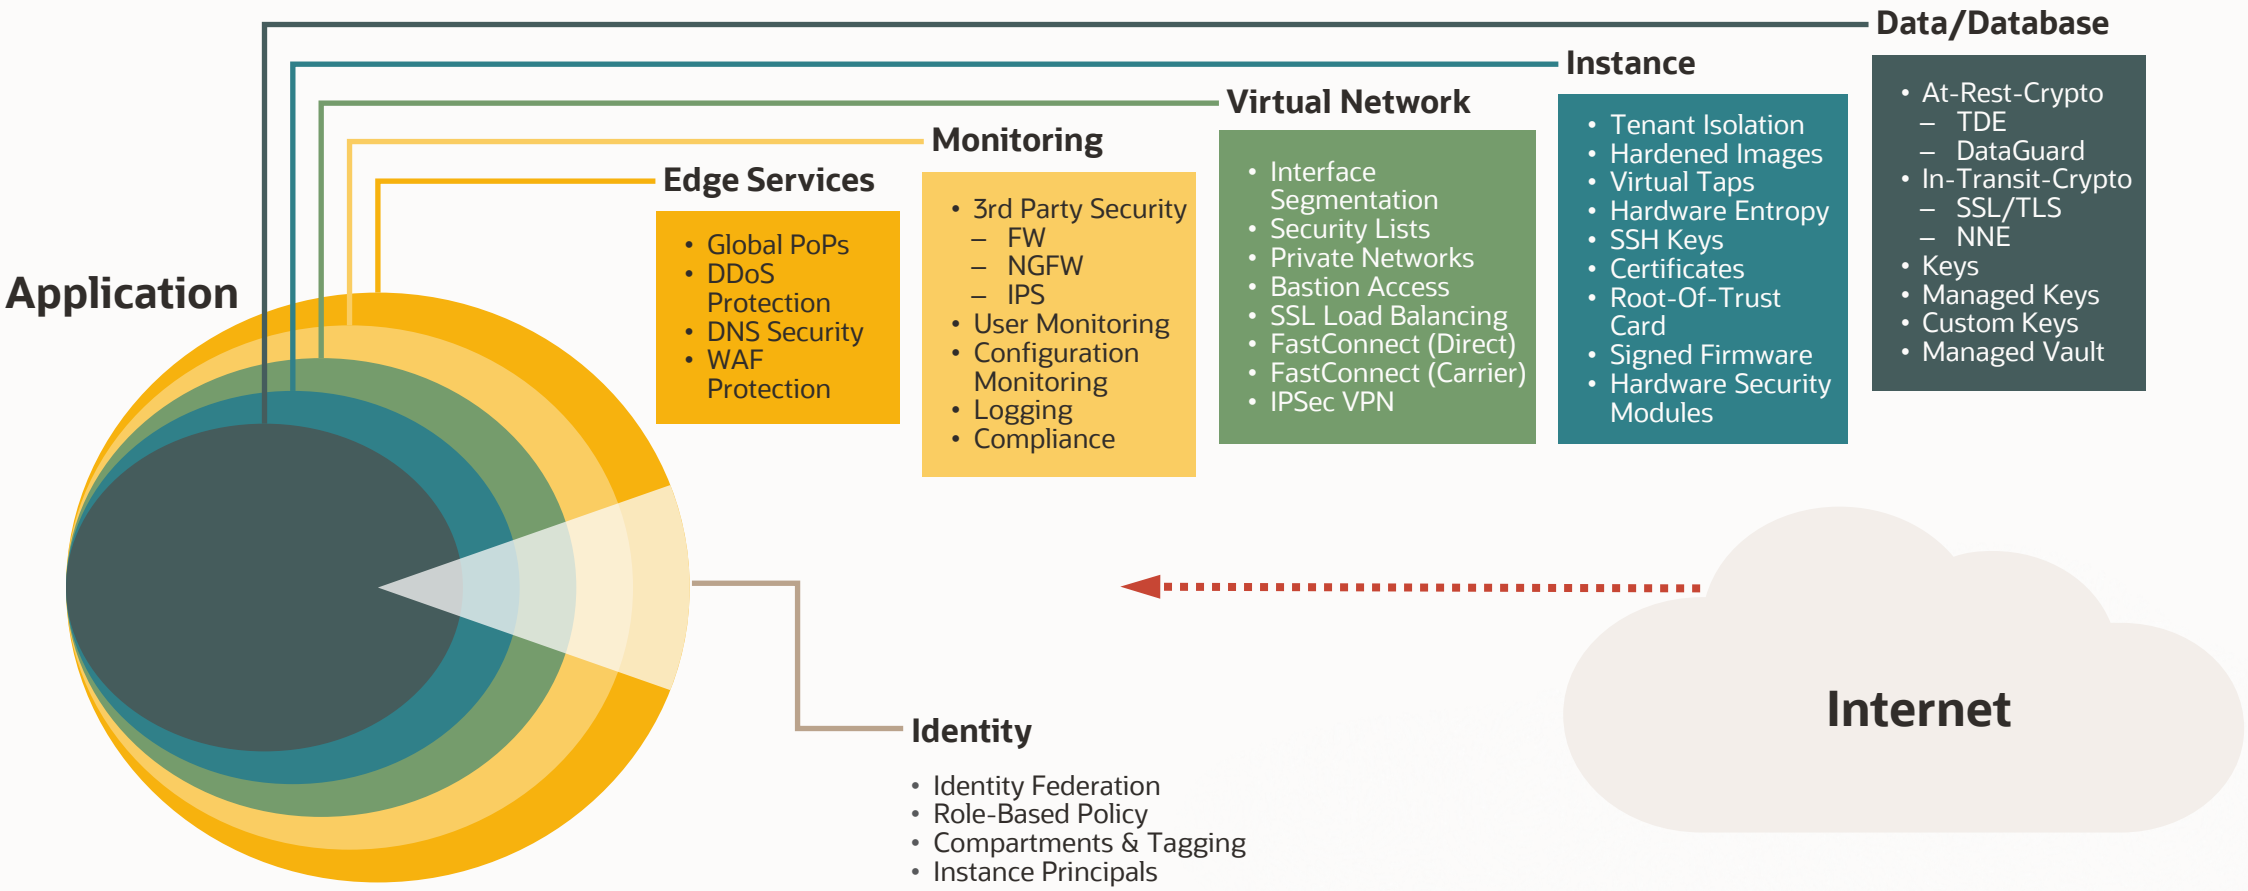

OCI compliance: Current audit programs

Global

- SOC 1 : SOC 2 : SOC 3
- 27001 : 27017 : 27018
- Self-Assessment
- US Privacy Shield

Government

- DoD DISA SRG IL2
- Moderate – Agency ATO
- VPAT – Section 508
- G-Cloud 11 - UK
- Model Clauses - EU

Industry

- HIPAA
- PCI DSS
- FISC - Japan
- IG Toolkit - UK

Regional

- GDPR - EU
- BSI C5 - Germany
- TISAX - Germany
- PIPEDA - Canada
- Cyber Essentials Plus - UK
- My Number - Japan
- Cloud Security Principles - UK

OCI - October 2017

Partner Ecosystem

FORTINET Microsoft

Next Layer Services

Dev Ops	Big Data & Analytics	Container Ecosystem	Hybrid Services	Migration / Hybrid
Cloud Shell	Big Data / Data Science	API Gateway	Dedicated Regions	VMware Solution
Dev Tools & Events	Data Integration / Data Catalog	Events / Functions	Azure Interconnect	Data Migration
Resource Manager	Streaming, Data Flow	K8S / Container Registry	Exadata Cloud@Customer	Identity (AD, SSO, LDAP)

Data Services

DBaaS	ATP	ADW	ExaCS	Analytics
VM and Bare Metal	Dedicated, Serverless	Autonomous Data Warehouse	Exadata Cloud Service	Oracle Analytics Cloud

Core Services

Compute	Storage	Networking	Security	Governance and Management
Bare Metal	Block	VCN	KMS, Virtual Vaults	Compartments, Tags
VMs, Dedicated VMs	Object	Cluster Networking	Policy, Data Safe	Logging / Monitoring
GPU shapes	Object Archive	Load Balancer	Identity, Secrets Mgmt	Audit
High frequency CPUs	File	Service Gateway	CASB	Email / Notifications
Autonomous Linux	Import Appliance	FastConnect, w/ IPSec	WAF	Cost Management
	Storage Gateway	DNS	DDoS	

Regions

Phoenix, USA | Ashburn, USA | Frankfurt, Germany

OCI - Now (October 2019)

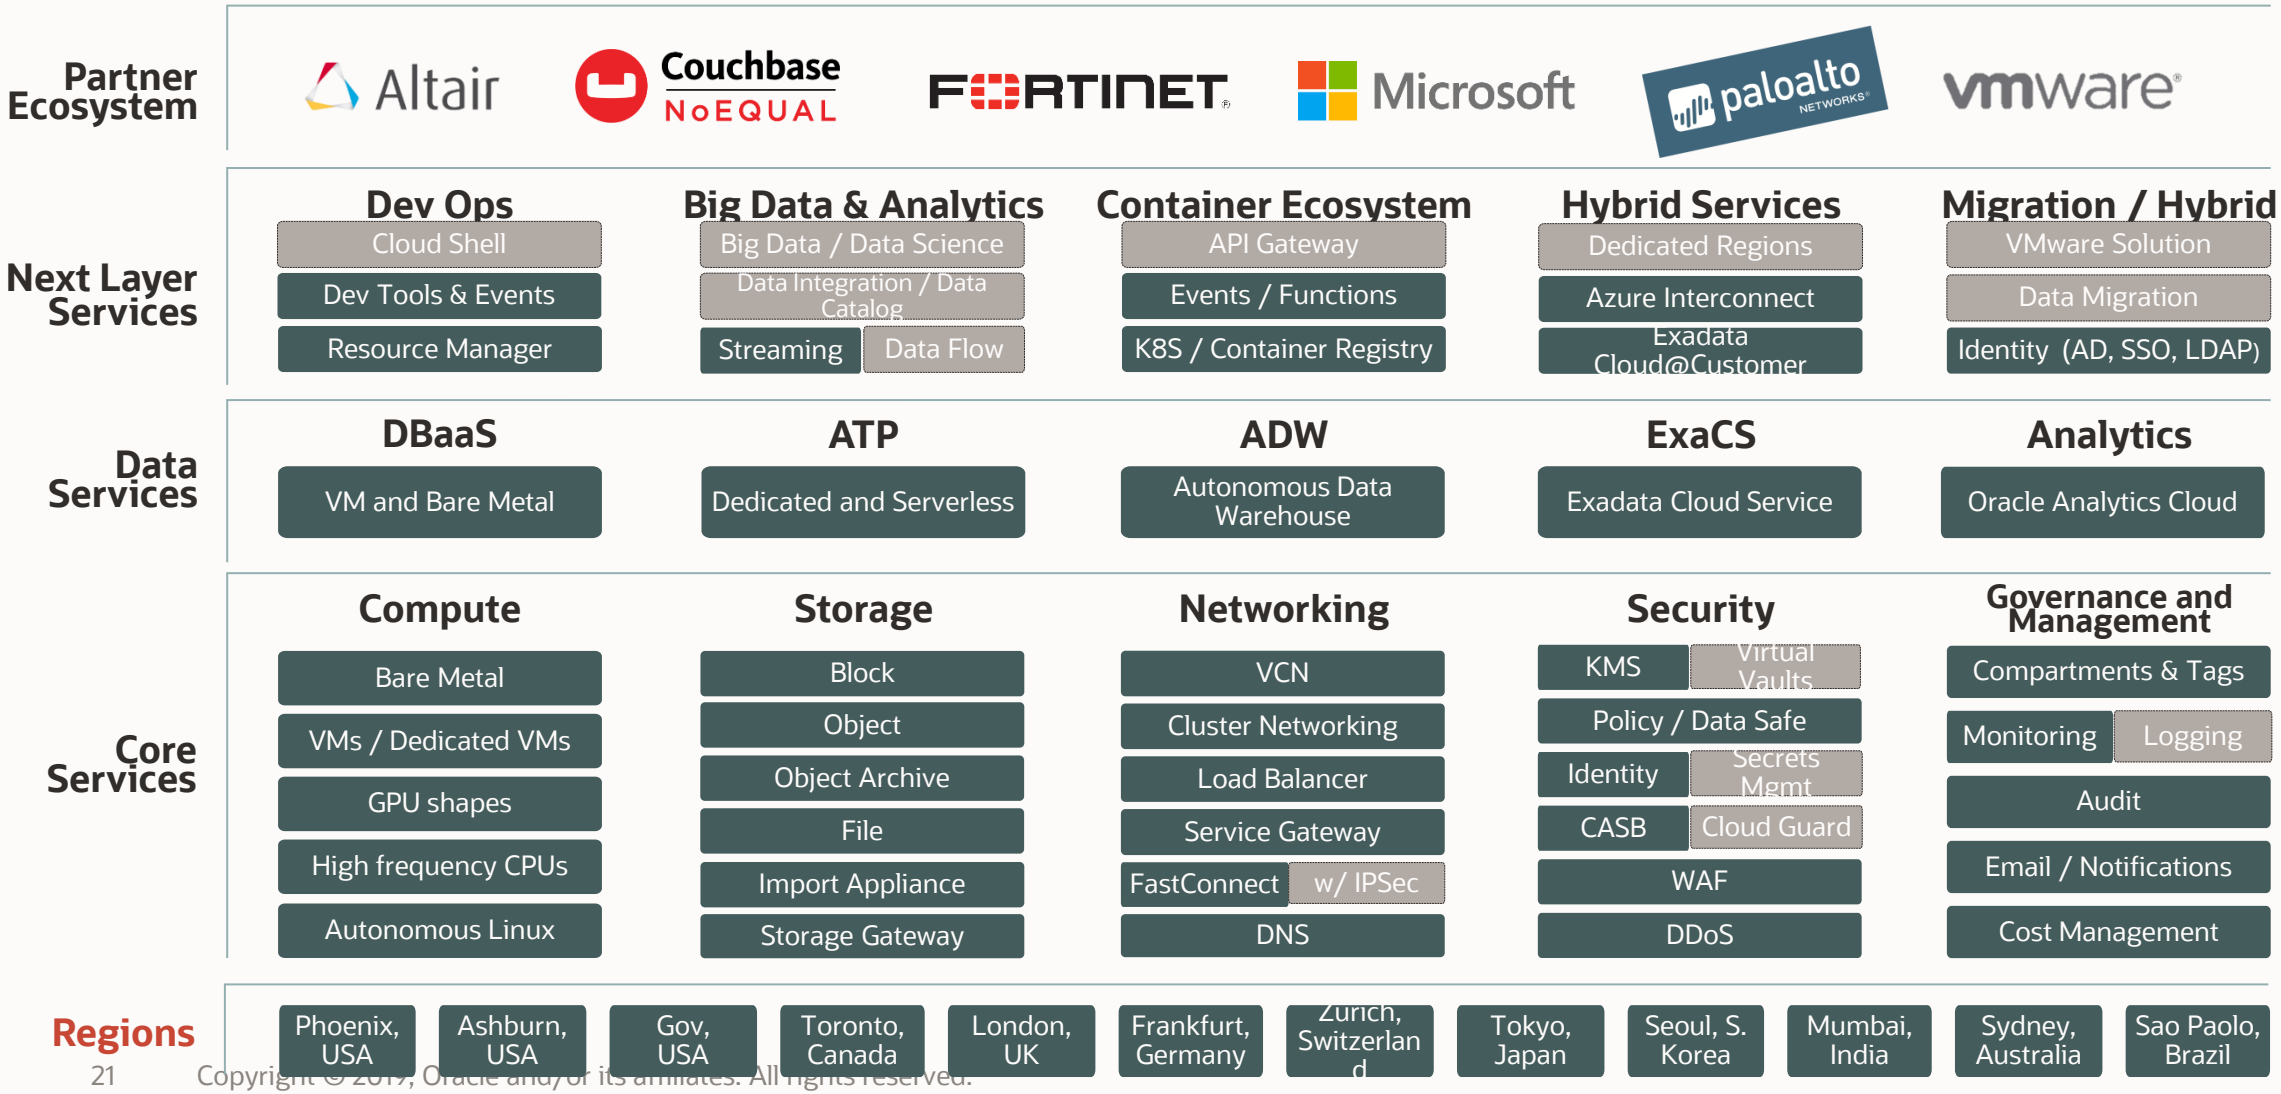

OCI - By End of CY 2020

Partner Ecosystem

Altair · Couchbase NoEQUAL · FORTINET · Microsoft · paloalto NETWORKS · vmware

Next Layer Services

Dev Ops	Big Data & Analytics	Container Ecosystem	Hybrid Services	Migration / Hybrid
Cloud Shell	Big Data / Data Science	API Gateway	Dedicated Regions	VMware Solution
Dev Tools & Events	Data Integration / Data Catalog	Events / Functions	Azure Interconnect	Data Migration
Resource Manager	Streaming / Data Flow	K8S / Container Registry	Exadata Cloud@Customer	Identity (AD, SSO, LDAP)

Data Services

DBaaS	ATP	ADW	ExaCS	Analytics
VM and Bare Metal	Dedicated and Serverless	Autonomous Data Warehouse	Exadata Cloud Service	Oracle Analytics Cloud

Core Services

Compute	Storage	Networking	Security	Governance and Management
Bare Metal	Block	VCN	KMS / Virtual Vaults	Compartments & Tags
VMs / Dedicated VMs	Object	Cluster Networking	Policy / Data Safe	Monitoring / Logging
GPU shapes	Object Archive	Load Balancer	Identity / Secrets Mgmt	Audit
High frequency CPUs	File	Service Gateway	CASB / Cloud Guard	Email / Notifications
Autonomous Linux	Import Appliance	FastConnect w/IPSec	WAF	Cost Management
	Storage Gateway	DNS	DDoS	

Regions

Phoenix	Ashburn	Toronto	UK/Gov	Frankfurt	Israel/Gov	Dubai	Jeddah	Tokyo	Seoul	Mumbai	Sydney	Sao Paolo	Chile
Bay Area	US Gov	Montreal	Amsterdam	Zurich	S. Africa	UAE 2	Saudi Arabia 2	Osaka	Chuncheon	Hyderabad	Melbourne	Singapore	Belo Horizonte

An economical and powerful foundation

Key resources	Performance	Price	Reliability
Compute	Most powerful instances available	Up to 52% less than the competition	The only Compute Manageability SLA
Networking	Up to 2x25 Gbps from the instance No oversubscription	Lowest price for getting or serving your data	The only Networking Performance SLA
Block Storage	Superior IOPS/GB Superior IOPS/instance	Lowest price for guaranteed performance (up to 7,900% less)	The only Block Storage Performance SLA
Database	Superior TPS	Superior price / performance	RAC / Exadata 99.995% with Autonomous Database

Compute

Wide range of compute services for any enterprise use case

Bare Metal	VMs	Containers	Functions
Instance isolation Highest IOPS High throughput Low latency	Security- hardened hypervisor Flexible sizing Dense IO and Dedicated host option	Bare metal performance Self-healing clusters	Pay per use Serverless Container-native Open source

AMD EPYC	Intel Xeon	NVIDIA GPUs

Local Attached Storage	Remote Attached Storage
NVMe SSDs Up to 51.2 TB Millions of IOPS	NVMe Block Volumes 32 TB / volume 60 IOPS / GB

Storage

Comprehensive, best-performing storage services for enterprise workloads

Local	Block	File	Object
NVMe SSDs Up to 51TB Millions of IOPS 10-100 µs latency	NVMe SSDs 32 TB / volume 60 IOPS / GB <1ms latency	HA, distributed file system Start with KBs Scale to Exabytes	Distributed, HA Self-healing Unlimited scalability

Archive	Storage Gateway	Data Transfer
Durable object storage at 90% lower cost	Local NAS-like performance Configurable cache	Move petabyte scale data Option for appliance, disk No cost to transfer data

Networking

High fidelity virtual networks and connectivity

VCN

Fully configurable subnets, routing, firewalls
Default IPSec VPN
25Gb network infrastructure

FastConnect

Dedicated, SLA backed connectivity
No data transfer charges
34 carriers, 5 in Japan

Load Balancing

Choice of TCP, HTTP, HTTP/2
End-to-end SSL
TLS encryption

Service Gateway

Private access without traversing internet
Full range of IaaS/PaaS services covered

DNS

<30ms response time
Global load balancing
Traffic management
Network health checks

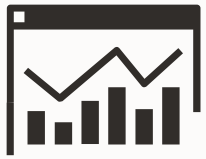

Governance & Management

Architected from the ground up for maximum isolation and protection

Access Control

Integrated IAM for all services

Simple role-based policies

Identity federation

Resource principals

Resource Governance

Flexible compartment structure

Built-in automation ensures tagging integrity

Cost Management

Cost analysis dashboard

Budgets

Resource quotas

Detailed, extensible usage reports

Cost tracking tags

Audit

Rich history of all events

Query API

Bulk export

Custom retention period

Monitoring

Fine-grained out-of-the-box metrics

Robust, custom metrics

Alarms

Notifications

Fully managed pub-sub

Built-in integrations for popular messaging protocols

Database

The most comprehensive, resilient, high performing database services

Customer managed

Oracle DB on Compute

VM/Bare Metal
Data Guard
Auto TDE

Semi-managed

Oracle Database Cloud Service

VM/Bare Metal
RAC
Data Guard
Auto TDE
Automated backup, patching

Oracle Exadata

Extreme performance
Base – Full rack
RAC
ADG
IORM
Cloud Service/
Cloud@Customer

Fully-managed

Oracle Autonomous Database

ADW

Auto-scaling
Auto-tuning
Auto-patching
Serverless
Spatial, ML
Document
SQL Developer

ATP

Auto-scaling
Auto-tuning
Auto-patching
Serverless or dedicated
APEX

Manageability

The future of comprehensive data platform services

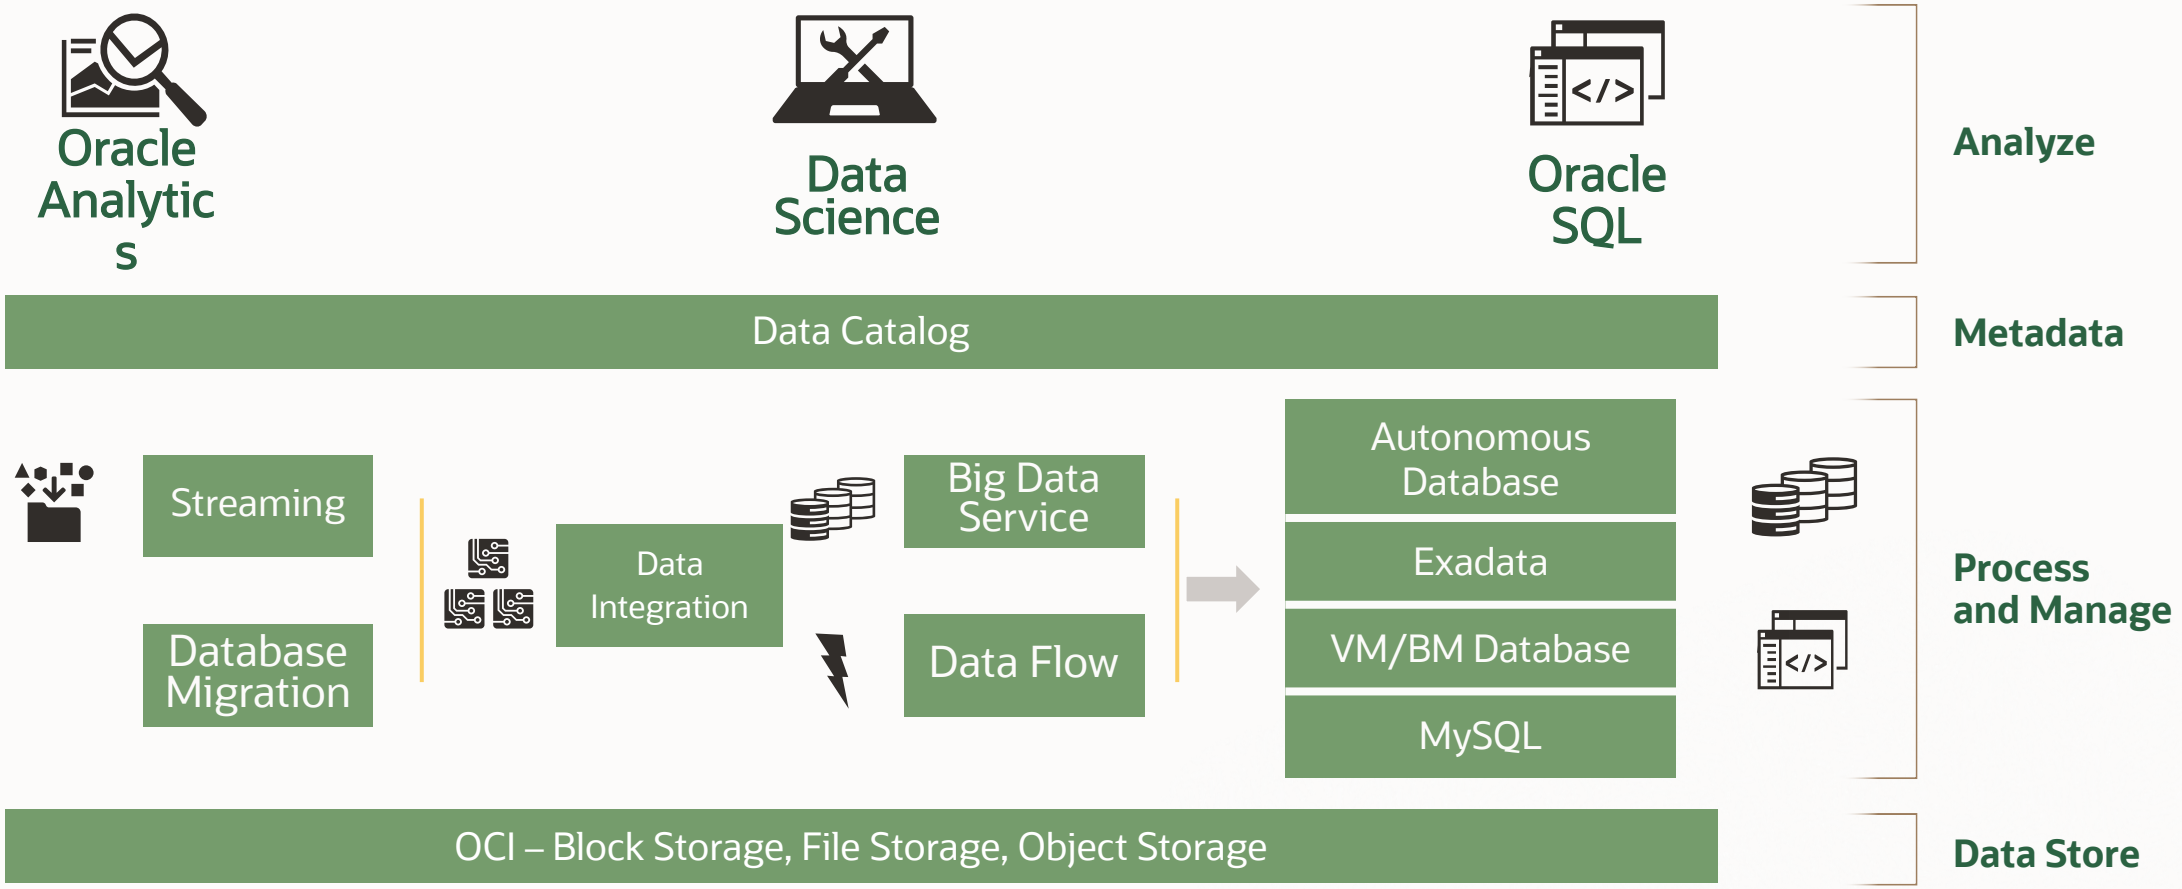

Oracle Cloud Native Services

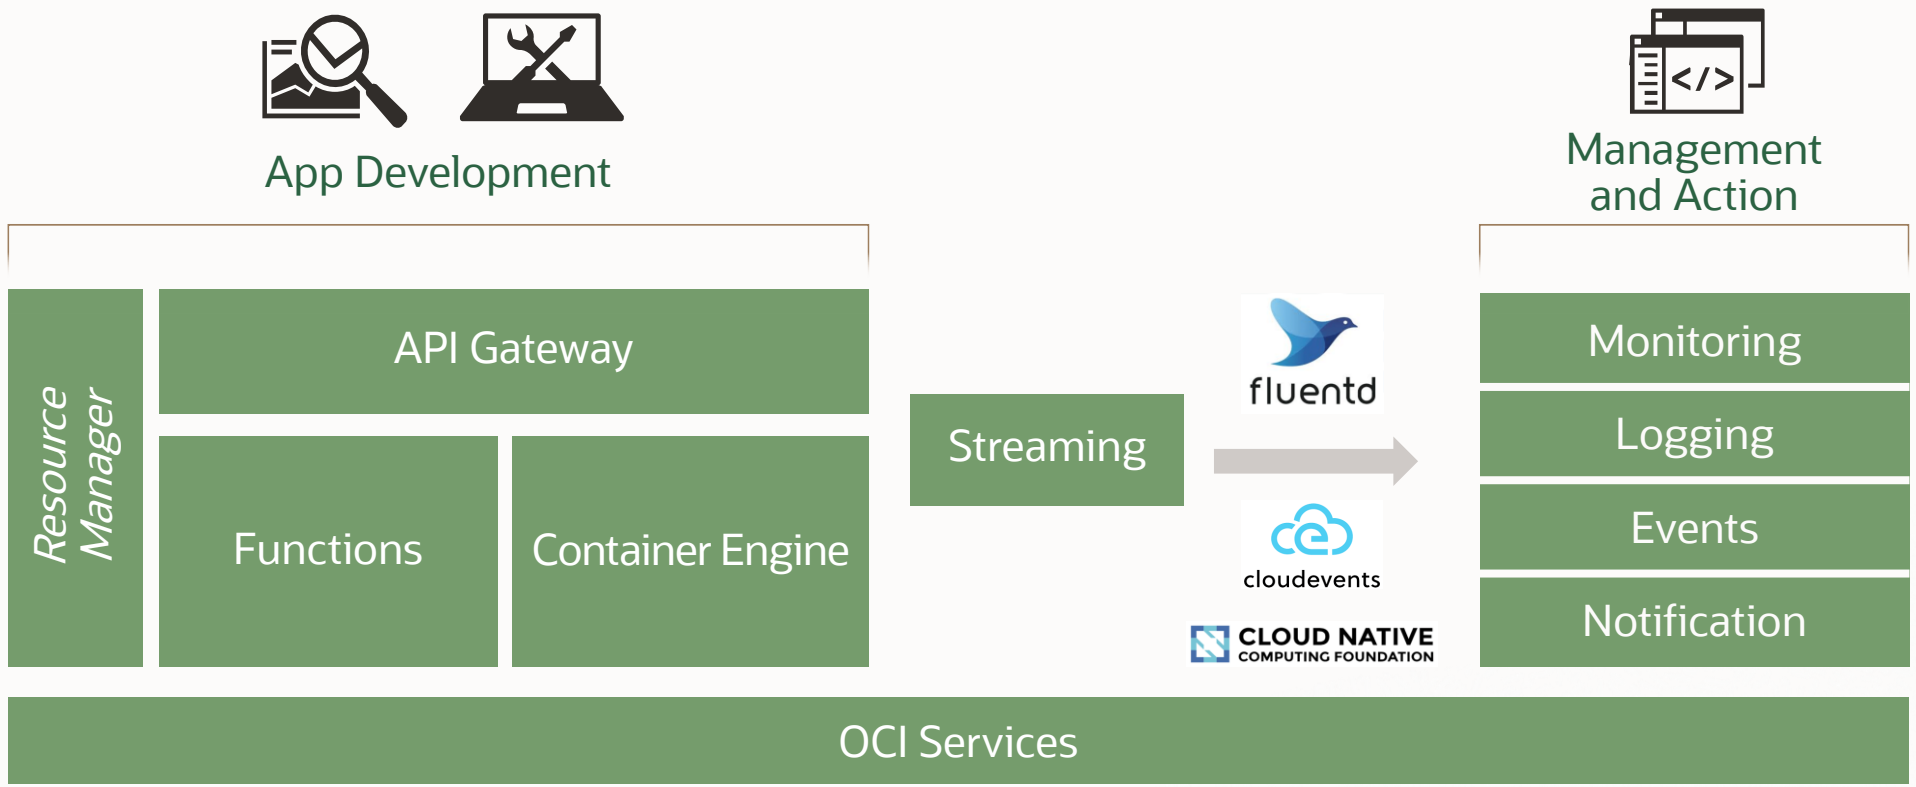

Container engine for Kubernetes and Registry

Cloud Native
- Standard Docker and Kubernetes
- Registry Integration
- Integrated with virtual cloud networking and storage

Developer Friendly
- Streamlined workflow
- Full REST API
- Built in cluster add-ons
- Open standards

Enterprise Ready
- Simplified Cluster Operations
- Full Bare Metal Performance and Highly Available IaaS
- Team Based Access Controls
- Autonomous Clusters

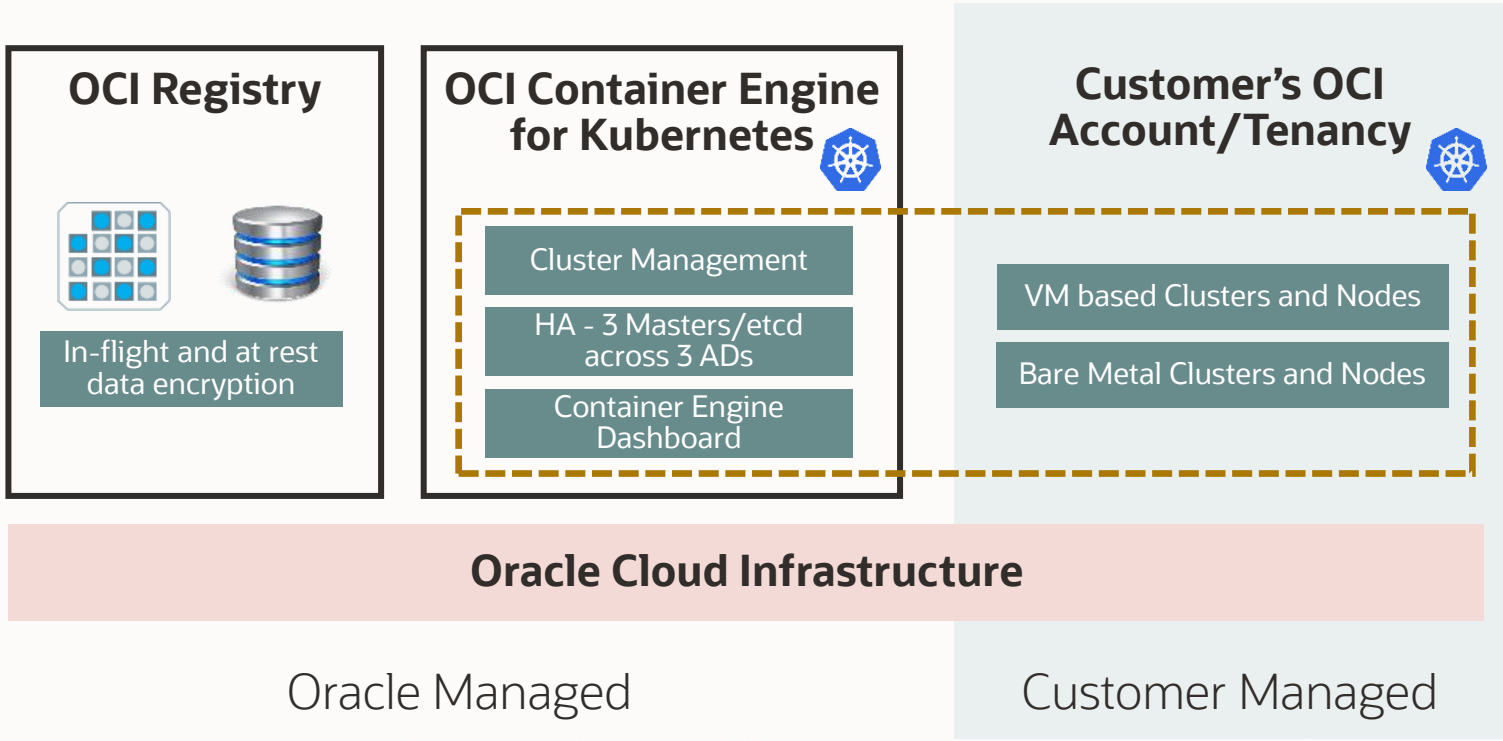

Physics laboratory migrates WebLogic infrastructure

- **Less focus on maintenance**, more on development
- Can move disaster recovery to the cloud
- **2 minutes to deploy** WebLogic clusters in Oracle Cloud Infrastructure Container Engine for Kubernetes instead of half day

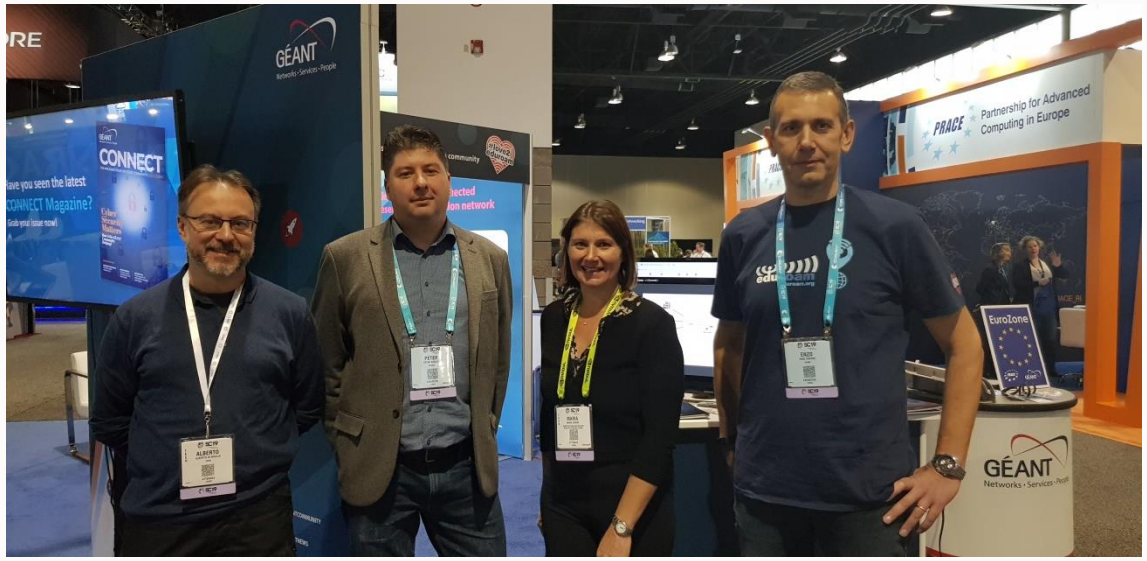

CERN openlab

"Time to deploy a cluster had decreased from a half day to 2 minutes"

1-10G FastConnect Service to OCI

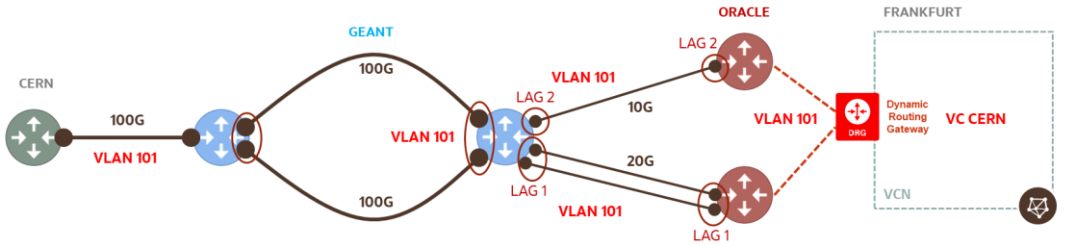

- Switch-over is needed between LAGs but automatic load balancing can be done between cross-connects in LAG 1
- Bandwidth aggregation is possible in LAG 1

- LAG 1 cross-connect group consists of two cross-connects

Oracle Cloud ecosystem

Certified Oracle and partner solutions ready to deploy at a click

Oracle	OS, Virtualization	Apps, Migration	Networks, Security	Data	Analytics, AI/ML, HPC
ORACLE	ORACLE Linux	blueprism	A10	Couchbase NoEQUAL	Altair
E-Business Suite	Microsoft	RMM Rackware	AVI Networks Now part of VMware	Microsoft SQL Server	compellon
GoldenGate	stromasys legacy server emulation	Solution-Soft	aviatrix	ORACLE Database	H2O.ai
JD Edwards	SUSE	ZERO DOWN SOFTWARE	CIS Center for Internet Security	SCYLLA.	kinetica
PeopleSoft	ubuntu		FORTINET	sesame SOFTWARE	PYRAMID ANALYTICS
Siebel	vmware		paloalto NETWORKS		
WebLogic			SHIELDIO		

Spend Universal Credits on partner solutions in our marketplace

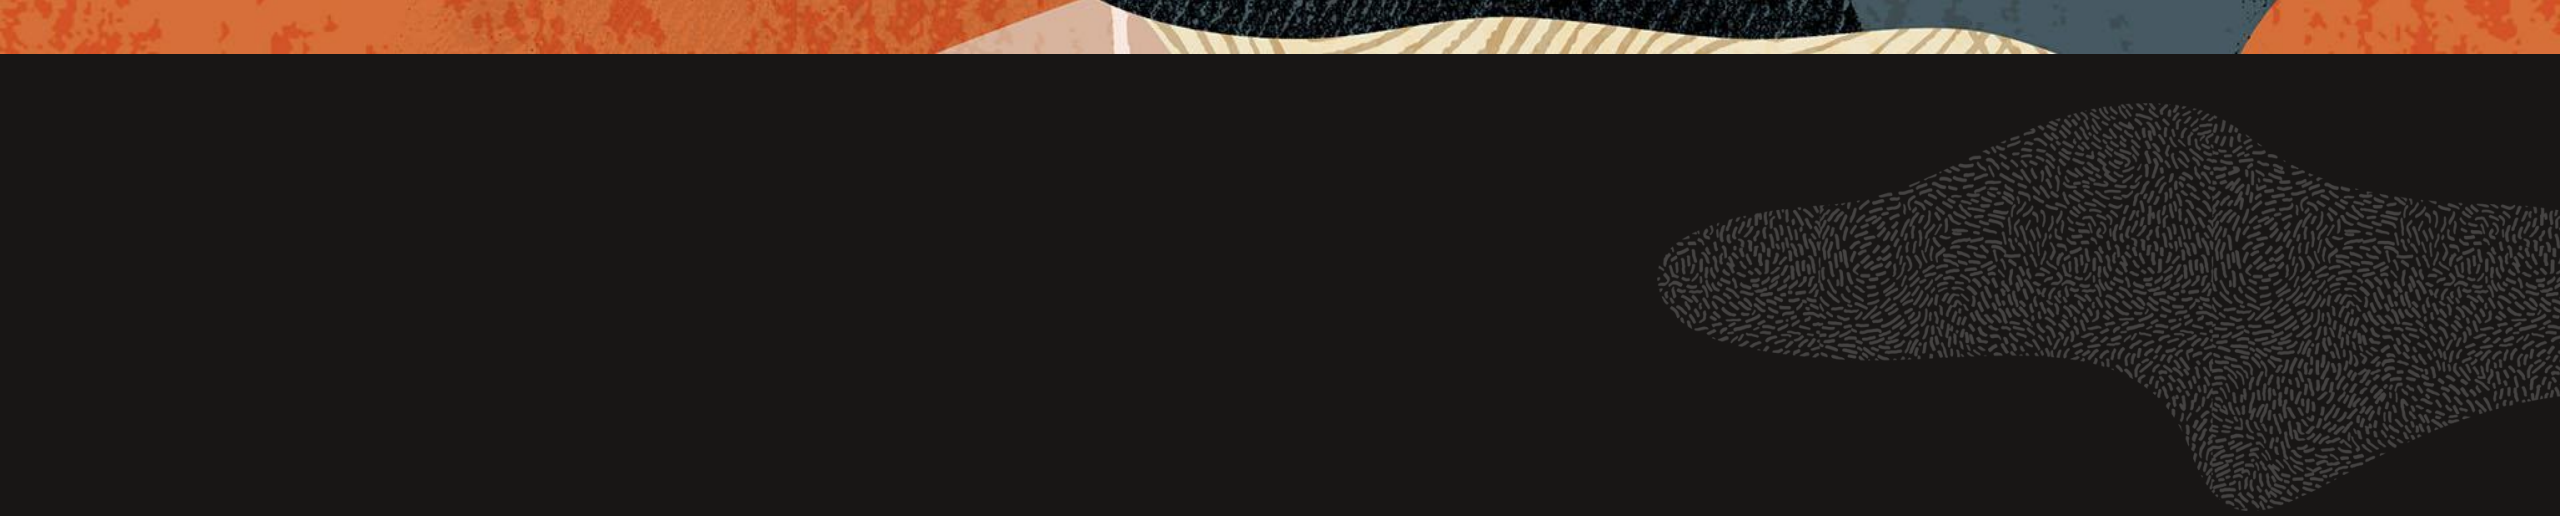

ORACLE + vmware®

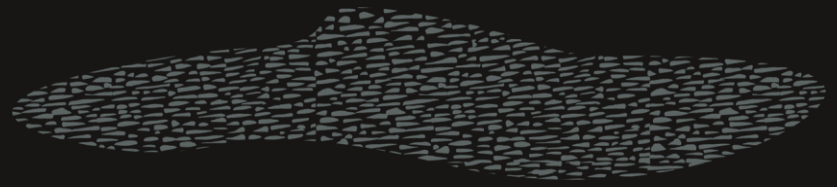

The partnership: Improving the experience of joint customers

- ✓ Oracle is developing a new product: Oracle Cloud VMware Solution
- ✓ Seamlessly migrate and extend VMware fleet to Oracle Cloud
- ✓ A familiar experience with full configurability and management
- ✓ Oracle joins VMware Cloud Provider Program
- ✓ Oracle Cloud VMware Solution will be sold by Oracle and Oracle partners
- ✓ Oracle will provide support for VMware on Oracle Cloud Infrastructure
- ✓ Oracle will also support Oracle Database and applications deployed on VMware

Diagram of Oracle Cloud VMware solution

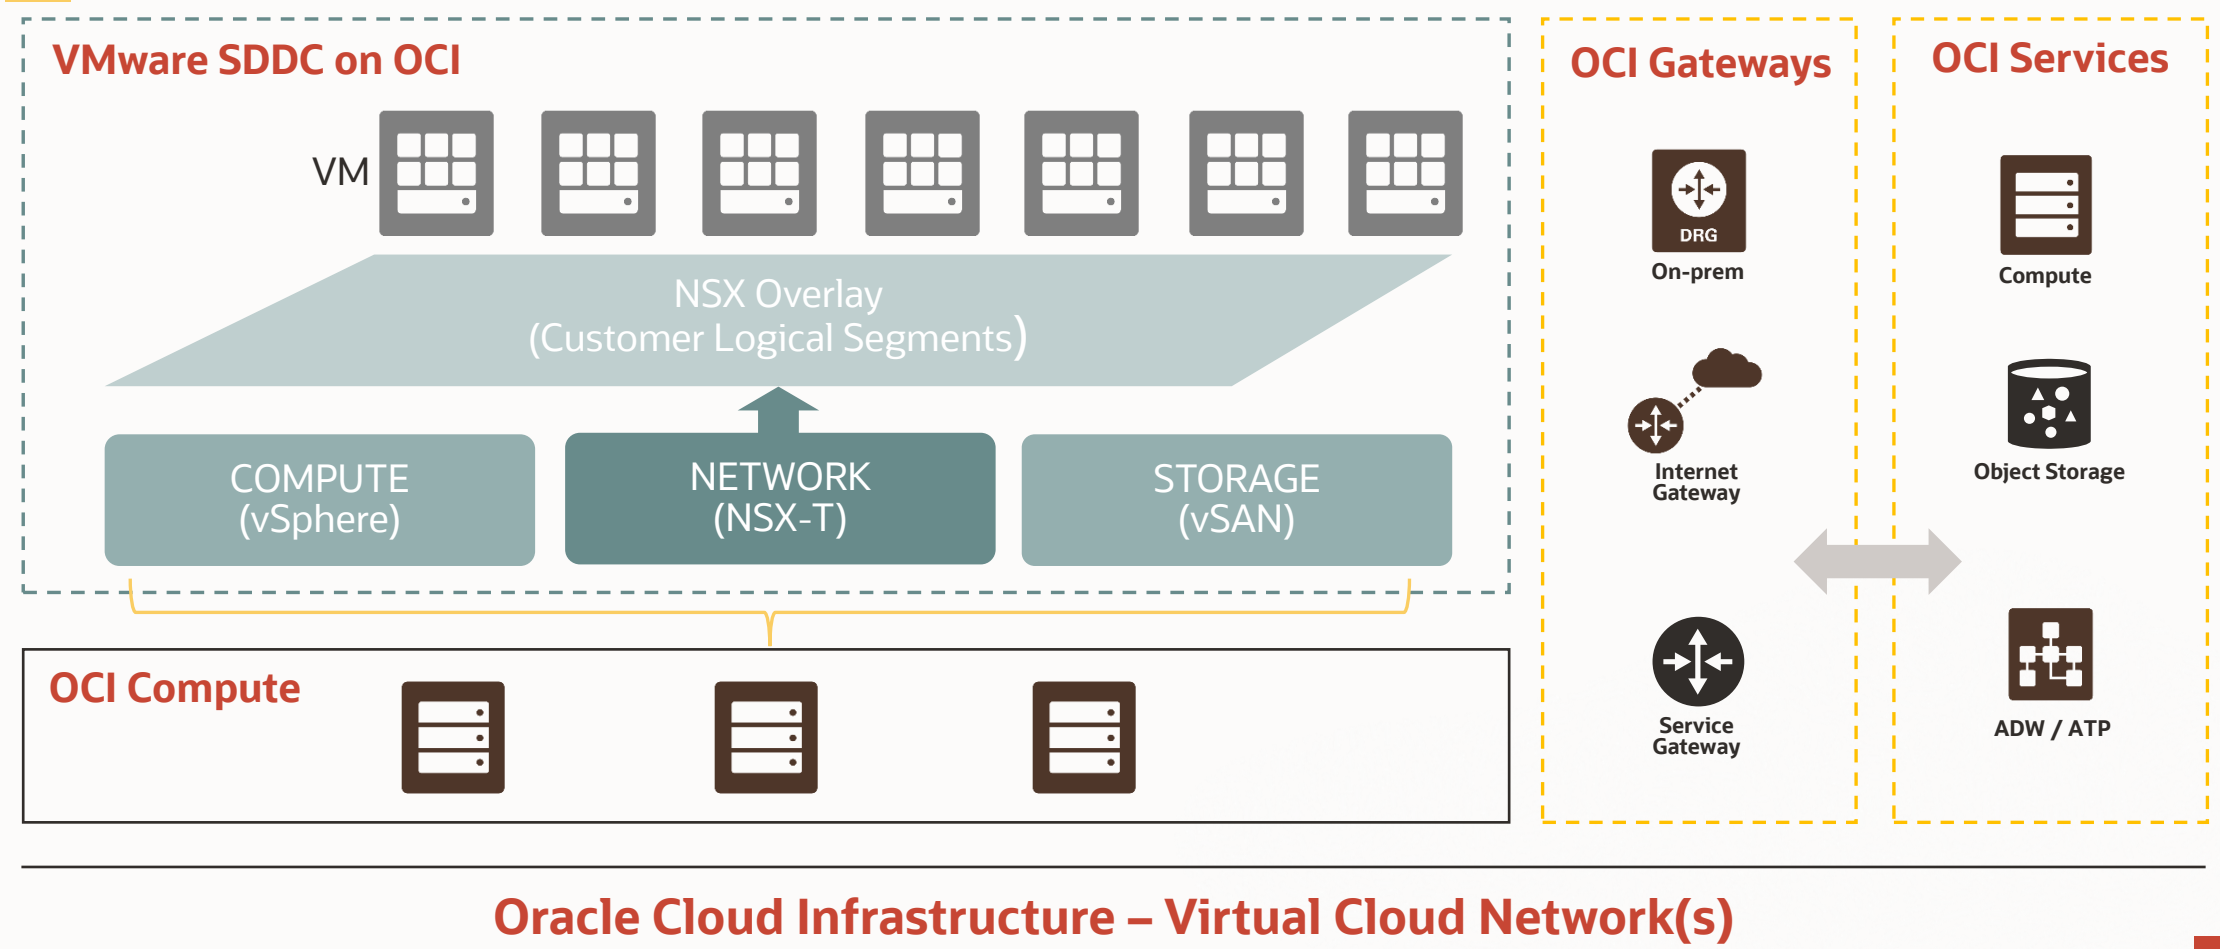

Use cases

Datacenter extension

Extend your Data Center foot print to the cloud, using VMware on one of the growing list of Oracle Cloud Infrastructure regions.

Disaster recover on Oracle Cloud

Enhance your Business Continuity strategy, by creating a Disaster Recovery environment for VMware on OCI

vmware®

+

ORACLE Cloud Infrastructure

Datacenter migration

Seamless Migration of your VMware workloads to Oracle Cloud, without the need for re-architecting your applications.

On-demand capacity

Spin-up additional compute capacity as you need it, when you need it.

ORACLE + Microsoft

Interconnected Multi-Cloud Solutions for Enterprise

- ✓ Microsoft Azure and Oracle Cloud are interconnected today, so you can migrate and run mission-critical enterprise workloads across clouds
- ✓ Unified identity and access management via single sign-on with automated user provisioning to easily manage resources across clouds
- ✓ Collaborative support of custom and Oracle Applications on Azure with Oracle Database on Oracle Cloud – connect best-in-class services across clouds

  - Available Now: US East, London
  - Coming Soon: US West, Government, Asia, and Europe regions

Enterprise Cloud Interoperability Partnership

- Oracle Cloud Infrastructure
- Oracle Autonomous Database
- Oracle Exadata
- Oracle Applications
- Oracle RAC
- Oracle Analytics Cloud
- And other services…

- Azure DevOps
- Azure Stream Analytics
- Azure Databricks
- Azure Cognitive Services
- Azure IoT Hub
- Azure Kubernetes Service
- And other services…

Oracle-Azure partnership use cases

Cross-cloud SSO and interconnect

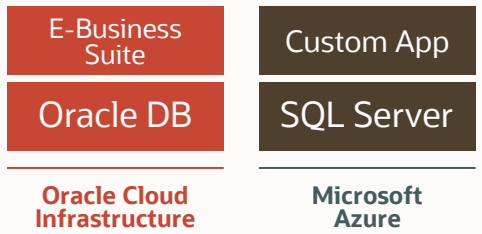

- Oracle Apps (PSFT, JDE) on Azure using Oracle Database on OCI
- Custom .NET application on Azure using Oracle Database on OCI
- Custom Cloud Native applications on Azure using Oracle Autonomous DB

Cross-cloud SSO and interconnect

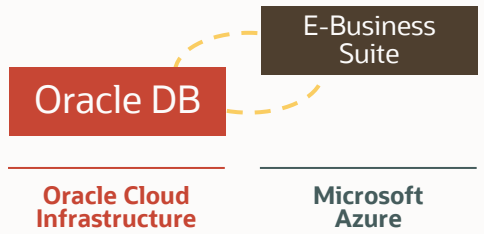

- Full stack Oracle or custom applications on Oracle Database on OCI and Full stack apps on Azure that interoperate and share data

Cross-cloud SSO and interconnect

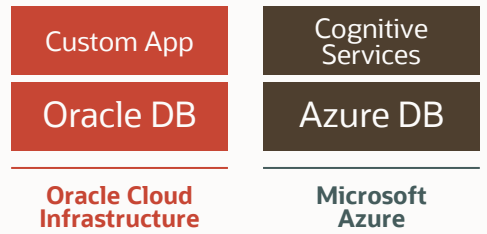

- Applications/Oracle Database in OCI, Azure Data Lake for analytics & Cognitive Services for AI
- SQL Azure, SQL Server, SQL DW on Azure and Oracle Analytics Cloud, Data Science service on OCI

What do we define as HPC?

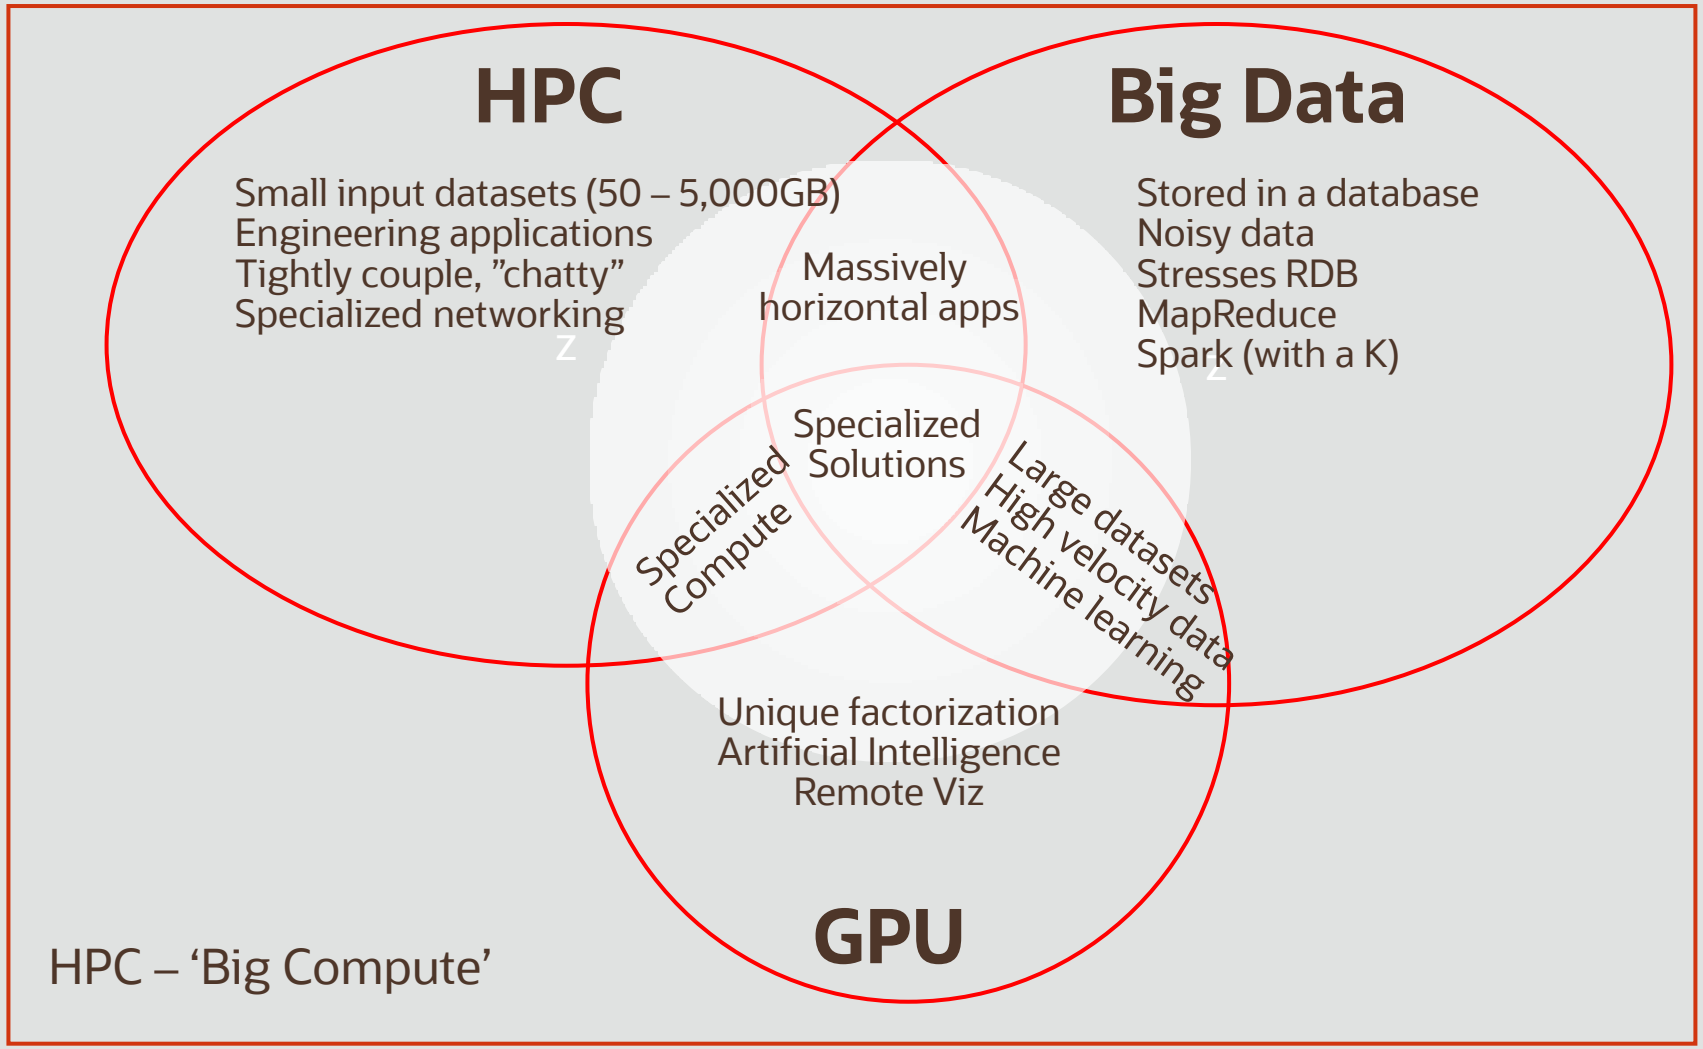

Who is investing in HPC today?

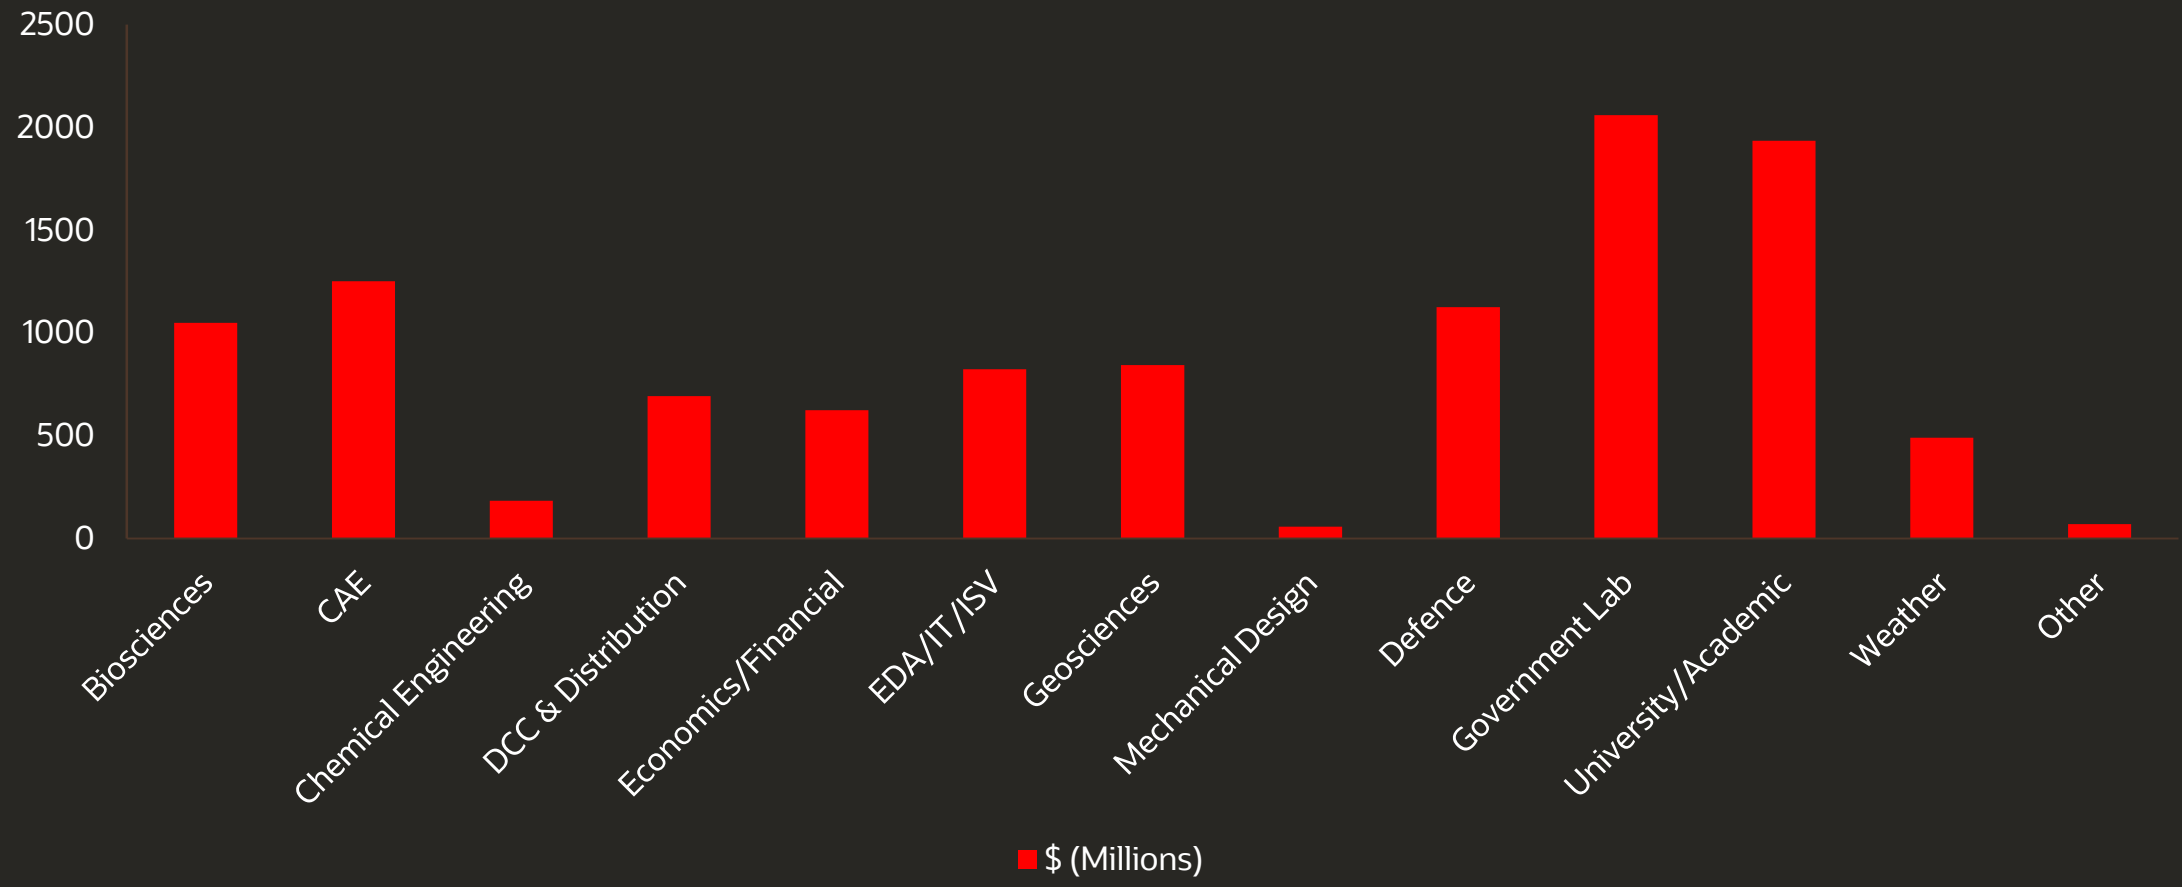

ORACLE

Reference -* IDC 2016

3.7 GHz, providing the highest performing processor in the cloud

36 Cores of Intel 6154 High Performance

384 GB of memory and 6.4 TB of local storage

20k over 20,000 cores in a single cluster

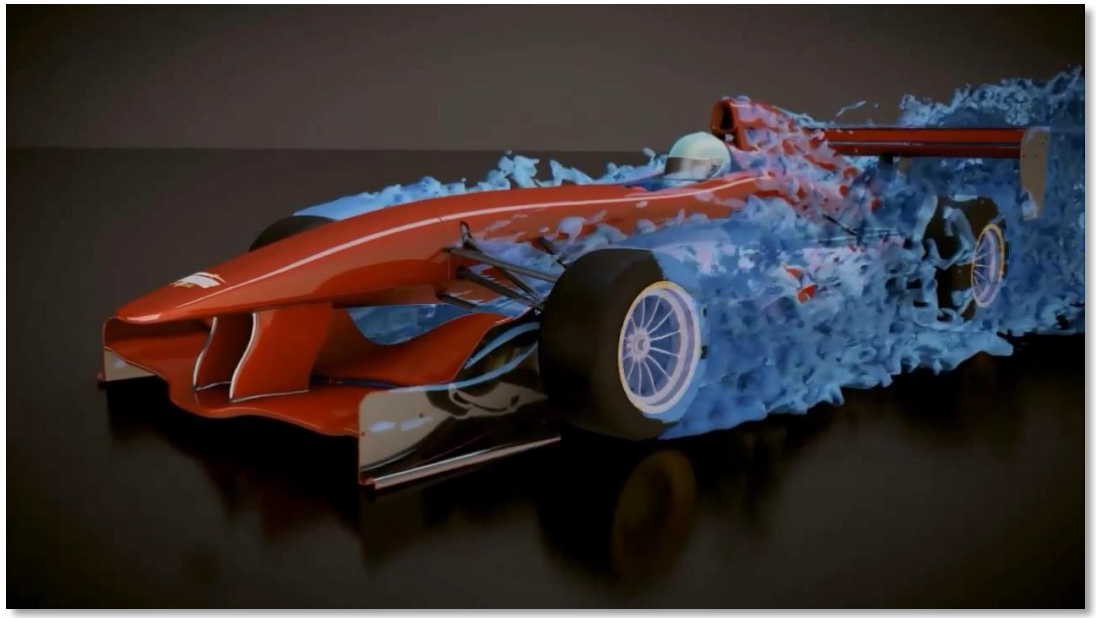

intel® HPC

64 Physical threads on a single machine
128 total threads

3 ¢ per core per hour

most cost efficient physical processor

↑ Highest memory bandwidth for cloud

BM.Standard.E2.64 best price performance in the cloud

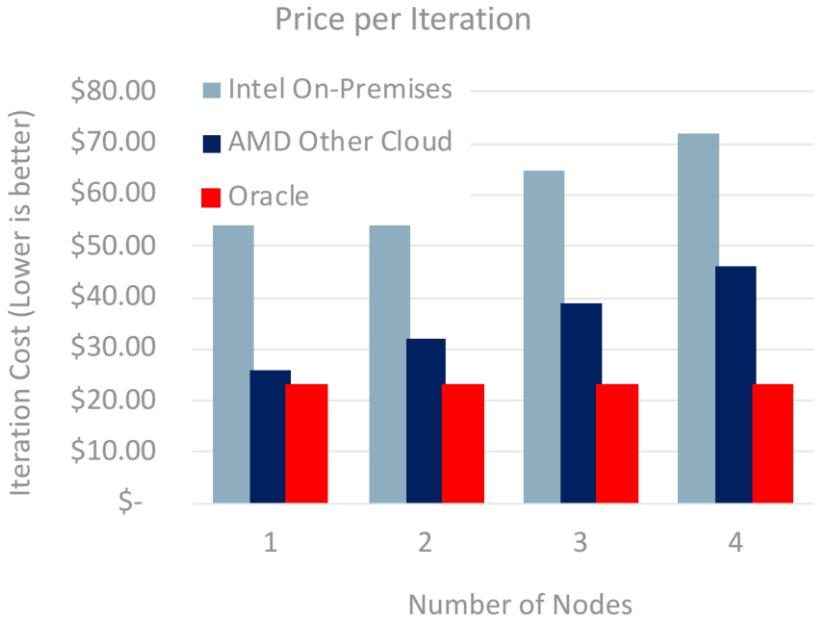

8 BM.GPU3.8 – 8 Way NVIDIA Volta SXM2, NVLINK, HGX-1 Open Compute 52 Skylake CPU cores, 50 Gbps

2 BM.GPU2.2 – 2 Way NVIDIA Pascal Perfect for Visualization or Inferencing

1 Attach up to 1 PB of NVME backed block storage to “Feed the Beast”

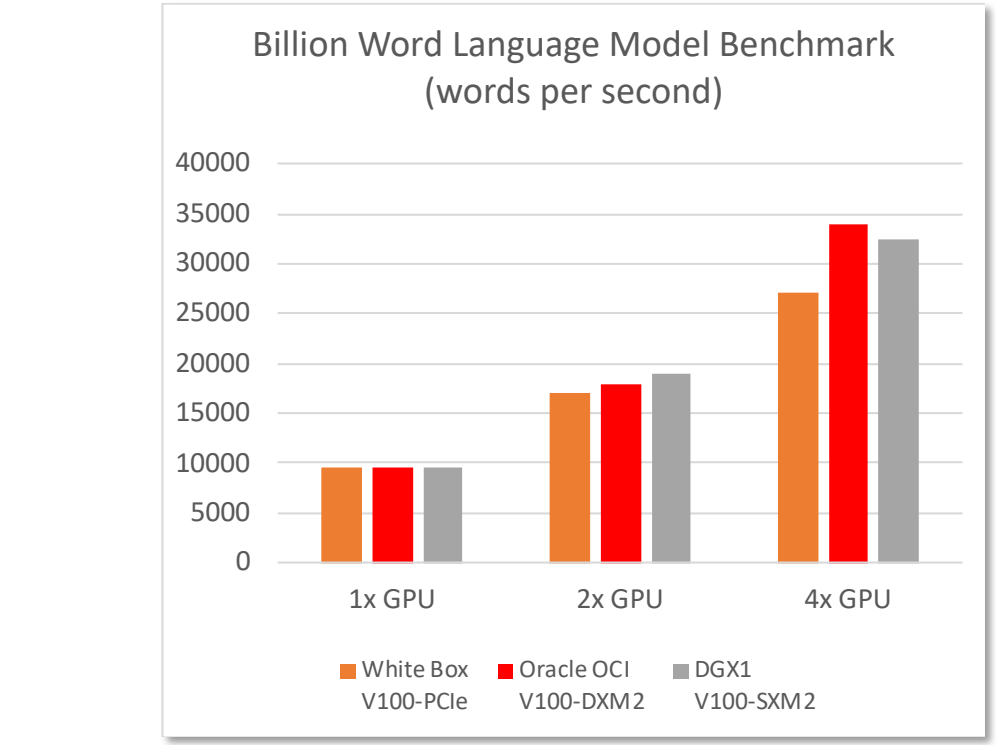

1.5 µs latency, over ROCEv2 - on-premises performance

12 GB/s transfer rate between nodes over the RDMA network, bare-metal application performance

125 Gbps full line rate connectivity

The only hardware point-to-point network connection in the cloud

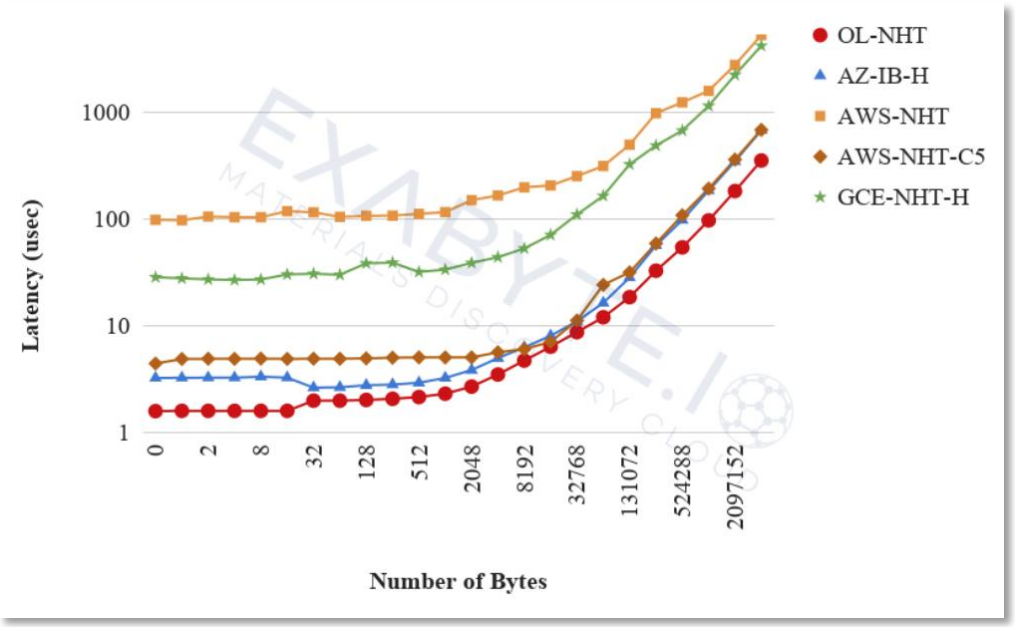

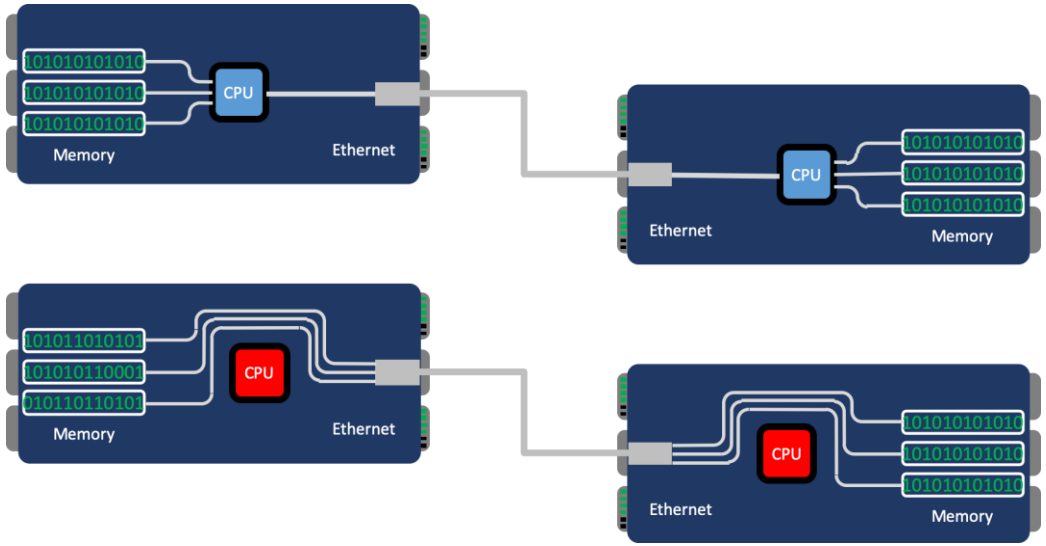

How are we different?

- Bare Metal HPC, on-premises like performance
- No 'jitter' in our HPC hosts or on our network
- Secure, isolated RDMA network with no oversubscription
- More than 20,000 cores in a single RDMA cluster
- Latency under 2µs vs. 15 µs on AWS

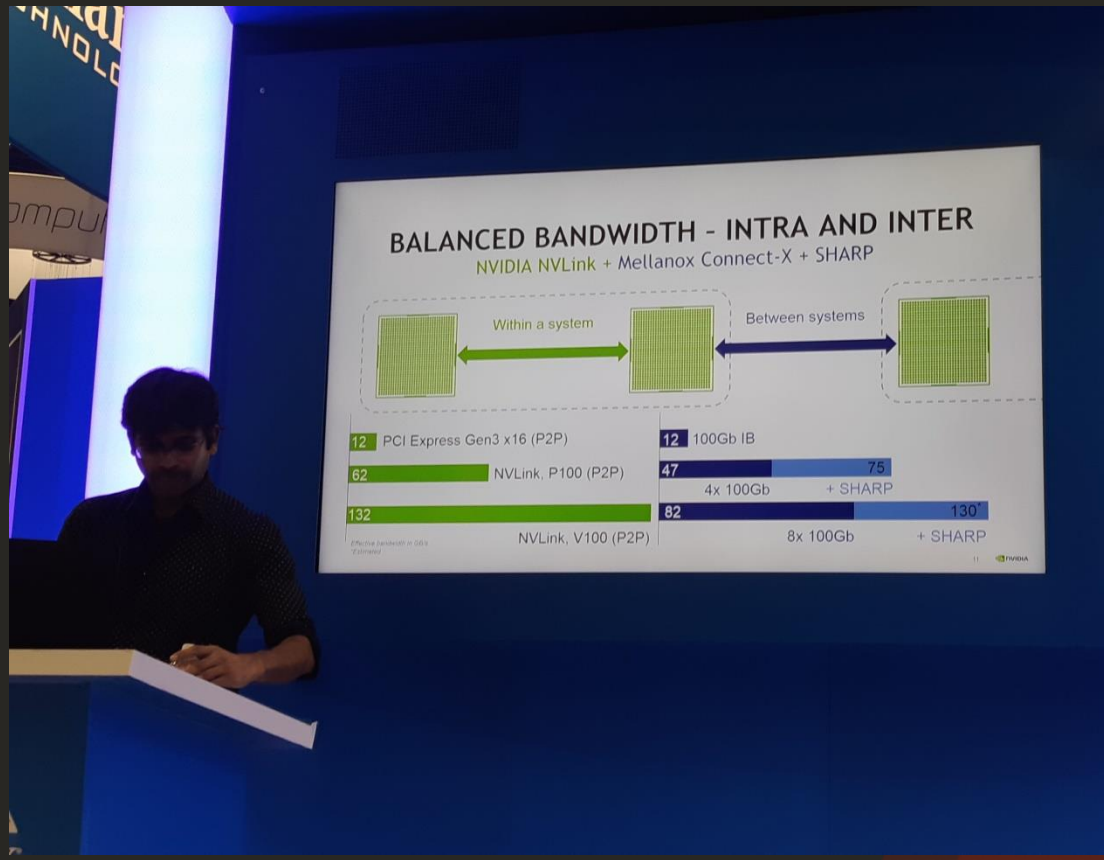

Performance

Latency

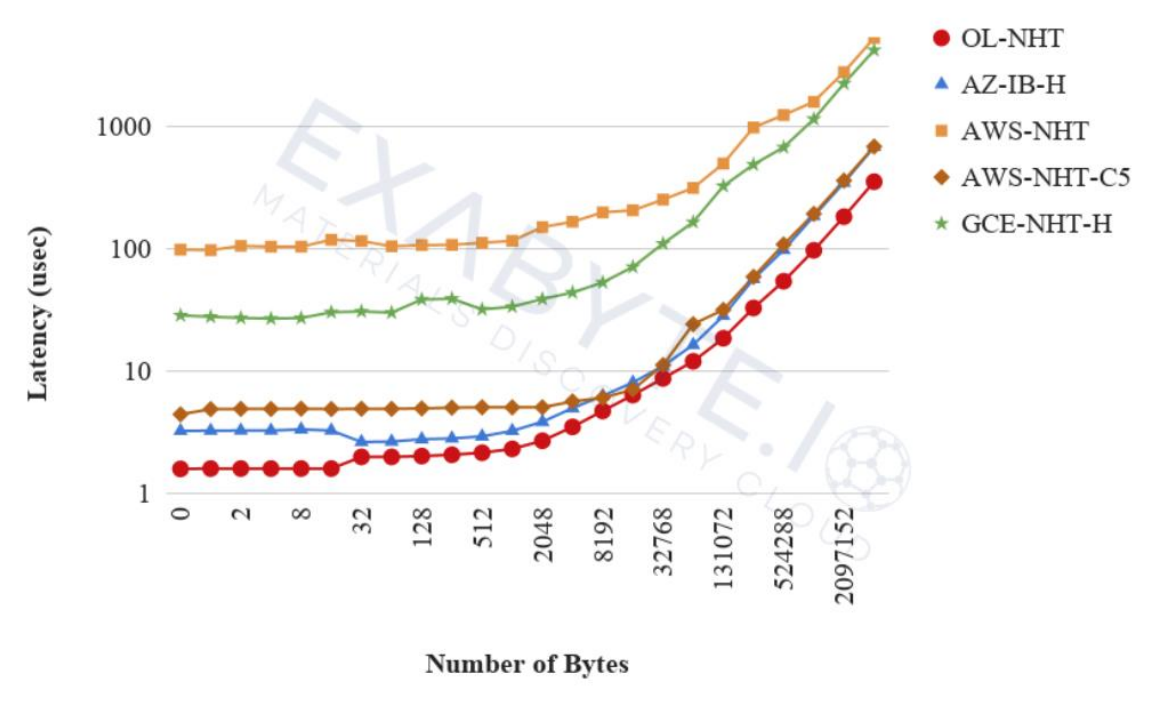

Application Performance

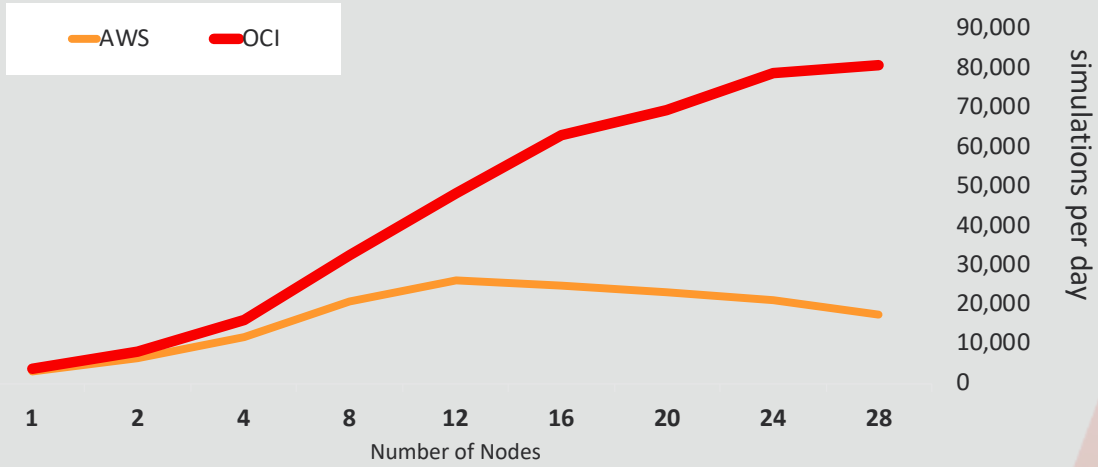

Simulation Tools Ecosystem

Visual Effects Rendering

Artificial Intelligence & Deep Learning

Manufacturing Automotive Aerospace

Open Source HPC Applications

Simple, Rapid Deployment

3 Deployment Models

Command Line

Use Command Line and API's for integration to automated workflows

Stacks

Use Resource Manager Stacks for interface based Terraform Deployment

Marketplace

Use Marketplace for deploying partner provided infrastructure and applications

Objective

&

ORACLE

The collaboration between GÉANT and Oracle aims towards providing all NRENs and connected institutions with Cloud solutions which are innovative, complete, supported and affordable to guarantee a smooth journey into the Cloud.

Oracle's Educational Quotation

- A preferential quotation, negotiated and rubberstamped by GEANT, based on Oracle's standard terms and conditions, available for GEANT Members and their connected institutions.
- Not covered by EU public procurement tender!
- Must be used with caution. Full compliance with the EU/National public procurement rules and regulations is the responsibility of the purchasing institution.

Benefits via the GEANT Cloud Catalogue

- Governmental pricing and volume discounts apply to all GÉANT members and their connected institutions.
  - Pre-approved for the services listed in the GEANT Cloud Catalogue
- Oracle enables single-sign-on with institutional accounts via federated access using SAML2 protocol.
  - Oracle's access management platform is also federated with Microsoft Azure.
- Oracle offers direct network interconnection, for both public and private FastConnect peering, with GÉANT.
  - Currently tested by CERN, other institutions can be on-boarded later

Pricing Model Options

Universal Credits

Pay As You Go (PAYG)	Monthly Flex
• No upfront commitment • Pay only for what you use • Pay in arrears based on usage	• 1 year minimum term • Agreed to monthly spend
• List Price • Built for land and expand • Best when usage is uncertain • Elastic payments based on usage	• PaaS savings vs PAYG start at 33% • Additional discounts based on size of deal and term of deal • Predictable spending • Spend more, save more

Governmental

- Multi-year term
- Agreed to monthly consumption
- Separate named SKUs

- No transfer of credits between SKUs to allow named services
- Commitment to hourly consumption (744h/month) to prevent credit loss
- Stop creating new instances to prevent overage
- Volume discounts

Partnering with Academia

- Oracle for Research
  - offers researchers, scientists and university-associated innovators access to Oracle Cloud technology and a global community working to address complex problems and drive meaningful change in the world. This program provides free Oracle cloud credits and technical resources to qualified applicants.
- Oracle Labs
  - seeks to identify, explore, and transfer new technologies; investing in research collaborations with faculty, research directors, and principal investigators at universities, labs and non-profit research organizations worldwide.
- Oracle Academy
  - advances computing education globally by offering a complete portfolio of education resources for teaching and not-for-profit academic research to more than 6 million students in 128 countries.

** Internet2, EDUCAUSE, EUNIS, HEUG, IMS Global, PESC …*

ORACLE

Summary

*ORACLE Cloud **completes** the GÉANT Cloud catalogue perfectly. ORACLE Cloud gives all existing R&E customers the **opportunity** to start and **succeed** their journey into the Cloud, **no matter** what the **workloads** or the **deployment preferences are**. With ORACLE´s **know-how**, the R&E customers get a professional **support and service** on a **technological** and **customer relationship level**.*

ORACLE

Visit https://clouds.geant.org/oracle/

Oracle Education and Research

https://www.oracle.com/industries/education-and-research/