

ORACLE

# Oracle Cloud Infrastructure

HPC-Services in der Oracle Cloud & OCRE



**Dr. Rita Engemaier**

Account Executive Cloud Technology  
Hochschulen und Forschung  
Rita.Engemaier@oracle.com



**Werner Bauer**

Principal Cloud Engineer  
Werner.Bauer@oracle.com



**Wolfgang Dreyer**

OCI Specialist Director, EMEA CoE  
Wolfgang.Dreyer@oracle.com




# Agenda

1. HPC on Oracle Cloud Infrastructure (OCI)
2. OCI GPU Capabilities
3. High Bandwidth, Low Latency RDMA Cluster Networking
4. « Cluster per job » strategy | use case scenarios & references
5. OCRE (Open Clouds for Research Environment) & GÉANT
6. Q&A



# HPC on Oracle Cloud Infrastructure (OCI)



## Ice Lake processors

- BM HPC 3.36  
0,075 USD core Hour
- Cores: over 20k cores using Cluster Network
- 200G bandwidth & RDMA support with under 1.5  $\mu$ s latency
- NVMe storage for local MPI checkpointing
- Demanding workloads (i.e. CFD, EDA, Crash simulations) expected to deliver 30% more performance compared to existing X7 generation of instances at same price per core
- **Available NOW**
- Announcement [video](#)

## Next Generation AMD Instances

- E4.128
- 0,025 USD core Hour
- Powered by AMD's Milan generation processors
- BM or VM shapes available via flexible shapes
  - Pick and choose core and memory split based on workload characteristics
- Available through cloud native services
- **Available NOW**
- Announcement [video](#)

## ARM Compute Instances

- Powered by Ampere's Altra Processors
- Cores: BM or VM shapes – very dense, up to 160 cores single threaded performance (3.3Ghz/core)
- Available through cloud native services
- Best price performance for general purpose/ massively parallel workloads with up to 30% performance improvement over other x86 compute instances per core
- **Available NOW**
- Announcement [video](#)

## GPU bare Bare-Metal Instances

- NVIDIA Ampere Architecture
- Both Bare-Metal and VMs
- GPUs: up to 8 (for BM) A100 Tensor Core GPUs
- Memory: Up to 2 TB
- 1.6 Tbps RDMA
- For intensive applications like genomics, AI deep learning training and inference, data analytics, scientific computing, edge video analytics and 5G services, graphics rendering, cloud gaming...
- Announcement [video](#)

## File Server options

- Get the performance you need
- Choose from a wide set of parallel file systems in the market place: IBM Spectrum scale, BeeGFS, Lustre, Quobyte and more
- Achieve 60 -290 GB/sec throughput for HPC parallel file systems
- Unmanaged
- Oracle ranked N°7 in the [IO500 benchmark](#)



# OCI GPU Capabilities




## **NVIDIA Pascal P100 16GB**

Up to 2 GPUs per Instance

1x GPU via Virtual Machine  
2x GPU via Bare Metal

28 CPUs, 192GB, 50Gbps




## **NVIDIA Volta V100 16GB**

Up to 8 GPUs per Instance

1x, 2x, 4x GPU via Virtual Machine  
8x GPU via Bare Metal

52 CPUs, 768GB, 50Gbps



## **NVIDIA Ampere A100 40GB**

Up to 8 GPUs per Instance

8x GPU via Bare Metal

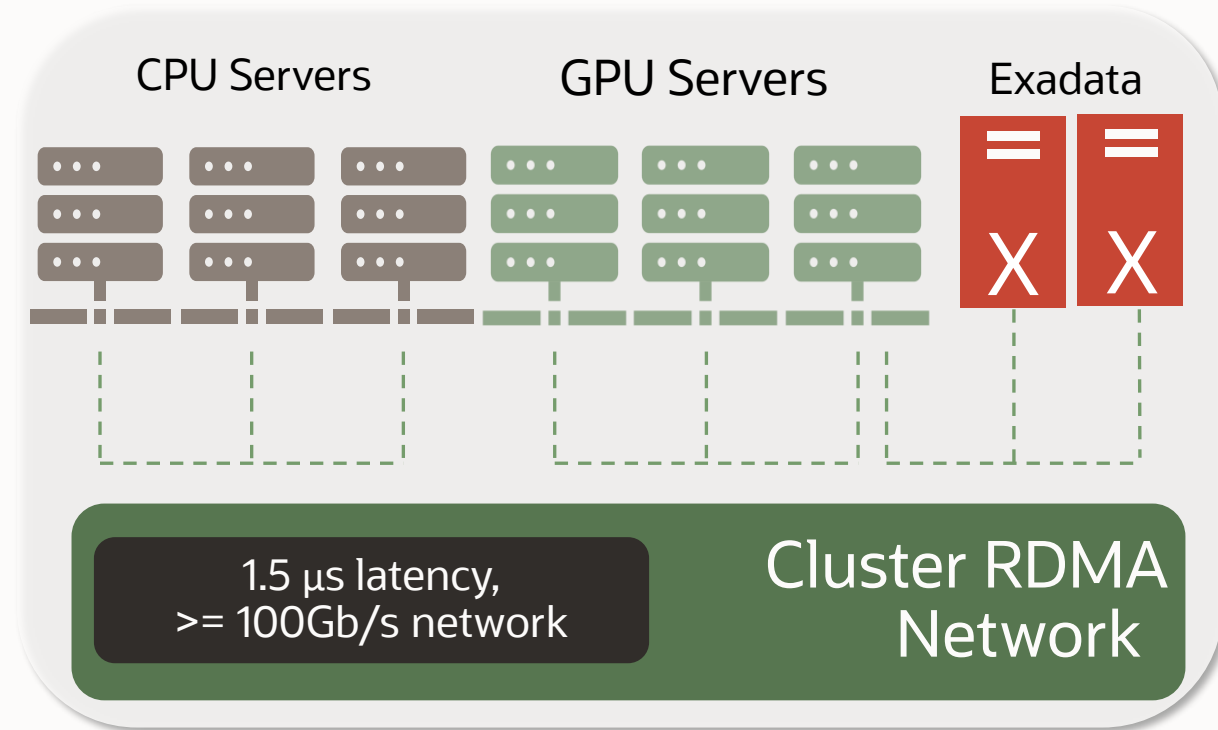
64 AMD CPUs, 2048GB,  
1600Gbps

4 x 6.4 TB SSDs



# High Bandwidth, Low Latency RDMA Cluster Networking

- Oracle connects the servers directly to the RDMA switch
- Up to 20,000 HPC cores in a single RDMA cluster
- Up to 512 GPUs in a single RDMA cluster
- No hypervisor, no virtualization, no jitter bare metal HPC
- 1.5  $\mu$ s latency,
- Local NVME disc for Temp or aggregated Filesystem usage
- 100Gb/s for HPC,  
8 \* 200Gb/s for GPU  
X9 will come with 200Gb/s



# Oracle Cloud Infrastructure – built for complex enterprises with sustainability in mind




**100% renewable energy** use at Oracle Cloud data centers in **Europe**

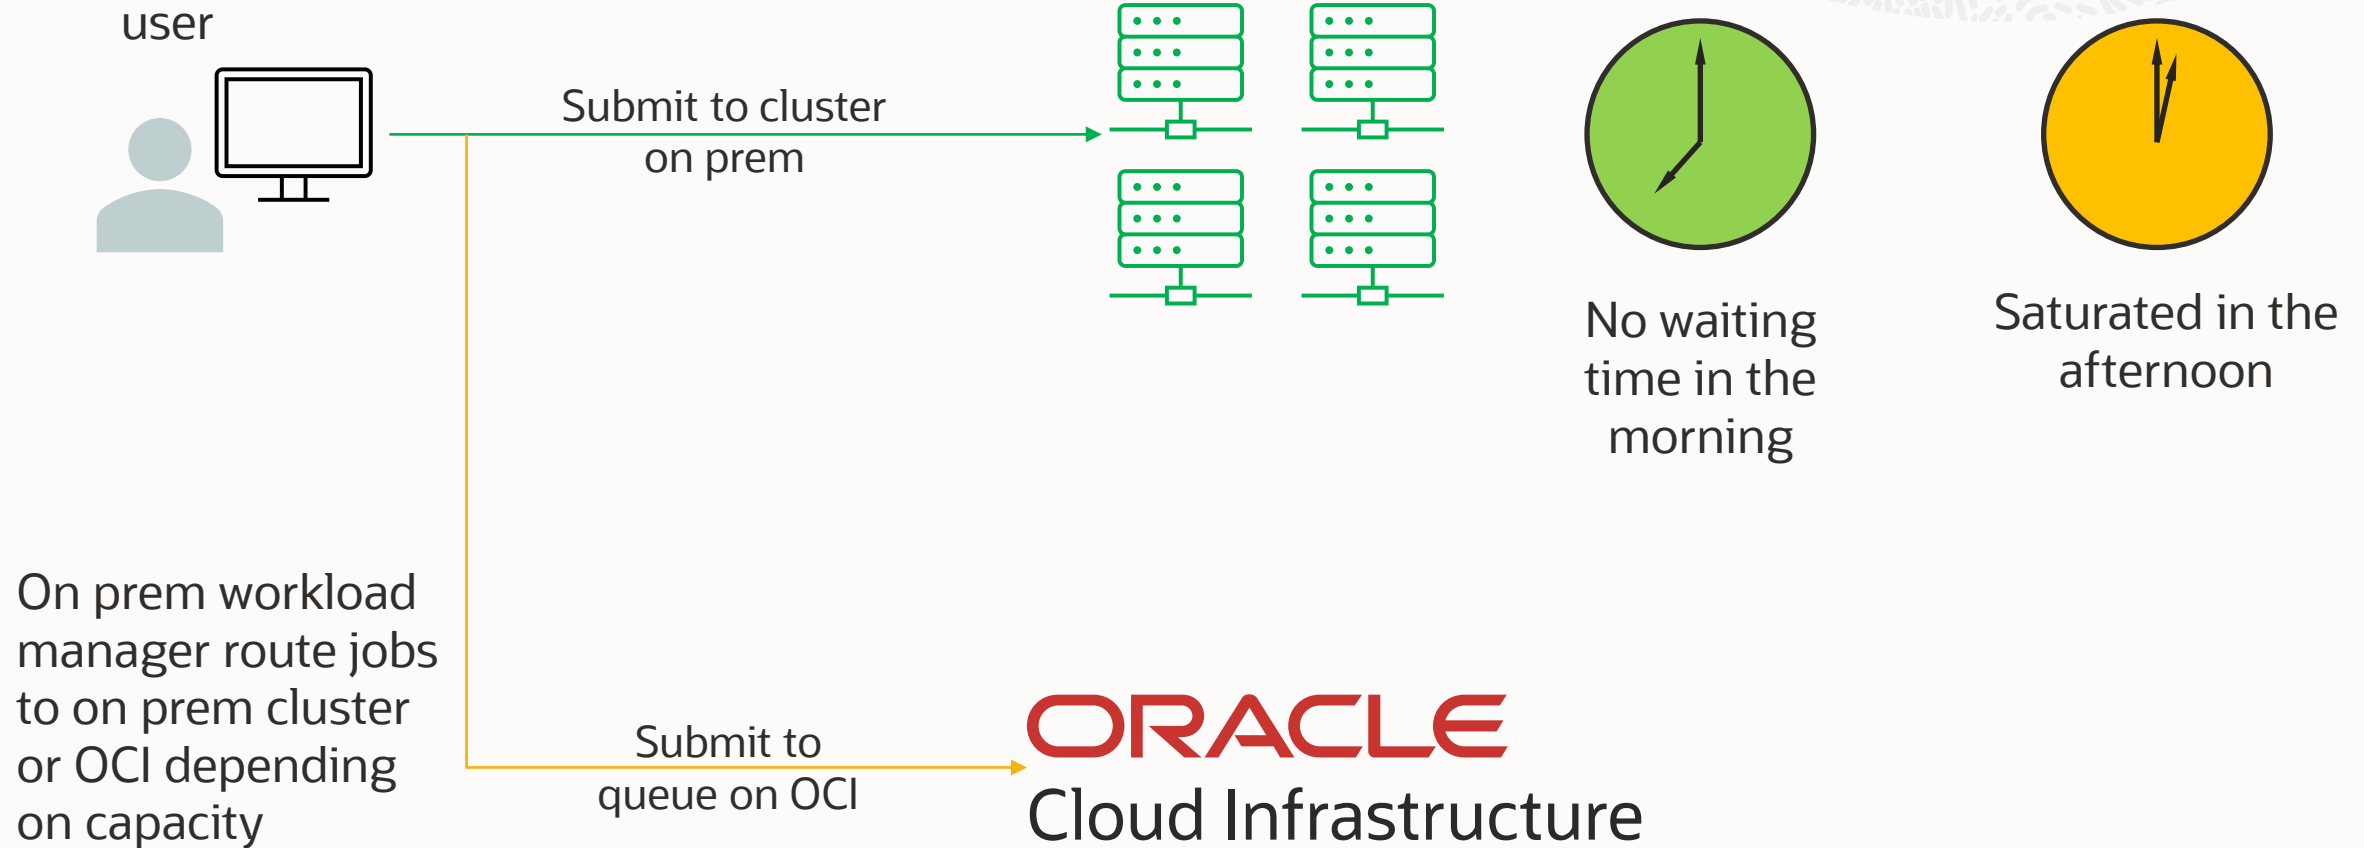
London - 100%, Amsterdam - 100%,  
Zurich - 100%, Frankfurt - 100%



**99.6%** of retired hardware collected for recycling, reuse or manufacturing in FY20

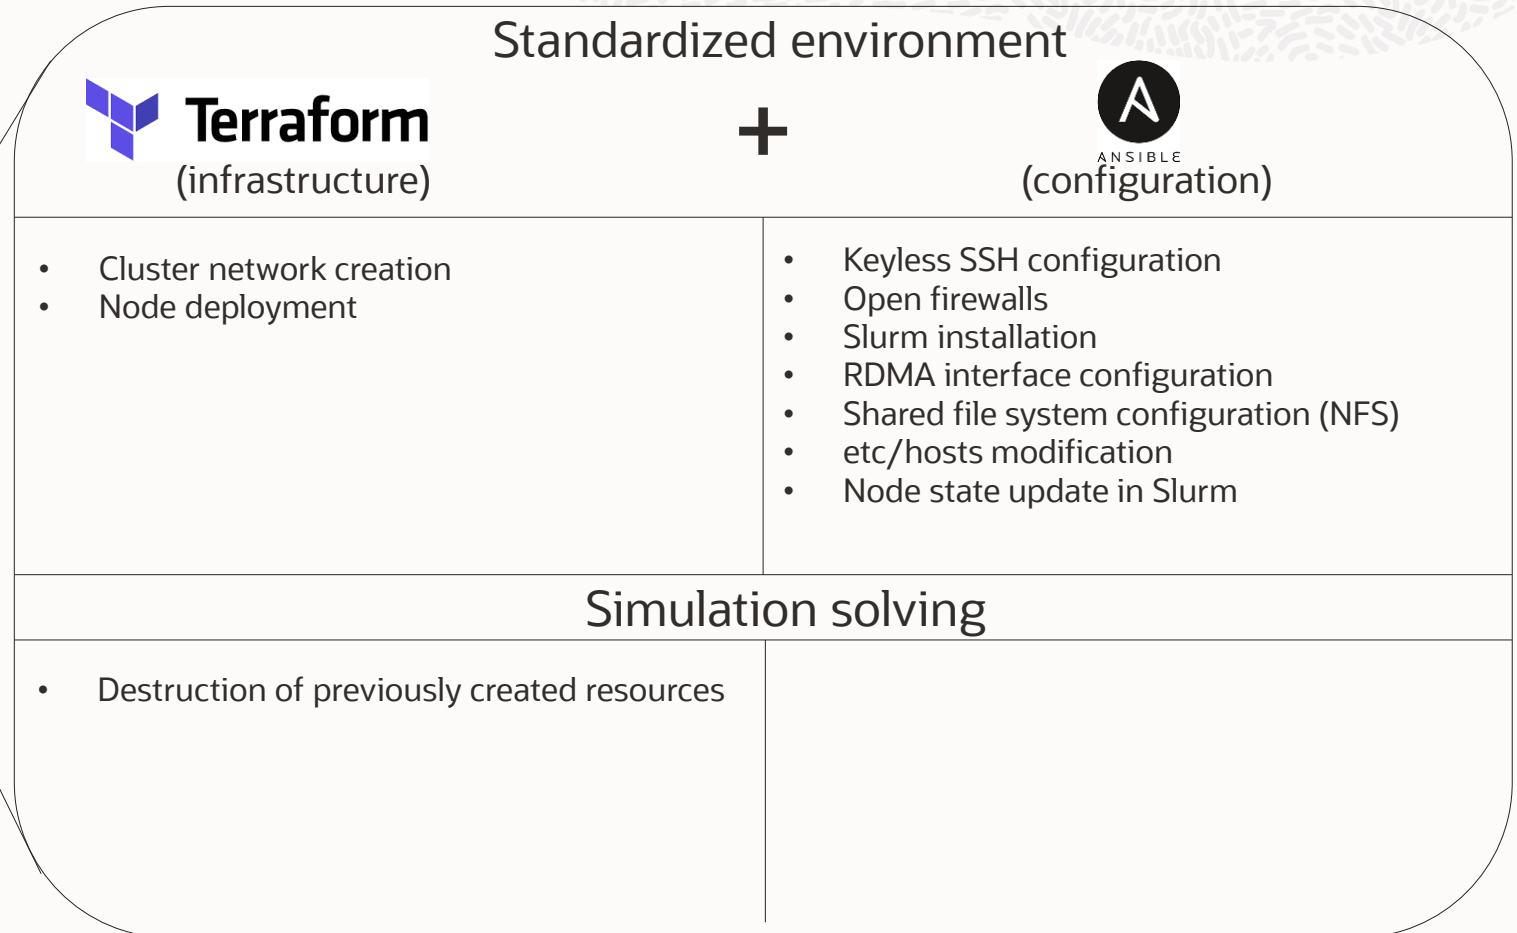


# HPC Bursting Scenarios



# « Cluster per job » strategy (On-demand Provisioning)

Job  
submission





# References


Please contact your account team to arrange individual briefing calls with Reference representatives!

Email [wolfgang.dreyer@oracle.com](mailto:wolfgang.dreyer@oracle.com) for latest version



- LifeScience
- Financial Services
- Manufacturing, Automotive and Aerospace
- Animation- and Media- Industry
- Energy and O&G
- HealthCare
- GPU-usage
- Artificial Intelligence and machine learning

## The Race to Combat Infectious Disease



ADONER Synthetic multipeptide-display scaffold for next generation vaccines. We developed ADONER, a next generation synthetic vaccine that mimics features of Chikungunya virus, to efficiently prime the immune system to protect against this pathogen. <https://doi.org/10.1016/j.cels.2020.100000>

Imphoron Ltd  
Biotechnology - Bristol, Bristol


### Use case

- design a novel vaccine delivery system that is easy to produce in high volumes
- envisioned a vaccine delivery system from a lab-produced, thermostable protein molecule that could be readily manufactured at low-cost and did not require refrigeration to retain viability—important for vaccine shipment to and storage in warmer climate locations.

### Cloud Performance Characteristics benefited for this research

- ✓ Best suited to heterogeneous high-throughput tasks
- ✓ Pipelines needing different node type(s) for different parts
- ✓ Can be much more specific than an average on-premise cluster
- ✓ Always have access to the latest hardware
- ✓ Nodes are only switched on and paid for while jobs are running
- ✓ Nodes are switched off automatically when idle

4 Confidential - © 2019 Oracle Internal/Restricted/Highly Restricted




enabled the team to "dramatically speed up our analyze the COVID-19 virus to design the vaccine candidate".

University, Australia Flinders UNIVERSITY

computer models of the spike protein and its human ACE2 to identify how the virus was infecting human cells

flinders.edu.au/blog/2020/04/03/flinders-targets-covid-19-vaccine/  
Copyright © 2020 Oracle and/or its affiliates.



"Reducing the time to results from 90 to under five days using Oracle HPC is transformational from a research perspective."

University of Bristol



<https://www.oracle.com/au/ocm/obcs/research/higher-education/hpc-combating-epidemics-cs.pdf>

Source: <http://www.bristol.ac.uk/chemistry/news/2020/smoking-cessation.html>

## Nissan supports massive simulation needs with Oracle Cloud

- Japan Automotive | CFD & Crash Simulations

- One of the largest automobile manufacturers in the world with over \$100 Billion in annual revenue
- Required powerful HPC cloud vendor to run complex crash simulations
- Performance was 30% faster on Oracle Cloud Infrastructure vs on-premise at same core count
- Performance was 90% faster when bursting to 1044 cluster networking



## 3 \* Multi National Car Manufacturer Computational Fluid Dynamics and Crash Simulation business on Oracle Cloud

- These companies are big multi national Tier-1 car manufacturer based in Detroit (FCA), Japan and Europe. Their own HPC resources suffer from agility and long processes of renewal. Maintenance and administration bind resources which could be used in car-design.
- The companies can initiate and scale jobs in minutes. State of the art compute resources are available for their engineers at a mouse click
- Engineering simulation of crash or air flow enables the company to design cars faster, with greater fidelity and equal or even lower cost than using their own on-premise compute resources

5% to 20% performance gain on the ideal scaling up to 20,000 cores comparable to major purpose-built on-premises compute clusters

ORACLE

## BJSS moves traditional on-prem trading systems to Oracle Cloud

United Kingdom Consulting Services | Consulting & Delivery

- One of UK's leading IT and business consultancies, winner of the Queen's Award for Enterprise
- Financial services client needed to **move on-premises trading system to the cloud**
- Migrated to Oracle Cloud Infrastructure in less than a week
- Saw **5x increase** in high performance computing speed
- Mean latency on bare-metal HPC was **25% faster** than competing cloud providers

Copyright © 2020, Oracle and/or its affiliates. All rights reserved.



## Rapid Genomic Processing

Oracle Cloud Infrastructure Supports COVID-19 Researchers with NVIDIA Parabricks Genome-sequencing software for GPUs



Oracle Cloud Infrastructure (OCI) is working with NVIDIA to accelerate researchers' understanding of the COVID-19 virus. Oracle OCI has released an updated **NVIDIA GPU Cloud Machine Image** that enables researchers to use NVIDIA's **Parabricks** software suite to perform genome-sequencing of the virus in conjunction with NVIDIA's **accelerated OCI GPU license** of Parabricks to any researcher in the fight against COVID-19.

Copyright © 2020, Oracle and/or its affiliates.

## Updated NVIDIA GPU Cloud Machine Image

The **NVIDIA GPU Cloud Machine Image** on Oracle OCI contains GPU-optimized software that gives researchers, data scientists, and developers an environment to gain insights, build business value, and develop new solutions.

This image takes advantage of cutting-edge Oracle cloud computing infrastructure. You can run the **NVIDIA GPU Cloud Machine Image** across the full portfolio of GPU options on Oracle Cloud, from virtual machines with one or more P100 or V100 GPU chips to bare metal GPU instances with 8 V100 chips attached.

### Parabricks

Researchers can test possible outcomes and responses related to this new and largely unknown virus using **Parabricks** software, designed to perform next-generation sequencing of DNA data.

Parabricks can analyze whole human genomes in ~45 minutes, compared to ~30 hours for 30x WGS data. If you're a researcher working on COVID-19, [click here](https://www.oracle.com/au/oci/parabricks) to request a Parabricks license.


We're excited for our cloud resources to be a small part of our society's response to this pandemic, and hope to see progress in research that benefits our world in the face of this challenge.

ORACLE

## AgroScout helps growers bring food security

Israel Technology Company | AI/Machine Learning

- Uses autonomous drones and AI to detect early diseases and pests in field crops
- Went from static application to dynamic microservice-based app
- Time to deploy a cluster decreased from half day to 2 minutes
- New version deployment reduced from 24 hours to minutes
- Can query thousands of images in seconds



Copyright © 2020, Oracle and/or its affiliates. All rights reserved.

## YellowDog provides on-demand 3D rendering solutions

United Kingdom Technology Company | Cloud 3D Rendering

- Provides rendering options for small to medium media houses
- Sought to expand rendering options and capabilities for customers
- Increased performance 10x vs prior solution and 2x vs AWS
- Moved from first gen Oracle Compute to second gen in less than a year at no incremental cost
- Reduced infrastructure costs resulted in savings passed to customer

Copyright © 2020, Oracle and/or its affiliates. All rights reserved.



"We benchmarked Oracle Cloud against our production deployment with fantastic results that blew anything we had previously used out of the water."

Gareth Williams  
CEO and Founder, YellowDog



# Customer Vision – delivering value quickly and effectively



## Accelerate your migration to cloud

- Dedicated teams to fast-track your Oracle workloads to the cloud
- Specialism by workload type for efficiency and effectiveness



## Best Practice Reference Architectures

- Based on Oracle best-practice deployment architectures
- Deployment architectures fine-tuned for you as an individual customer
- Delivered as a working system with all configuration and documentation



## Access to Specialist Oracle Teams

- Optimise and improve your applications as they migrate
- Direct access to product experts
- Address security and networking challenges
- Fast-track connectivity

## 4 Reason We Win

... if you only remember 1 slide

### We run Bare Metal with off-box virtualization

- Fastest performance in the cloud, comparable to on-prem and managed platforms
- 100% yours, no hypervisor, no virtualization, no software

### We are the only ones to run a flat & fast (CLOS) Network with RoCE


- Only Cloud with 100G RDMA over Converged Ethernet (RoCE), ~1.5µs cluster networking
- Flat, fast and never over-subscribed (every node is max - 2 hops away) with an SLA

### We run “NVMe Everywhere” and don’t charge you to boost performance

- 6.4T NVME (more than anyone) & 32TB Block ea.
- The fastest parallel file system performance

### We run the Lowest Cost for Compute, Storage + Network – Price it out ...

- Unmatched pricing for HPC compute, highest performing block storage and network
- Ease of purchase via OCRE/GÉANT with no need for complex procurement and competitive tenders



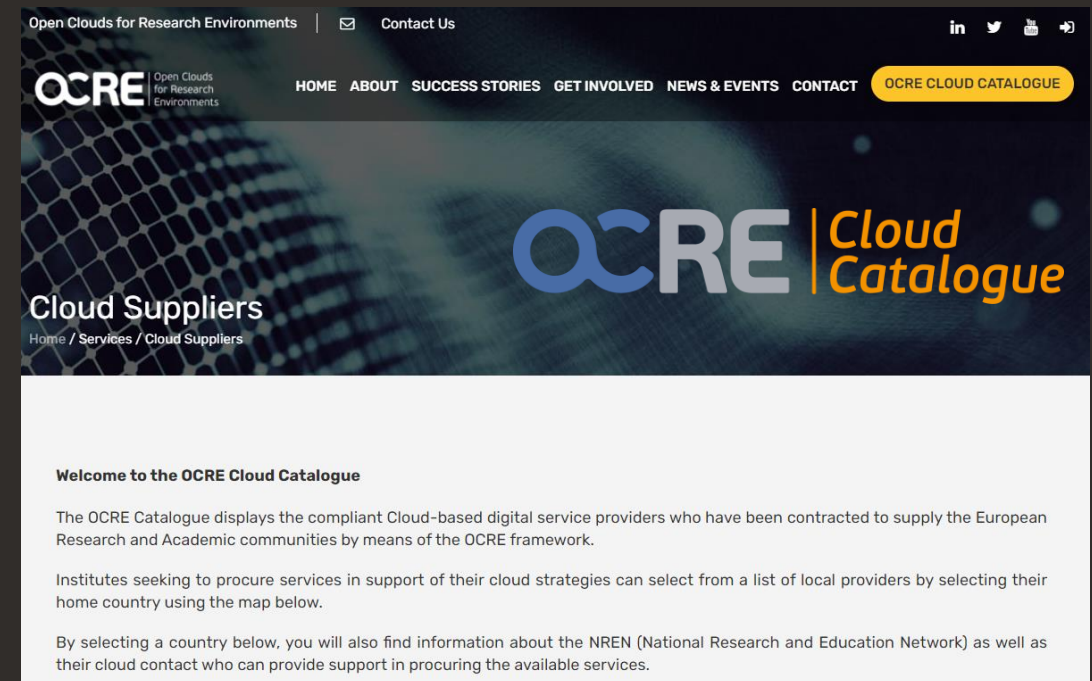
# Offering Cloud Services Through GÉANT

**OCRE** (Open Clouds for Research Environment) Project Consortium funded by the European Commission for 4 years (2021 – 2024)

- the Framework Agreement with the Cloud Service Provider or Reseller
- Scope of Services offered by the Cloud Service Provider (IaaS+)
- Pricing Information
- Call Off agreement between your institution and the Cloud Service provider

Education and research institutions in Europe can now directly consume cloud services without the need for public procurement.

Oracle Cloud services are available via the Oracle partners listed in the OCRE Cloud Catalogue.



# Open Clouds for Research Environments

## OCRE Framework Benefits



### Ease of procurement

By using this framework, research institutions do not need to run complex procurements and competitive tenders.



### Predictable pricing

PAYG and FAM with volume discounts, cloud credits won't be lost. Value-added services by implementation partner.

**Trusted advisory**  
Cloud Catalogue of services promoted by NRENs.  
EC-funded Cloud Vouchers by GÉANT.  
Demand aggregation and shared use.



### Meeting compliance

Network connectivity and identity (SSO) integration via GÉANT/NRENs (DFN).  
Data security and privacy (GDPR) compliance.



# Oracle Cloud Services für Géant

Oracle bietet mit seiner Oracle Cloud Infrastructure eine lizenzkonforme Cloud-Architektur, deren Services einfach, sicher, skalierbar und hochverfügbar sind. Logicalis unterstützt und berät Sie beim Weg in die Cloud, damit Ihre IT-Infrastruktur zukunftsfähig und sicher bleibt.

## Ihre Vorteile

- ✓ Schnelle und einfache Bereitstellung neuer Umgebungen
- ✓ Skalierbar für Wachstum und Engpässe
- ✓ Lizenzkonformität für Oracle-Software-Lösungen
- ✓ Kostenreduktion durch Bedarfs-gerechte Nutzung
- ✓ Höhere Verfügbarkeit und damit Einhaltung von SLAs

Überblick

Oracle Cloud Lösungen entdecken

Oracle Cloud Services für das DFN

## Überblick

**GÉANTCloudServices**  
CONNECTING RESEARCH AND EDUCATION TO THE CLOUD

## Kontaktieren Sie uns

Nachname \*

Geschäftliche E-Mail \*



Ihre Fragen an uns ...

**Vielen Dank.**



**Dr. Rita Engemaier**

Account Executive Cloud Technology  
Hochschulen und Forschung

[Rita.Engemaier@oracle.com](mailto:Rita.Engemaier@oracle.com)



**Werner Bauer**

Principal Cloud Engineer

[Werner.Bauer@oracle.com](mailto:Werner.Bauer@oracle.com)



**Wolfgang Dreyer**

OCI Specialist Director, EMEA CoE

[Wolfgang.Dreyer@oracle.com](mailto:Wolfgang.Dreyer@oracle.com)

ORACLE