



Contribution ID: 149

Type: **Panel presentation**

Oracle Data Science Cloud

Monday 25 January 2021 13:30 (15 minutes)

Oracle Data Science Cloud is a collaborative platform for data scientists to build and manage ML models. Leveraging open source technology, it provides a scalable cloud-based platform for data scientists to explore data and train, save, and deploy models, while utilizing the rich Python ecosystem as well as Oracle's proprietary Python libraries.

- Data Analysts - Oracle Analytics Cloud delivers cutting-edge visualization, augmented analysis, and natural language processing through an easy-to-use interface. Powered by AI and machine learning, Oracle Analytics Cloud makes it possible for any level of user to generate deep insights and create forward-thinking reports.
- Data Engineers - Oracle Database's Machine Learning capabilities bring the latest automation and self-learning tools to the database space. The result is an experience that's both powerful and user-friendly. Using Oracle's tools, it's easier than ever to manage data and support application development on a secure and scalable infrastructure.

The platform combines three key components:

- Infrastructure - systems tasks (like spawning servers) are abstracted and handled automatically so data scientists can focus on the substance of their work.
- Tools - open source tools (like Jupyter, R Shiny, or modeling libraries) that data scientists need are integrated into a centralized place.
- Workflow - automation for tasks, collaboration, and communication that let data science teams effectively deliver on their mission.

The presentation is going to include a live showcase of the Oracle Data Science Cloud platform using a simple user scenario!

Oracle Cloud Infrastructure offers exceptional performance, security, and control for today's most demanding high-performance computing (HPC) research workloads. The Oracle HPC cloud consists of:

- Bare metal instances support applications requiring high core counts, large amounts of memory, and high memory bandwidth. Users can build cloud environments with significant performance improvements over other public clouds and onsite data centers. Bare metal compute instances provide researchers with exceptional isolation, visibility, and control.
- NVIDIA-based offerings provide a range of options for graphics-intensive workloads, along with the high performance demanded by AI and machine learning algorithms.
- Virtual cloud networks enable researchers to easily move their existing network topology to the cloud. Standard bare metal servers support dual 25 Gbps Ethernet for fast front-end access to your compute clusters. Oracle's groundbreaking, back-end network fabric lets them use Mellanox's ConnectX-5, 100 Gbps network interface cards with RDMA over converged Ethernet (RoCE) v2 to create clusters with the same low-latency networking and application scalability that one would expect on premise.

Oracle provides several high-performance storage options suitable for HPC workloads.

- Local NVMe SSD: High-speed local flash storage ideal for large databases, high performance computing (HPC), and big data workloads such as Apache Spark and Hadoop.
- Block volumes: Standard block storage services offering 60 IOPS per Gb, up to a maximum of 32,000 IOPS per volume, backed by Oracle's highest performance SLA.
- Parallel file systems: HPC requires larger data sets and higher performance than standard enterprise file servers can provide. Customers often build their own parallel file systems either on premise or in the cloud using open source software such as Lustre.

Author: SZEGEDI, Peter (Oracle)

Presenter: SZEGEDI, Peter (Oracle)

Session Classification: On-premise, hybrid or cloud?

Track Classification: Main session