

ORACLE®




Cloud Infrastructure

GPUs for the Public Cloud

Ben Lackey – Director, Big Data Partners

Taylor Newill – Product Management, HPC/GPU



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.

Our Journey

ORACLE®

Copyright © 2019 Oracle and/or its affiliates. All rights reserved.



ORACLE®
SEATTLE

CLOUD(Y) CITY




Three Strategic Pillars for Oracle Cloud Infrastructure



Best place to run
specialized computing



Highest, most consistent
performance at the
lowest cost



Open Cloud

digital

DEC
1993



Sun and Mellanox

2010




BARE METAL CLOUD

Oracle Open World 2016



ORACLE CLOUD INFRASTRUCTURE

October 2017



NVIDIA TESLA P100 INSTANCES

GENERAL AVAILABILITY

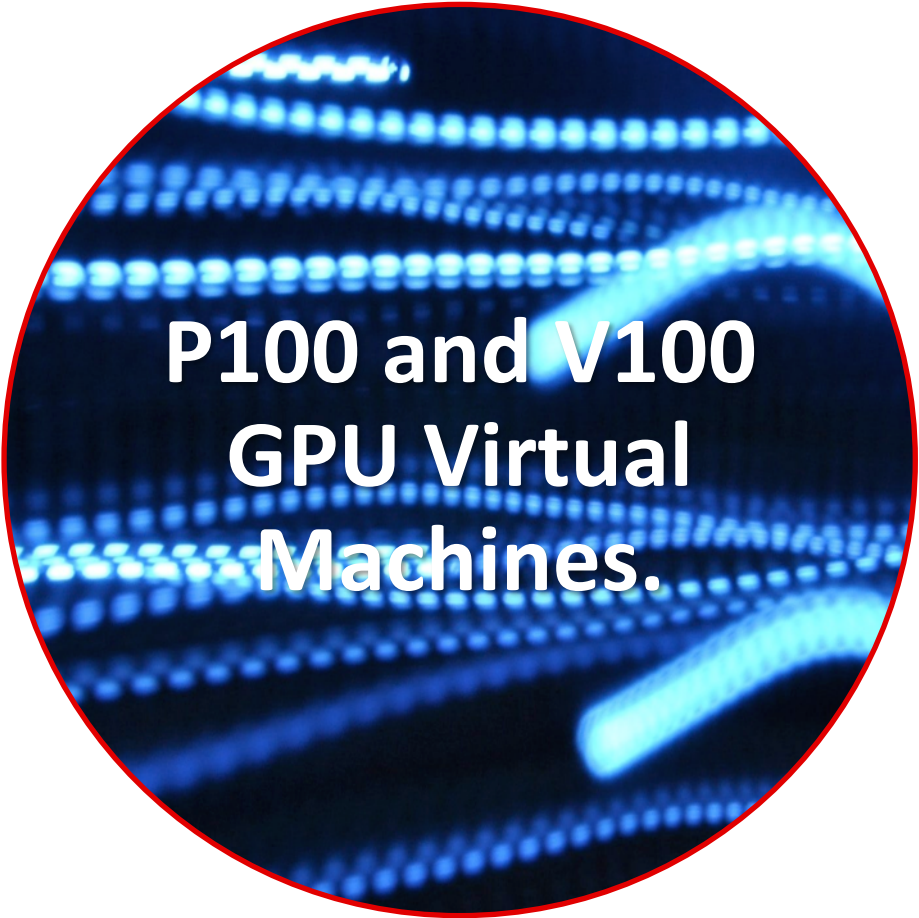
Oracle Open World 2017



Bare Metal
V100 Volta
instances

GENERAL AVAILABILITY

March 2018



P100 and V100 GPU Virtual Machines.

TOKYO
September 2018

ACCELERATING ORACLE CLOUD INFRASTRUCTURE WITH NVIDIA RAPIDS


*are runs seamlessly on the Oracle Cloud, allowing
support all their HPC, AI and data science needs,
g advantage of the portfolio of GPU instances
Oracle Cloud Infrastructure."*

Software Development, Oracle Cloud Infrastructure

ORACLE



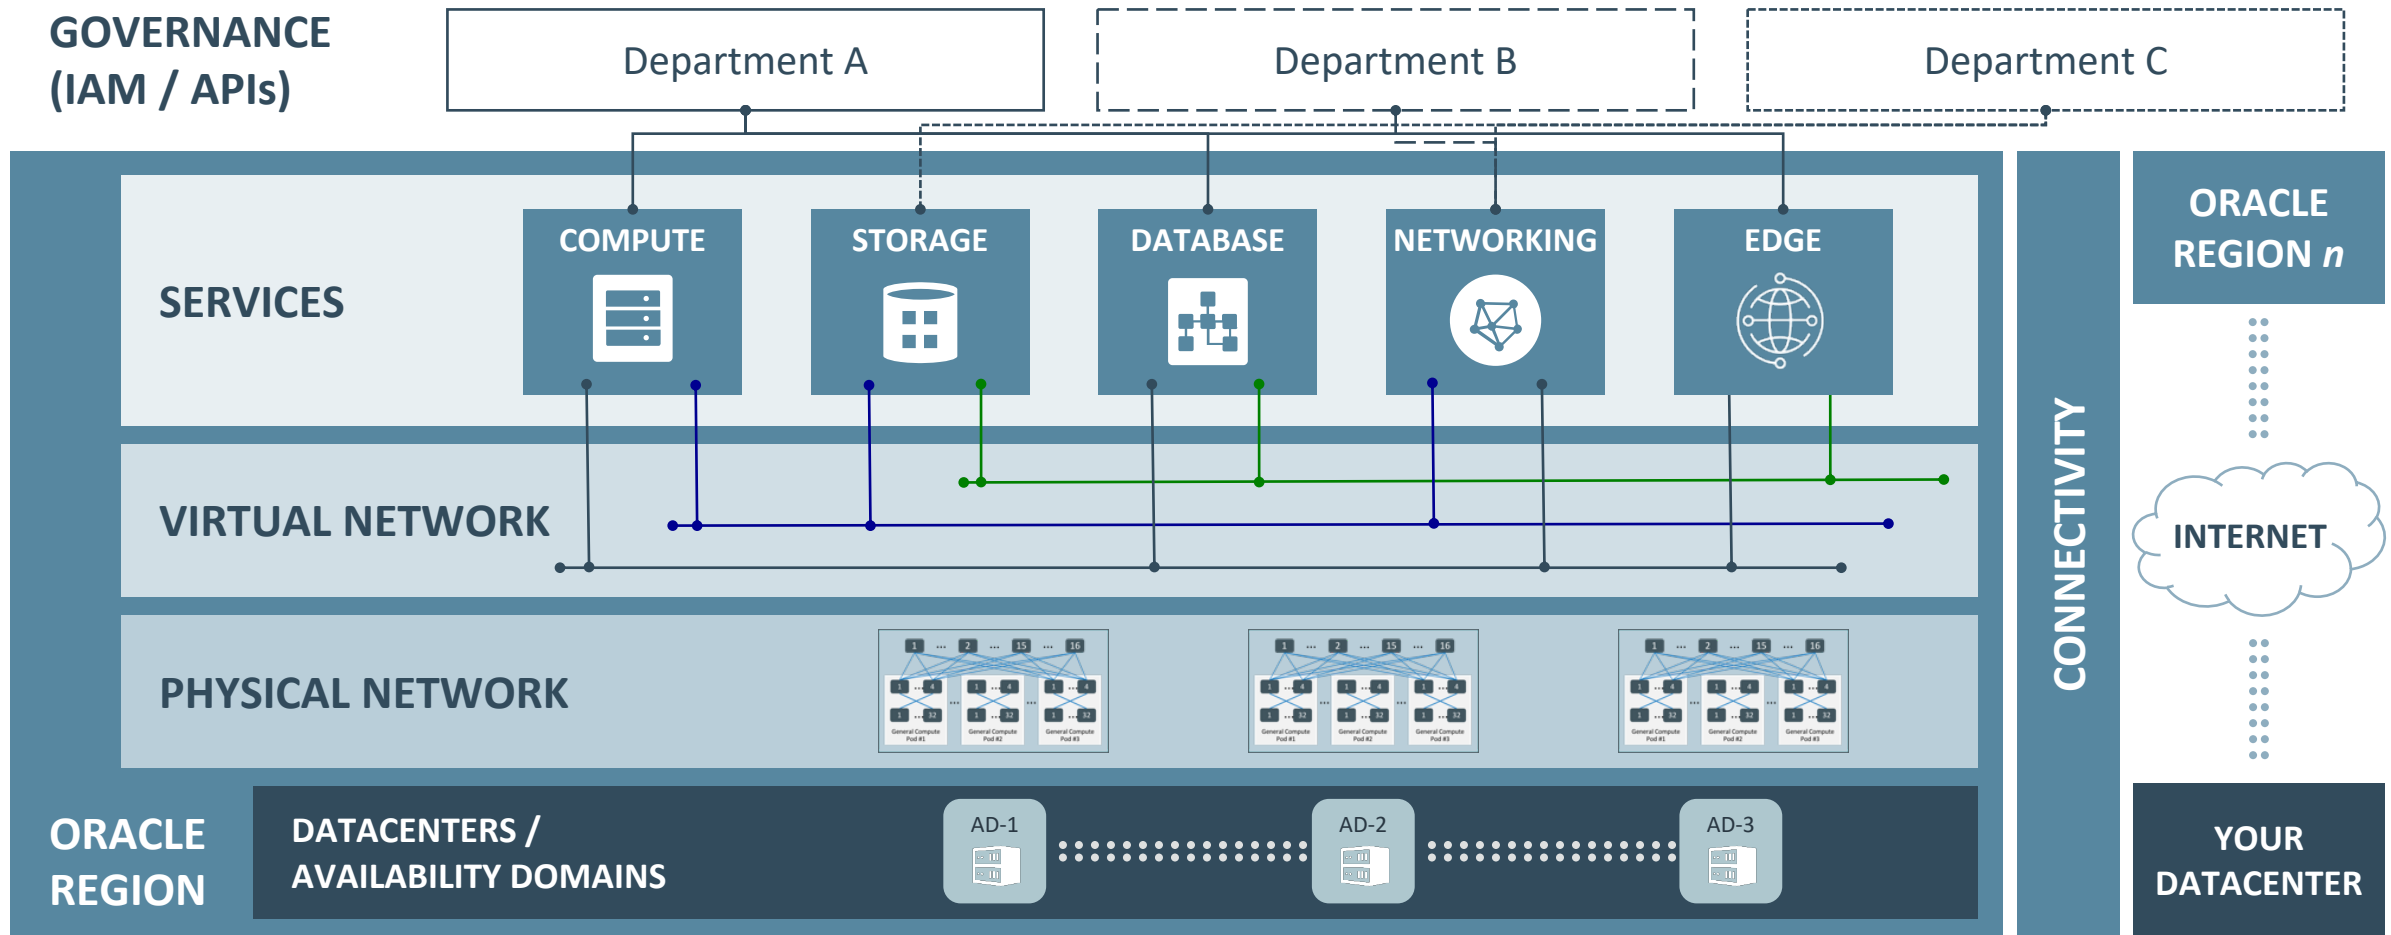
THE METAL IN BARE METAL



ORACLE®

Oracle Cloud Infrastructure

High performance compute, storage, database, edge on the same flexible virtual network





Virtual Machine GPU
1 P100 GPU
12 OCPU | 104 GB RAM
up to 1PB Block Storage
\$1.275 GPU/h




Virtual Machine GPU v2
1 – 4 V100 GPU
6 – 24 OCPU
104 to 360 GB RAM
up to 1PB Block Storage
NVLINK
\$2.25 GPU/h



Bare Metal GPU
28 Cores, 192 GB RAM,
2x Tesla P100 GPUs
\$1.275 GPU/h



Bare Metal GPU V2
52 Cores, 768 GB RAM,
8x Tesla V100 GPUs
NVLINK
\$2.25 GPU/h



Oracle Cloud Infrastructure GPU


COMING SOON



Bare Metal GPU v3
48 OCPU – 3.5GHZ all core turbo
768 GB RAM
HGX-2
8 * V100 32 GB SXM3 GPU
NVLINK
Up to 1PB Block Storage



Virtual Machine GPU v3
5 - 22 OCPU – 3.5GHZ all core
turbo
90 - 360 GB RAM
HGX-2
1 - 4 * V100 32 GB SXM3 GPU
NVLINK
Up to 1PB Block Storage



Oracle Cloud Infrastructure BM and VM




Bare Metal Standard
52 Cores, 768 GB RAM,
up to 1PB Block Storage
2x 25Gbe Network Interfaces

\$ 0.0638 / core/h



Bare Metal DenseIO
51.2 TB of local NVMe SSD
up to 1PB Block Storage
2x 25Gbe Network Interfaces

\$ 0.1275 / core/h



Virtual Machine Standard
1-16 and 1-24 Cores
7 – 320 GB RAM
up to 1PB Block Storage
Up to 25Gb/s Ethernet

\$0.0638 / core/h



Virtual Machine Dense
4 – 24 Cores
60 – 320 GB RAM
3.2 TB – 25.6 TB local NVMe SSD
up to 1PB Block Storage

\$0.1275 / core/h

Oracle Cloud Infrastructure Storage

\$0.0425 GB/month



File Storage Service
Up to 8 Exabytes of storage
Managed distributed file service
NFSv3 mount point
Pay for what you use

\$0.0425 GB/month




Block Storage
50 GB-16 TB volumes
Up to 25K IOPS per volume
400K IOPS per host

320MB/S per volume
Pay for Storage not IOPS



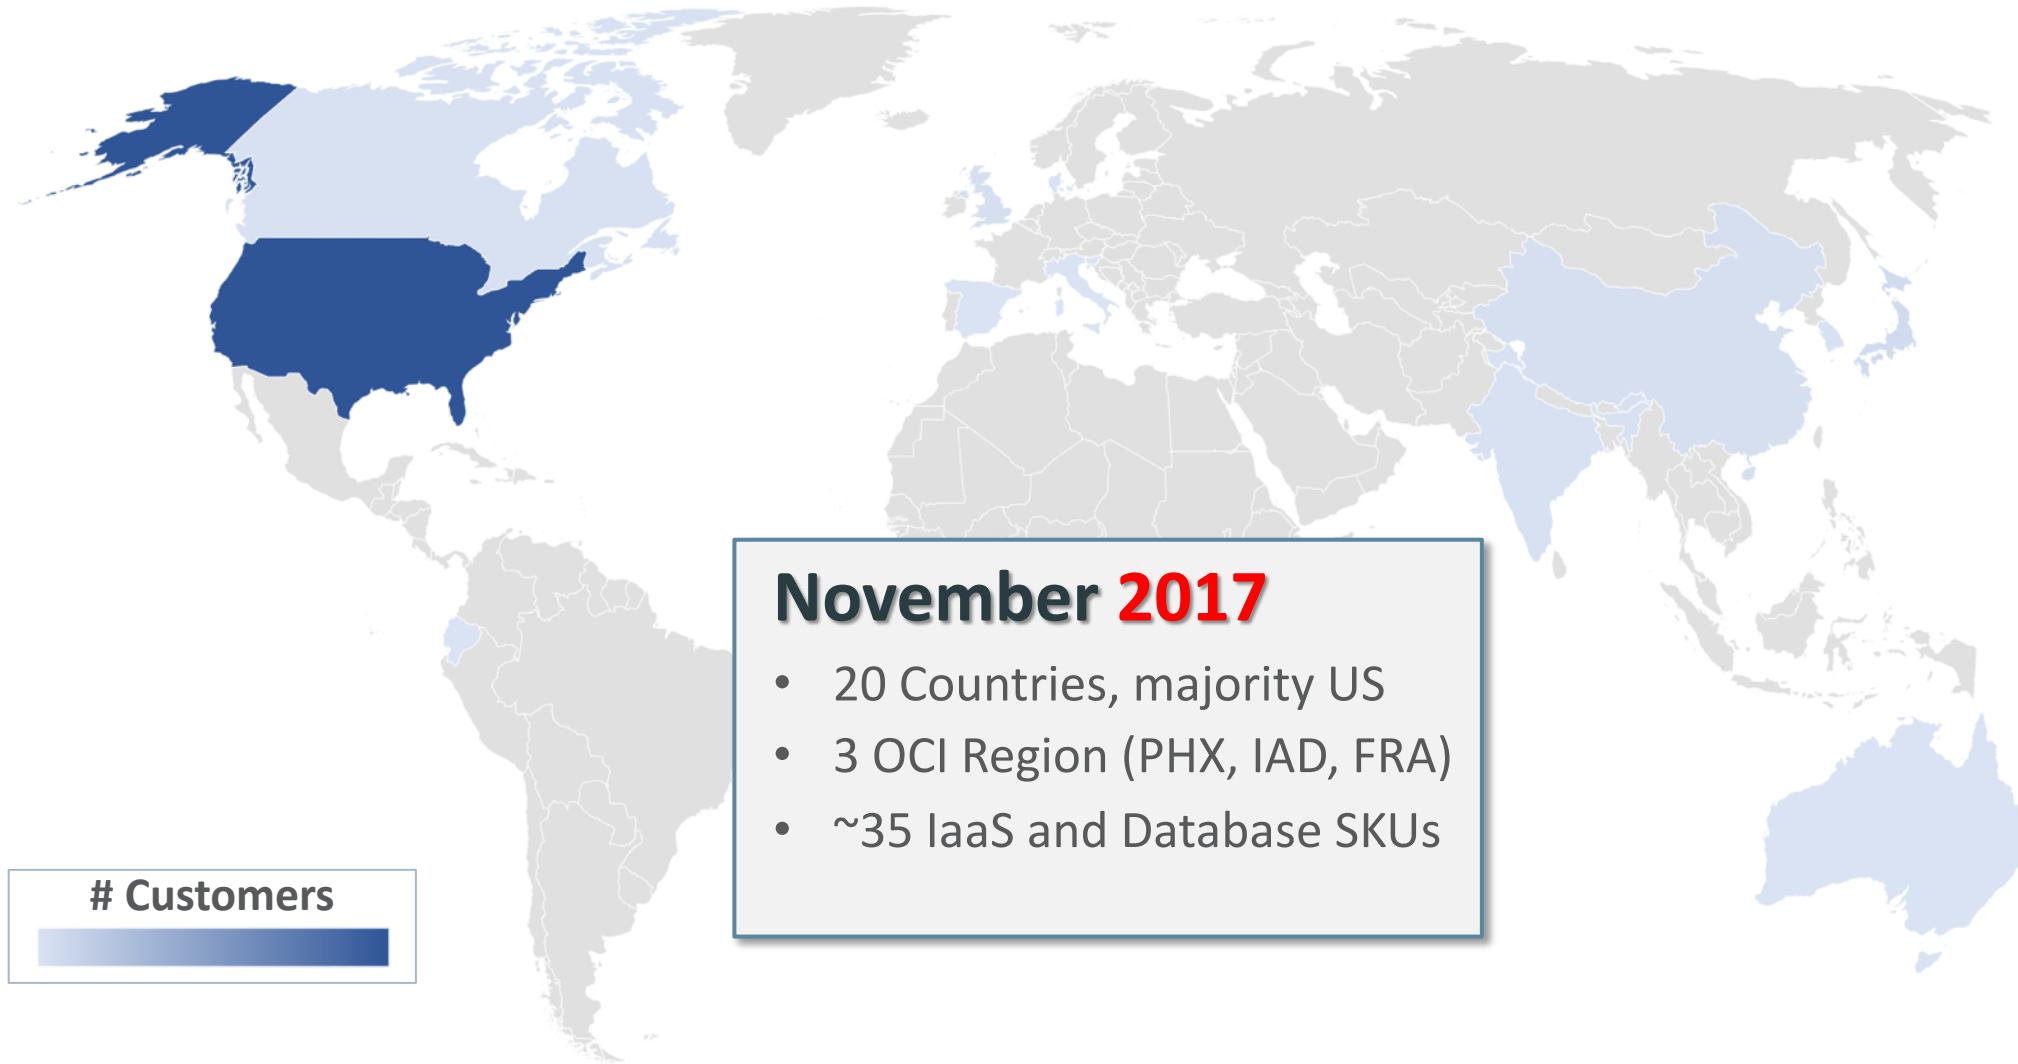
Local NVME
51.2 TB of local NVMe SSD
6.4-25.6 TB for VM



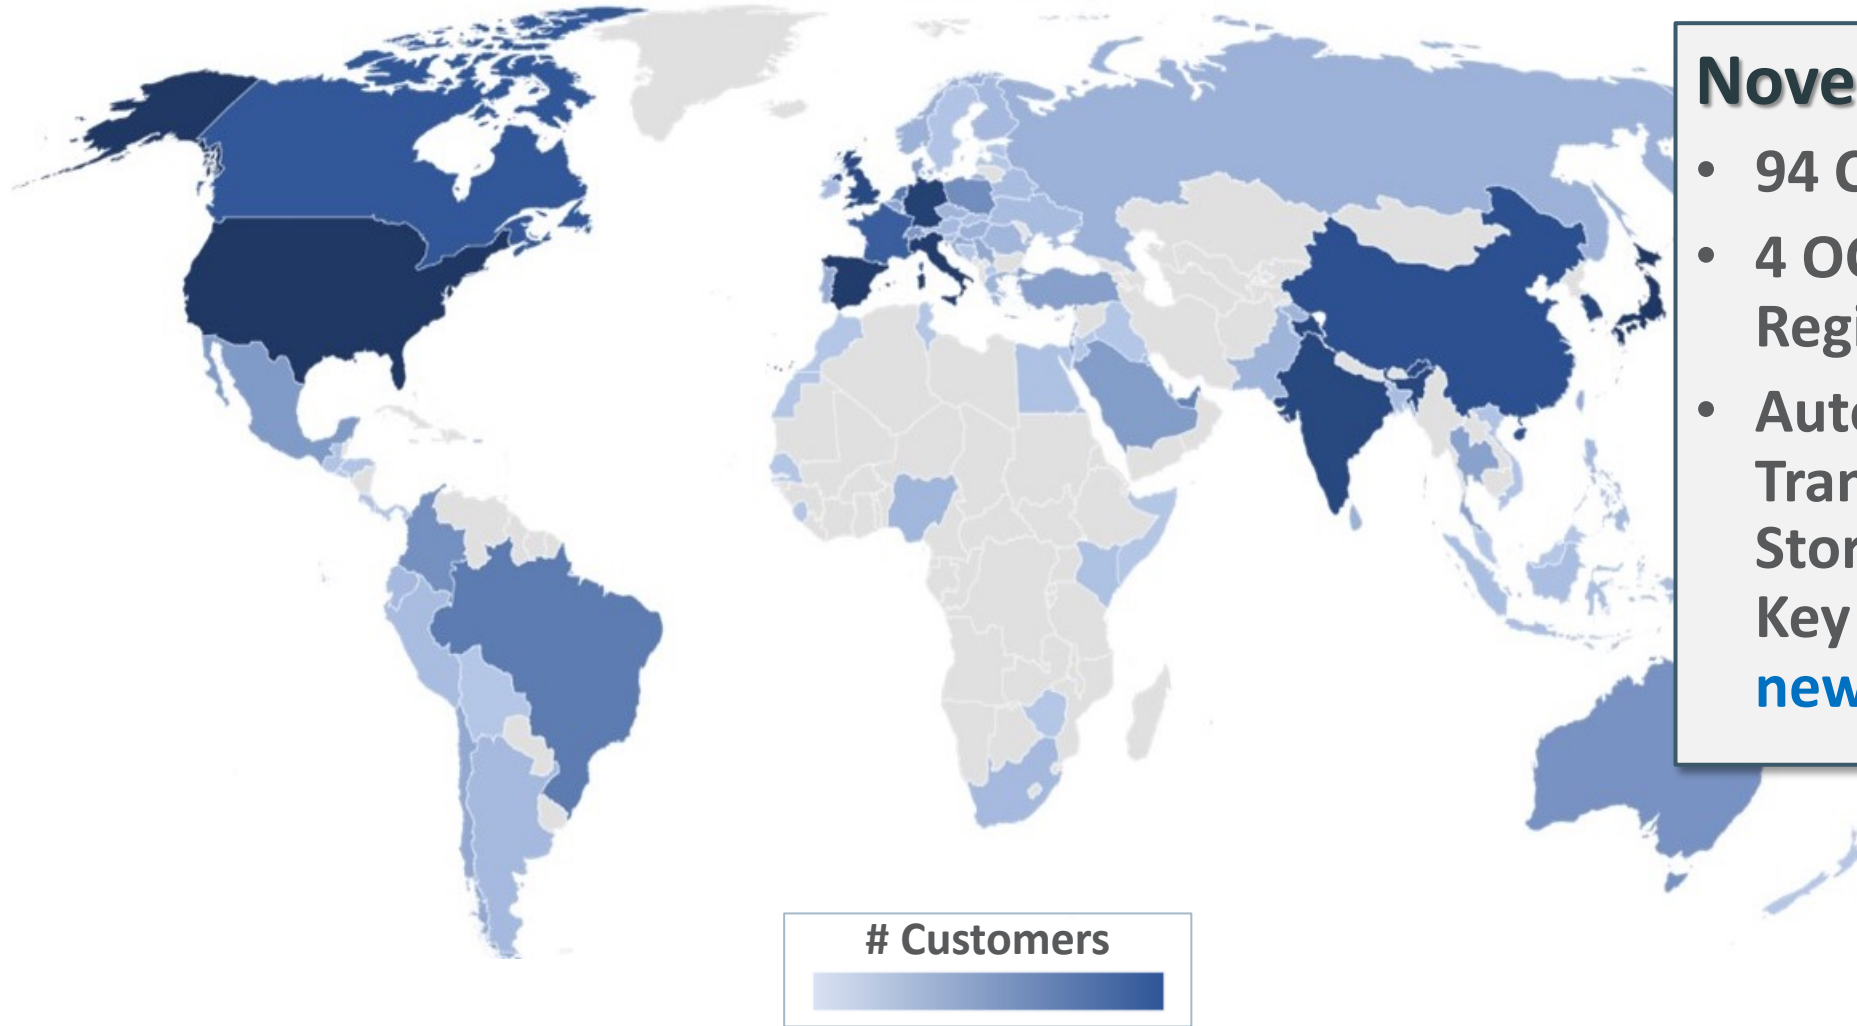
Object Storage
Up to 10TB per object
Highly durable

\$0.0255 GB/month
\$0.0034 10,000 req

OCI Customer Base – 1 year ago



OCI Customer Base - 20x YoY Growth

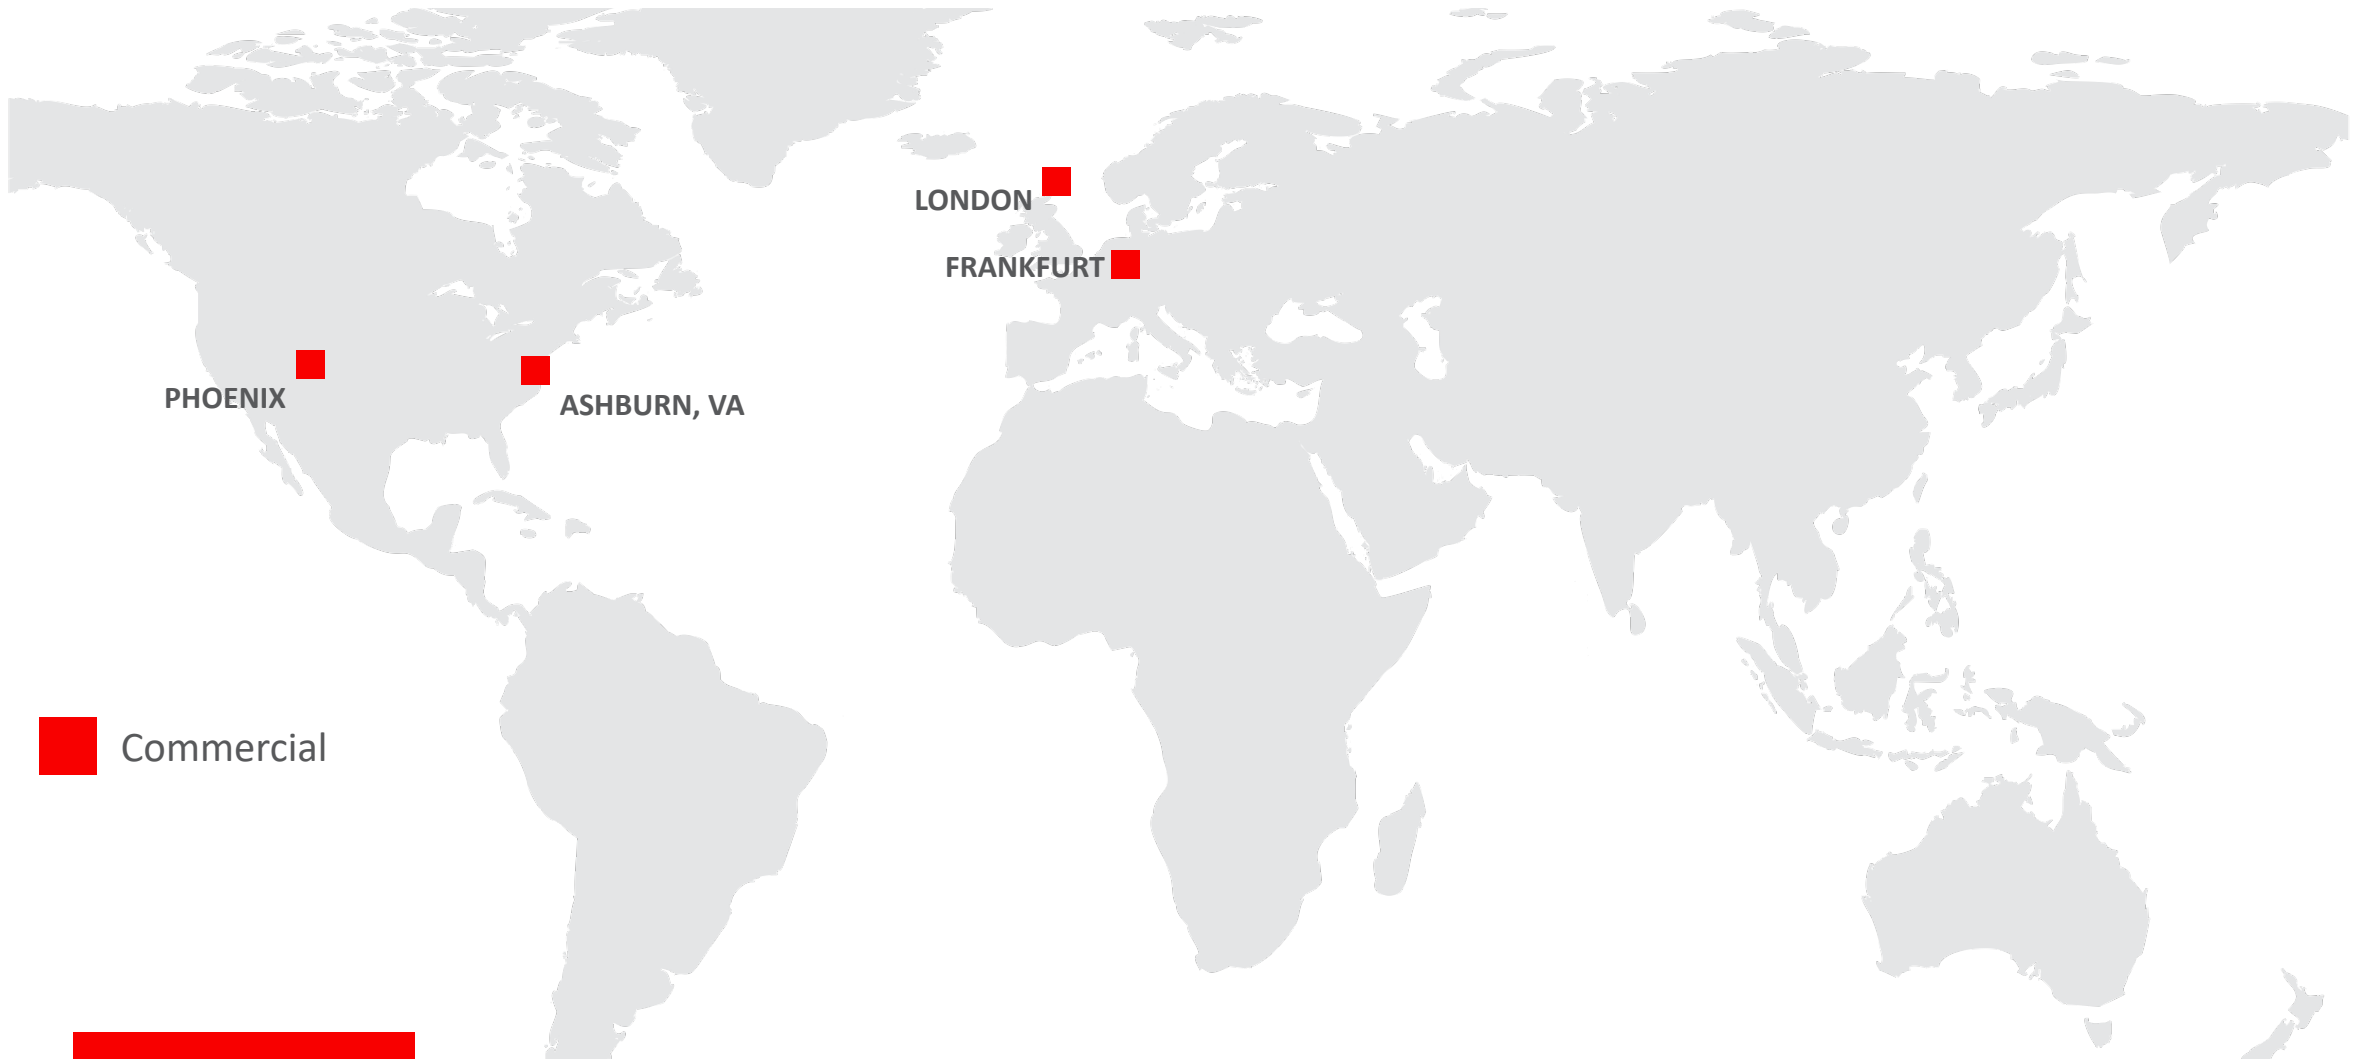


November 2018

- 94 Countries
- 4 OCI Commercial Regions
- Autonomous, Data Transfer, GPU, HPC, File Storage, Cost Tagging, Key Management, **200+ new services**

Oracle Cloud Infrastructure Global Footprint

Sep 2018



ORACLE®

Oracle Cloud Infrastructure Global Footprint

Dec 2019



- Commercial
- Government
- Edge Points of Presence

ORACLE®

kinetica



Workbench

Active Analytics

admin

.....

Log In

NEW YORK
HQCHICAGO
OFFICELONDON
OFFICEFRANKFURT
OFFICESHANGHAI
OFFICETOKYO
OFFICE

Assets Under Management



\$726B

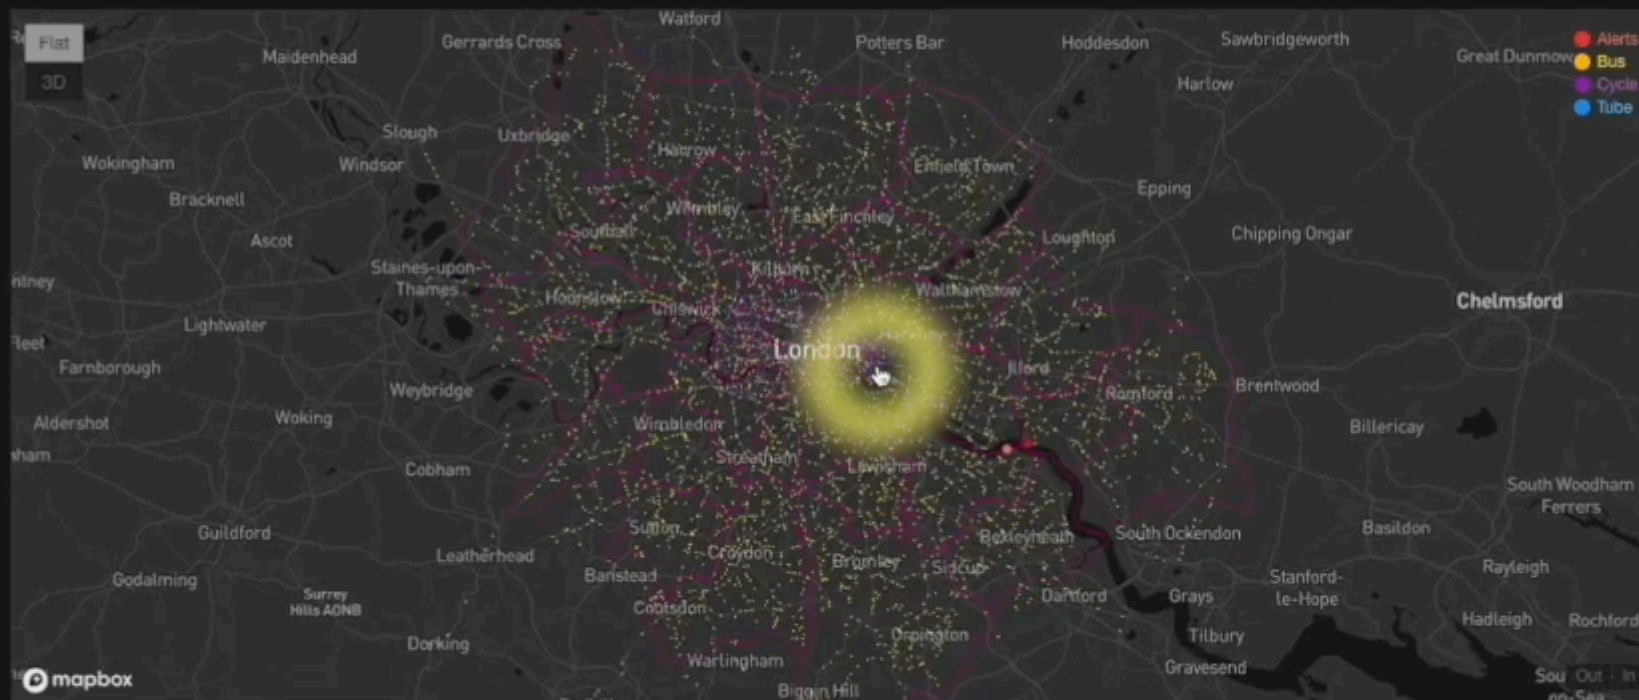
Streaming Value at Risk

0.2%

\$1.5B


Connected City Realtime ☆





kinetica

RESET ALL



30.1k
Bus Events

29.8k
Cycle Events

60.1k
Tube Events

8.00
Alert Events




Three Strategic Pillars for OCI



Best place to run
Oracle workloads




Highest, most consistent
performance at the
lowest cost



Open Cloud

Myth: GPUs don't perform well on the cloud

- **Truth:** Oracle Cloud Infrastructure GPUs match or exceed DGX1 performance
- Hardware refresh is faster and consistent on OCI. Jobs complete faster with each generation and the price performance increases
- Cost of full GPU server on-premises is over \$150,000 plus power, space, maintenance and support





HGX-2 vs. HGX-1



Myth: Cloud means I **HAVE** to run virtualized

- **Truth** OCI Bare Metal runs HPC workloads up to 18% faster than a VM
- Faster turn-around time for a simulation model leads to increased license ROI
- 15% can mean hours and days when running large simulations or AI workloads





Customers




Three Strategic Pillars for OCI



Best place to run
Oracle workloads



Highest, most consistent
performance at the
lowest cost

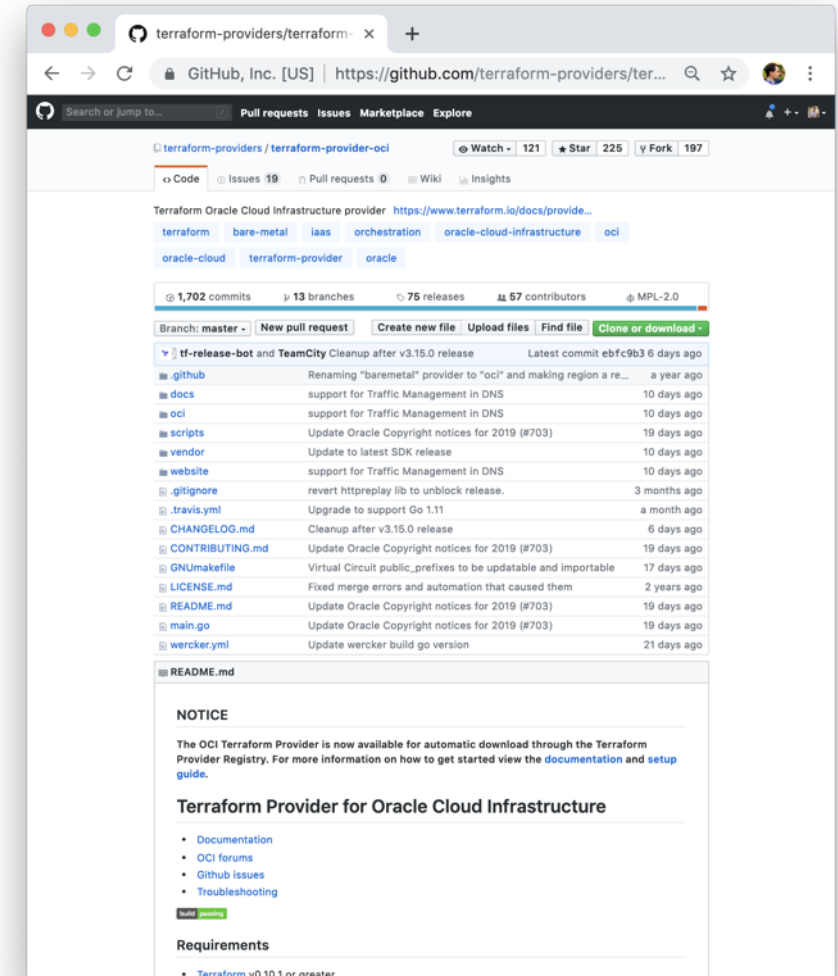


Open Cloud



Terraform

- Open Source Infrastructure as Code (IaC)
- Runs everywhere - OCI, AWS, Azure, GCP, Alibaba...
- OCI doesn't have a proprietary IaC layer, just Terraform!
- The OCI provider is open source and hosted in [Hashicorp's org](https://github.com/terraform-providers/terraform-provider-oci)



Big Data Ecosystem

Integration



BI



AI/ML



Data



Storage Services



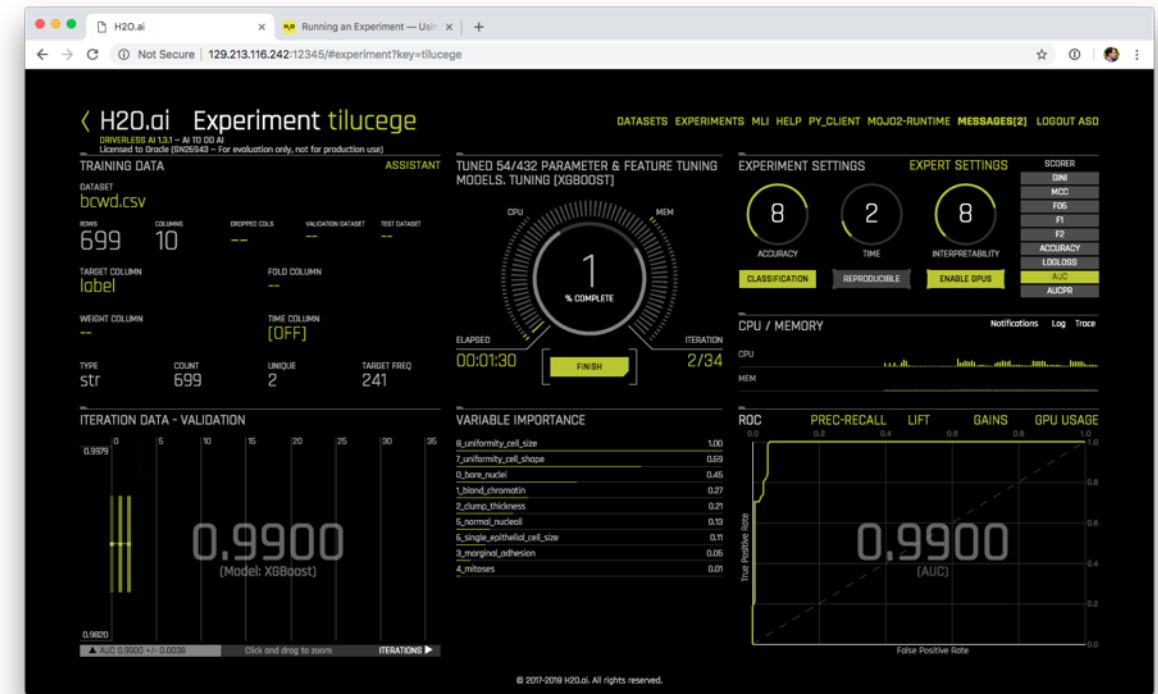
Foundational Services



ORACLE

H2O.ai

- H2O.ai Driverless AI (DAI) is a powerful machine learning tool for quickly creating models
- Ability to run on GPUs
- Automated feature engineering
- Time series support
- Cross validation
- Model export to Java and Python



ORACLE®