

The Yottabyte Research Cloud: A Case Study of a Private Data Science Cloud



UNIVERSITY OF
MICHIGAN



- Established 1817
 - 61,000 total enrollment
 - 7,300 regular faculty
 - Annual research budget of \$1.39 billion, largest among public U.S. universities.
-
- Based in Bloomfield Township, Mich.
 - Virtualization and cloud vendor offering public, private and hybrid cloud solution
 - yCenter
 - YottaBlox.

The Problem


- Like many Academic Institutions, the University of Michigan faculty are attacking research problems that are increasingly solvable only by analyses of larger and more complex datasets.
- Many of these datasets come bundled with varying levels of restrictions: HIPAA, ITAR, FISMA, etc.
- Researchers are in need of flexible compute and storage environments that provide consistency, performance and reliability.



U-M researchers are using YBRC to study health care claims data from across the country.

The Solution

- Yottabyte's virtualized datacenter solution provides a performant, dynamic environment to deploy and manage a variety of data science initiatives, including but not limited to:
 - Restricted data VDI desktop
 - On demand database provisioning
 - Complex architecture deployments (kubernetes, mesos, etc.)
- Yottabyte does this by providing out-of-the-box capabilities:
 - Performant: Based on cutting edge Intel components (NVMe storage)
 - Dynamic: Robust and feature rich software defined networking
 - Cost-Effective: Data deduplication for more efficient use of storage




Benefits for Research at U-M

- Yottabyte Research Cloud sub-services
 - FISMA capable Virtual Desktop Infrastructure (called “The Glovebox”)
 - Data Science microservices for research
 - Streaming Data Service for ingesting data from IoT devices
- Agile Deployment platform
 - Researchers can quickly spin-up complete virtual enclaves with little administrative overhead
 - Automating deployment tasks eases the burden of security auditing
 - Patching and audit logs become scheduled, automatic tasks
- Consistent Platforms
 - Consultants and support staff have a common platform to provide support



M | **ARC-TS** ADVANCED RESEARCH COMPUTING TECHNOLOGY SERVICES UNIVERSITY OF MICHIGAN

Advanced Research Computing – Technology Services (ARC-TS) provides U-M researchers across campus with the computational resources they need to produce ground-breaking scientific discoveries. ARC-TS designs, procures, and maintains cutting-edge resources for data-intensive and computationally demanding research.



Flux, U-M's campus-wide HPC cluster, consists of 27,000 cores.

M | ADVANCED RESEARCH COMPUTING

COMPUTATIONAL SCIENCE

DATA SCIENCE

TECHNOLOGY SERVICES

CONSULTING SERVICES

For more information

- arc-ts.umich.edu/ybrc
- yottabyte.com
- Jeremy Hallum, ARC-TS: jhallum@umich.edu
- Duane Tursi: duane@yottabyte.com