

# PILOT PROGRAM PROFILE UNIVERSITY OF MICHIGAN PILOTING FLEXIBLE, RESPONSIVE COMPUTING SERVICES

U-M Advanced Research Computing Service Pilots a  
Software-Defined Infrastructure Solution from Intel  
and Yottabyte



# The University of Michigan lives and breathes research and discovery.

Home to some of the world's most talented researchers and educators, U-M's capabilities and accomplishments span nearly every field of science, engineering, medicine, social sciences, management, education, and the humanities.

Supporting this thriving research community is Advanced Research Computing (ARC). ARC empowers data-intensive and computational research at U-M, providing high-performance computing resources, services, support, instruction, and consultation.

ARC currently serves 2,600 independent groups, providing them the computational power to research everything from genomics to primary education.

Each group has specific computing needs, and a large portion have stringent data security and compliance requirements, such as the Health Insurance Portability and Accountability Act (HIPAA), the Federal Information Security Management Act (FISMA), International Traffic in Arms Regulations (ITAR), and the Children's Online Privacy Protection Act (COPPA).

## The Challenge: How to Better Serve U-M Researchers

"Standing up a traditional IT infrastructure for every new project is labor-intensive and costly," says Brock Palen, ARC's associate director of Advanced Research Computing—Technology Services. "That limits our ability to respond to new research interests and opportunities."

Palen is looking to simplify and strengthen compliance, too, ideally by being able to enforce policy at a global system level rather than at individual project and user levels.

The ability to govern globally would also foster better collaboration, Palen notes, by allowing researchers to more easily share data outside their group.

"Our researchers should be able to work on what their expertise is. They shouldn't have to worry about MongoDB\* export policies and microservices and such," he says.

Then there's the matter of ever-expanding data sets.

Many projects, particularly genomics ones, already require massive data sets, and computing power and storage to match. The reference data set for the 1000 Genomes project, for example, is currently 54 terabytes. The next data set generated will be 10 times larger.

## The Solution: Data Center Modernization

Palen is in the early phases of piloting a software-defined infrastructure (SDI) solution for a select group of ARC-supported researchers. SDI integrates physical and virtual infrastructure, transforming the data center into a flexible, scalable, and programmable pool of stateless resources, ready to be provisioned for new apps and policies.

The ARC implementation is based on Intel® hardware—including the new Intel® Xeon® processor E5 v4 family—and Yottabyte\* yCenter software, which turns commodity hardware into a hyperconverged infrastructure.

SDI enables users to quickly and easily deploy applications, provision



### THE INTEL®-YOTTABYTE\*

pilot promises to open up access to ARC's advanced computing resources, by removing the cost barriers associated with building a traditional architecture.

---

# ARC'S SDI PILOT WILL ALLOW FOR THE ON-THE-FLY CONFIGURATION OF INDIVIDUAL ENCLAVES—ESSENTIALLY A PRIVATE DATA CENTER—FOR EACH RESEARCH PROJECT.

## EXPECTED TIME SPENT PROVISIONING NEW RESEARCH PROJECTS, BEFORE AND AFTER SDI

Before: two weeks  
After: minutes



virtual machines (VMs) and storage, and reconfigure the infrastructure without changing the underlying hardware. The elastic distributed platform also automatically adjusts when new resources are added, so there's no need for capacity planning and no risk of overprovisioning.

ARC'S SDI pilot will allow for the on-the-fly configuration of individual enclaves—essentially a private data center—for each research project. An enclave can consist of multiple services, such as a database service, high-performance computing on demand, or Apache Hadoop\* software on demand.

Each enclave will have its own set of policies based on regulatory requirements that are automatically enforced by the system. And each will be accessible only through stringent security measures such as two-factor authentication, next-generation firewalls, and intrusion detection systems.

“Every project has its own requirements, and when this solution is in place, we'll be able to basically tune these enclaves and give researchers their own private data center with storage, compute, and desktop,” Palen says.

---

### The Promise: An Agile, Cost-Effective, Software-Defined Data Center

Palen's group will no longer have to stand up a brand-new traditional IT architecture for every project. Provisioning time will drop from two weeks to mere minutes.

Using SDI, provisioning compute resources will become plug and play, he explains. “We'll have the ability to provision a project from our desk. There's no calling central IT and getting an IP and changing VLAN at



the switch, and doing bare metal and kick-starting it. We don't have to do any of that."

Provisioning will also cost less with SDI, because each project will no longer require an entire dedicated machine. "On the same physical footprint, we will be able to install three or more researchers at the same type of scale," says Palen. "And not be burning power and cooling and data center space and capital costs."

That means ARC expects to be able to fulfill more service requests, with fewer staff and no new hardware.

"In the past, we couldn't consider setting up an architecture unless a large number of faculty members were going to use it, and they were going to spend at least a million dollars a year," says Palen. "With this system, we should be

able to help more faculty, with a wider set of research interests and needs."

SDI can also make both managing compliance and enabling collaboration easier. ARC will be able to, in cooperation with its legal team, create policy for each enclave, and the system automatically enforces that policy.

Plus, Yottabyte storage technology provides near-zero-footprint snapshots of disk volumes. If adopted system-wide, that means no more duplicating petabytes of data. ARC will be saving on staff time, hardware, and storage.

Says Palen, "With the Intel-Yottabyte solution, we should be able to serve a larger group of faculty, in a more productive way, with a shorter turnaround time."



### The ARC SDI Solution Is Composed Of:

- Intel® server chassis
- Intel® Xeon® processor E5 platform, including the new Intel Xeon processor E5 v4 family
- Intel® Ethernet 10 Gigabit Converged Network Adaptors
- Intel® Solid-State Drives
- Yottabyte\* yCenter software



“WE SHOULD BE ABLE TO SERVE A LARGER GROUP OF FACULTY, IN A MORE PRODUCTIVE WAY, WITH A SHORTER TURNAROUND TIME.”

—Brock Palen, Associate Director of Advanced Research Computing—Technology Services, University of Michigan



**ADVANCED  
RESEARCH COMPUTING  
TECHNOLOGY SERVICES**  
UNIVERSITY OF MICHIGAN

### University of Michigan Research at a Glance

- Ranked as the top public research university in the United States by the National Science Foundation<sup>1</sup>
- 101 top 10–ranked grad programs by *U.S. News & World Report*<sup>2</sup>
- 422 new inventions in 2015
- \$1.3B in research expenditures

### Next Steps: Continuing the Transformation

While the ARC Intel-Yottabyte pilot is still in its early stages, the solution's potential value in cost-effectively expanding access to technology is clear.

Essentially, the Intel-Yottabyte SDI pilot promises to open up access to ARC's advanced computing resources by removing the cost barriers associated with building a traditional architecture. The potential benefit of this approach is that it can be realized not just at ARC and the U-M, but across any institution, organization, or enterprise.

“We expect to migrate bare-metal services over time to SDI, to further consolidate our data centers, freeing power and cooling for other projects,” Palen says.

“We're looking to really make the most of the hardware we have available, in a cost-effective way.”

And deriving maximum value from research dollars and existing infrastructure is increasingly critical as U-M ramps up its Data Science Initiative (DSI).

DSI will expand opportunities for student and faculty researchers to tap into big data analytics in the areas of transportation, health sciences, learning analytics, and social science.

ARC will be ready. “We expect to increase the impact on research by bringing more computational resources to attack research topics,” Palen says. ■

Increase data center agility, empower new capabilities, and reduce CapEx and OpEx with an Intel based SDI solution.

Learn more about software-defined infrastructure at [intel.com/cloud](http://intel.com/cloud).

Also, read [How Software-Defined Infrastructure Is Evolving at Intel](#).

Find solution briefs, proof-of-concept publications, reference architectures, and vendor connections at [Intel® Builders](#).



ADVANCED  
RESEARCH COMPUTING  
TECHNOLOGY SERVICES  
UNIVERSITY OF MICHIGAN



1 Rankings by Total R&D Expenditures. National Science Foundation.

<https://ncesdata.nsf.gov/profiles/site?method=rankingBySource&ds=herd>

2 "Best Colleges Rankings." *U.S. News & World Report*.

<http://colleges.usnews.rankingsandreviews.com/best-colleges/university-of-michigan-9092>

Copyright © 2016 Intel Corporation. All rights reserved. Intel, the Intel logo, the Intel. Experience What's Inside logo, Intel. Experience What's Inside, and Xeon are trademarks of Intel Corporation in the U.S. and/or other countries.

\*Other names and brands may be claimed as the property of others.

