

OOI Facilities Board Meeting
13-14 October 2020



Oregon State
University
OOI-CI DATA CENTER

Toward an Extensible Cyberinfrastructure for the Ocean Observatories Initiative

Anthony Koppers (Lead-PI, Professor, Associate Dean for Research)

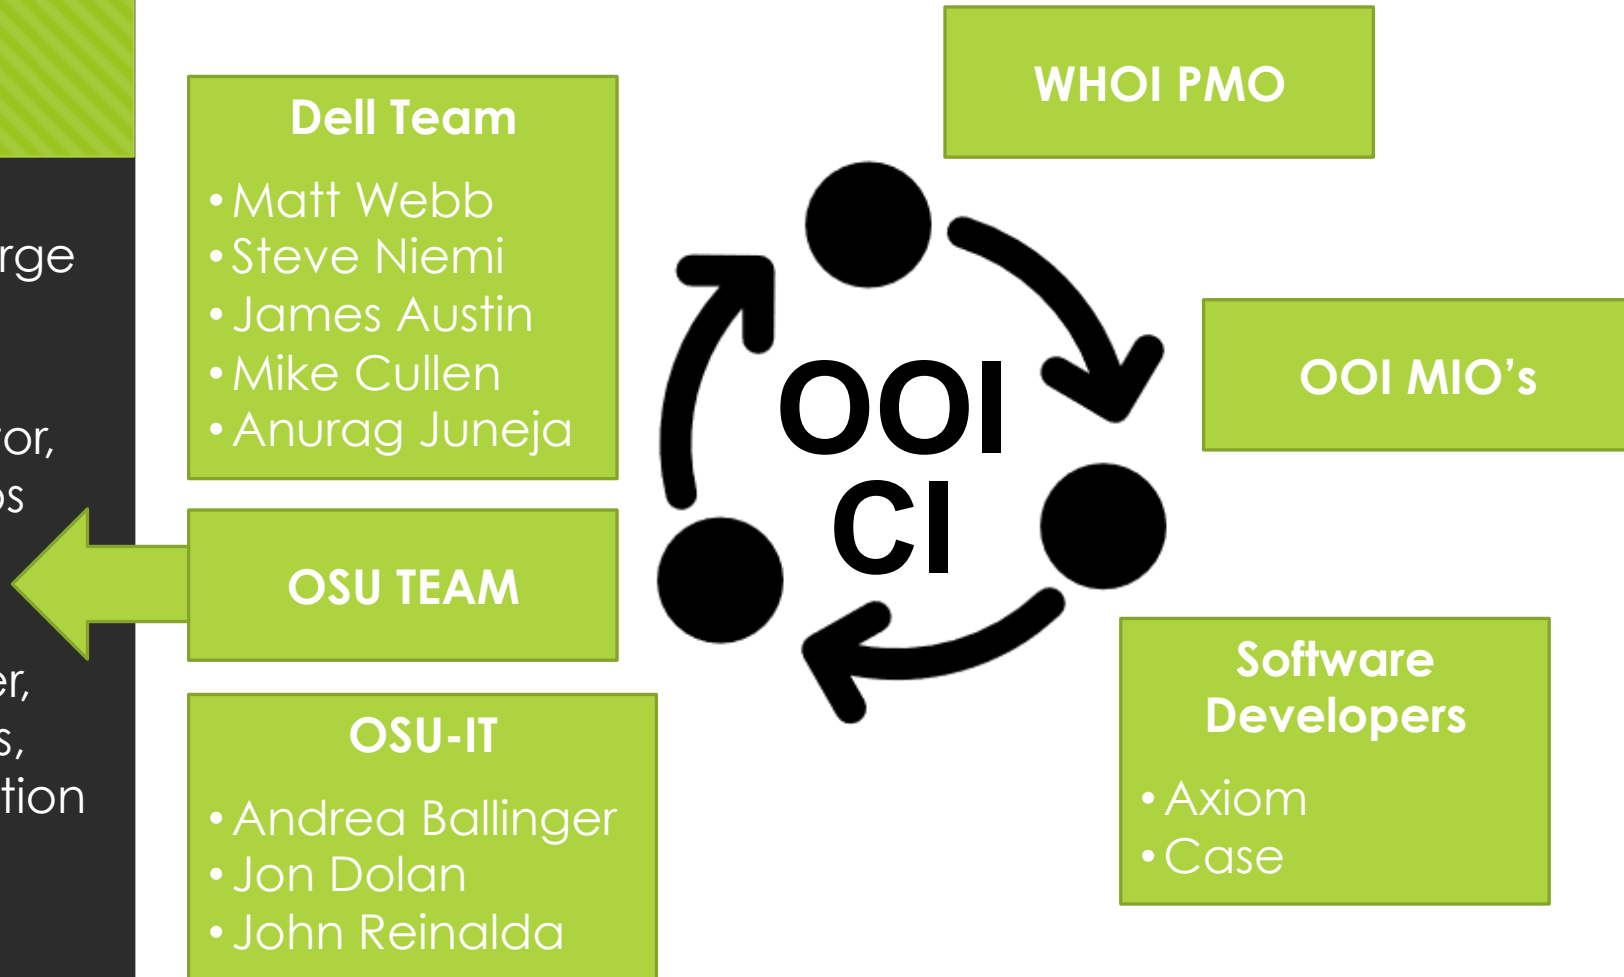
Mark Kever (Co-PI, Director of Research Computing)

Craig Risien (Operations Manager, OOI Endurance Array Engineer, Data Portal POC)

The Core OSU OOI-CI Data Center Team

- **Anthony Koppers** – Leadership, Large Facilities Partnerships, Data Governance, Marine Geology
- **Mark Kever** – Data Center Director, Supervision of Sys Admin & DevOps Engineers, Vendor Relations, Technological Vision, Reporting
- **Craig Risien** – Operations Manager, Reporting, Documentation, Tickets, Team Interfacing with OOI, Promotion using OOI Data
- **IT System Architect** – To be Hired
- **DevOps Engineer** – To be Hired

The Overall OOI-CI Structure



Anthony Koppers

Lead PI



- PhD 1998
- Professor in Marine Geology
- Associate Dean for Research
- MagIC Database

Anthony Koppers brings to the table extensive knowledge in seagoing marine geology, large facility leadership in the IODP, expertise in managing the EarthRef.org and MagIC geoinformatics consortia, and administrative knowhow based on his Associate Dean role.

He will provide leadership and thorough data stewardship for the new OOI CI data center.

Mark Keever

Co-PI



- MS 2007
- Director of Research Computing
- XSEDE, SIGHPC
- Storage Technology Development

Mark Keever has worked in research computing for over 20 years, joining CEOAS in 2019. Keever is the Director of Research Computing for CEOAS, and Director of Research Technologies for OSU.

He will provide project administration and management of the new OOI CI data center.

Craig Risien

Operations Manager



- MS 2002, 2006
- Senior Faculty Research Asst. II
- NANOOS, SCOW
- Endurance Array Technician

Craig Risien has worked as a faculty research assistant for CEOAS since 2006. His work focuses on data product develop and coastal oceanography. He is active as an OOI Endurance Array technician and data team member as well as data management for NANOOS.

He will assist with the build and operation of the new OOI CI data center.

Our Objectives

- Provide a low-risk and cost-effective OOI-CI Data Center.
- Introduce large increases in compute power, modernize storage solutions, and make use of cloud services for backup and disaster recovery.
- Achieve a seamless transition from the OOI-CI operations at Rutgers University to OSU in Year 1.
- Focus on extensibility in Years 2 & 3 by considering both bare-metal and/or cloud solutions based on OOI usages and needs.

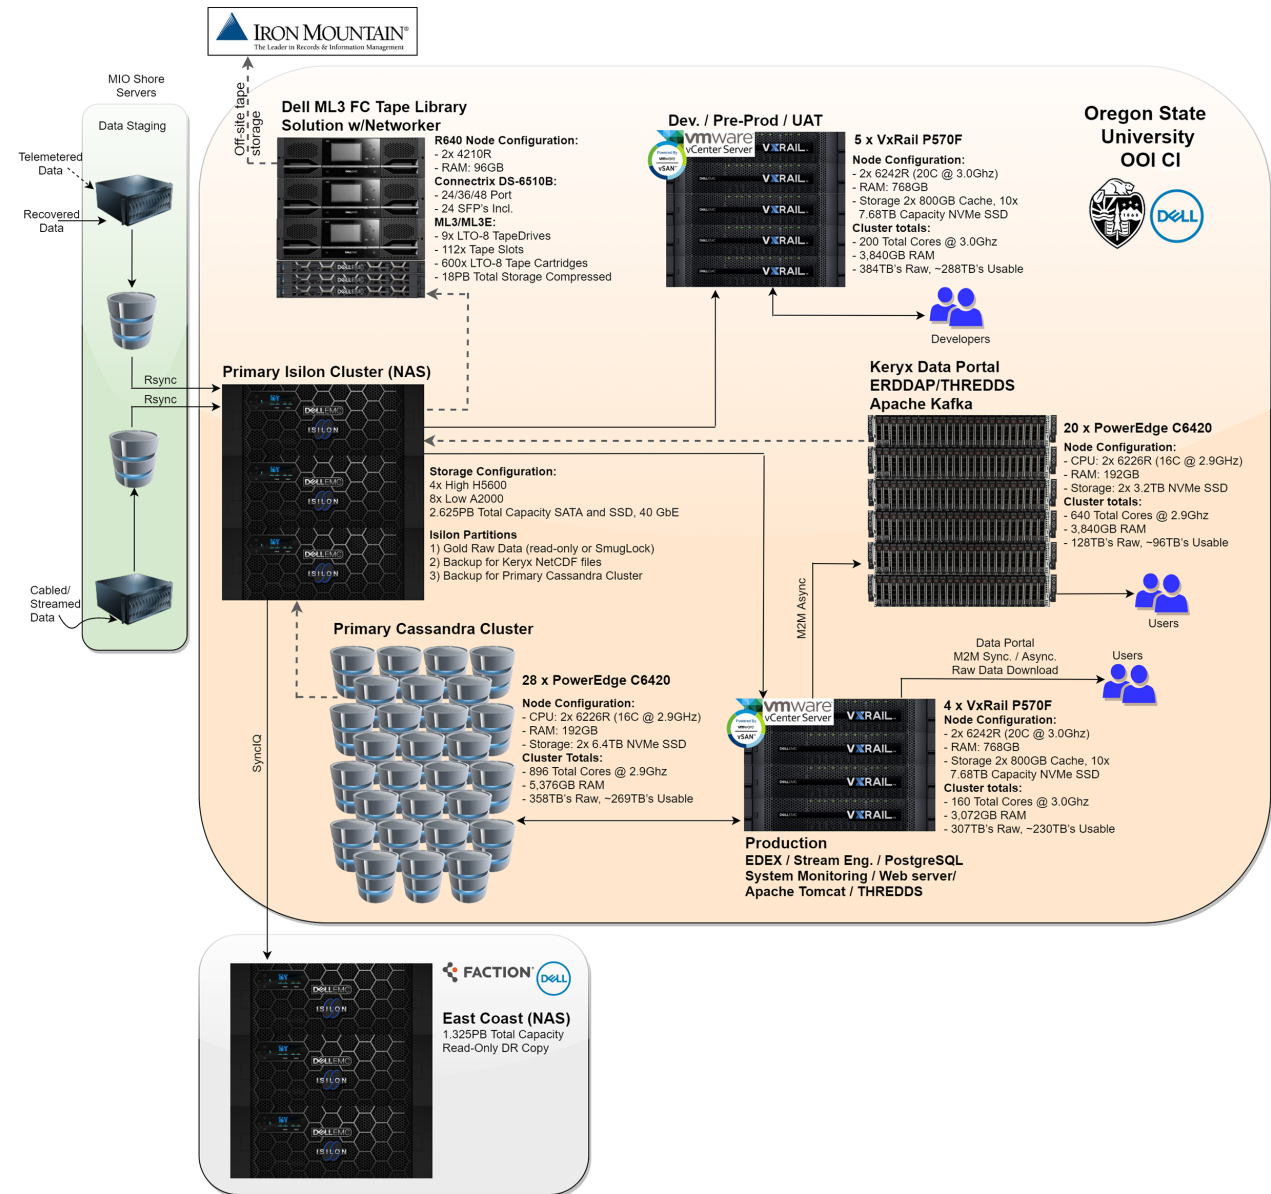
Recipe for a Successful Transition

- Install Dell Isilon and PowerEdge C6420 servers in a tightly coupled infrastructure, to provide a modern bare-metal secure scale-out NAS solution for storage, and fast, low-latency data processing.
- Hardware will be delivered completely installed by Dell EMC, except for setting up network tables and IP ranges.
- All prod and dev/test workloads will be virtualized with VMware running on two hyperconverged VxRail P570F platforms using vSAN/vSphere and these will allow us to replicate the Rutgers servers readily.
- All infrastructure uses dual 25 GbE connections against a 100 GbE spine with duplicative 100 Gbps connectivity to Internet 2.0 for low-latency replication of the Isilon storage on the East Coast with Faction Services.
- PI's have wide experience in international collaborative large facilities and geoinformatics, CEOAS and OSU IT systems, and domain expertise in OOI data operations.

Modernized, Simpler More Capacity, Faster

- Five Bare-metal Servers
 - Primary Isilon NAS (2.6 PB)
 - Prod VxRAIL VM Server (5 nodes)
 - Test/Dev VxRAIL VM Server (6 nodes)
 - Cassandra Cluster (28 nodes)
 - Keryx Data Portal (20 nodes)
- Secondary Isilon Server
 - East Coast NAS
 - Active:Active replica
- Tape Library (19 PB)
 - Incremental and Snapshot Backup
 - Offsite storage

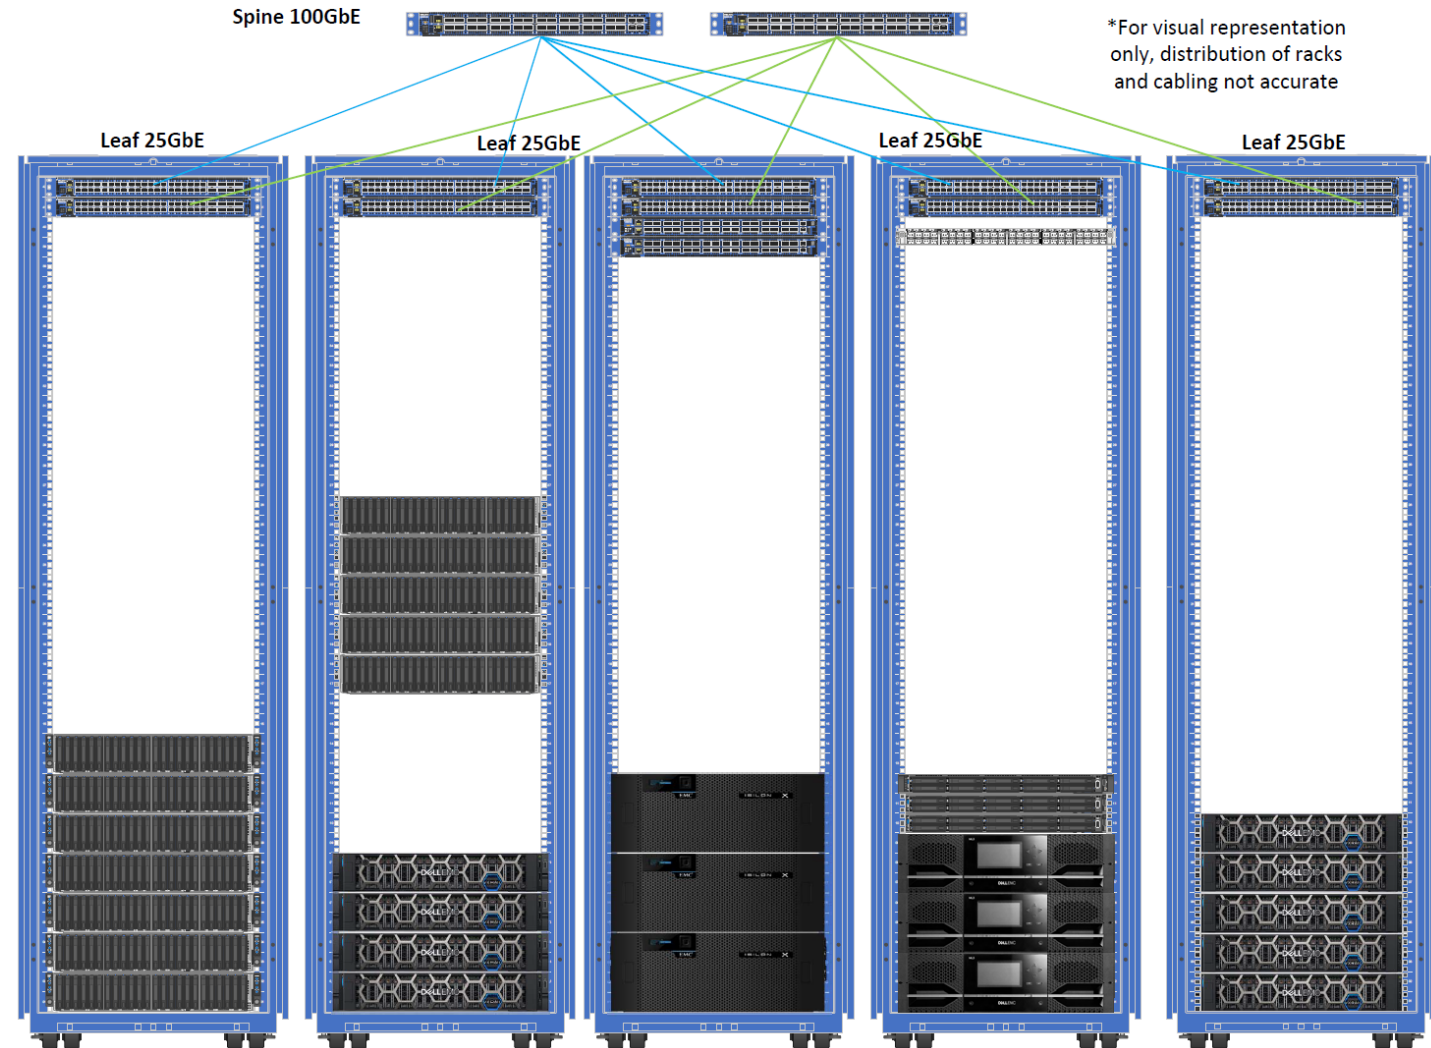
New OSU Architecture



Modernized, Simpler More Capacity, Faster

- Condensed Infrastructure
 - Fewer bare-metal servers
 - Single-point management control
- Modernized Storage
 - Emphasis on NVMe SSD Storage
 - Isilon 3x Scale Out Storage
- Larger Cassandra Cluster
 - More and faster nodes
 - Only 22% Initial Storage Capacity
- Faster Connections
 - Dual 25 GbE for 50 GbE throughput
 - LinkOregon Redundant Fiber to Internet 2.0
- All compute using VMware

New OSU Architecture



Data Stewardship and Allied Large Facilities

- All NSF-funded data:
 - Should be FAIR by making data Findable, Accessible, Interoperable, and Reproducible.
 - Should be attributed to its end users by minting DataCITE DOI's and authentication using ORCID IDs.
 - Should be stored for longevity for generations to come, e.g. via tape backup and offsite storage.
- All OOI data:
 - Should be of interest and used by users outside its own community.
 - Should be made available readily to enable science.

FAIR Data, DOI's and ORCID

- PI Koppers will bring cyberinfrastructure experience in setting up certified FAIR data repositories.
- He also provides experience in applying schema.org and geocodes to make data better discoverable.
- He has been minting DataCITE DOIs for online data sets.
- He has established a contract with ORCID in order to use their authentication API for user registration and log-in/out to manage data upload and ownership.
- Overall, OSU will contribute to the OOI-CI by providing transparent data governance ensuring adherence to NSF and NSB data policies and standards, quality control of OOI data stores, and adequate response to cyber-infrastructure issues arising among OOI stake holders and end users.

Data Stewardship and Allied Large Facilities

- All NSF-funded data:
 - Should be FAIR by making data Findable, Accessible, Interoperable, and Reproducible.
 - Should be attributed to its end users by minting DataCITE DOI's and authentication using ORCID IDs.
 - Should be stored for longevity for generations to come, e.g. via tape backup and offsite storage.
- All OOI data:
 - Should be of interest and used by users outside its own community.
 - Should be made available readily to enable science.

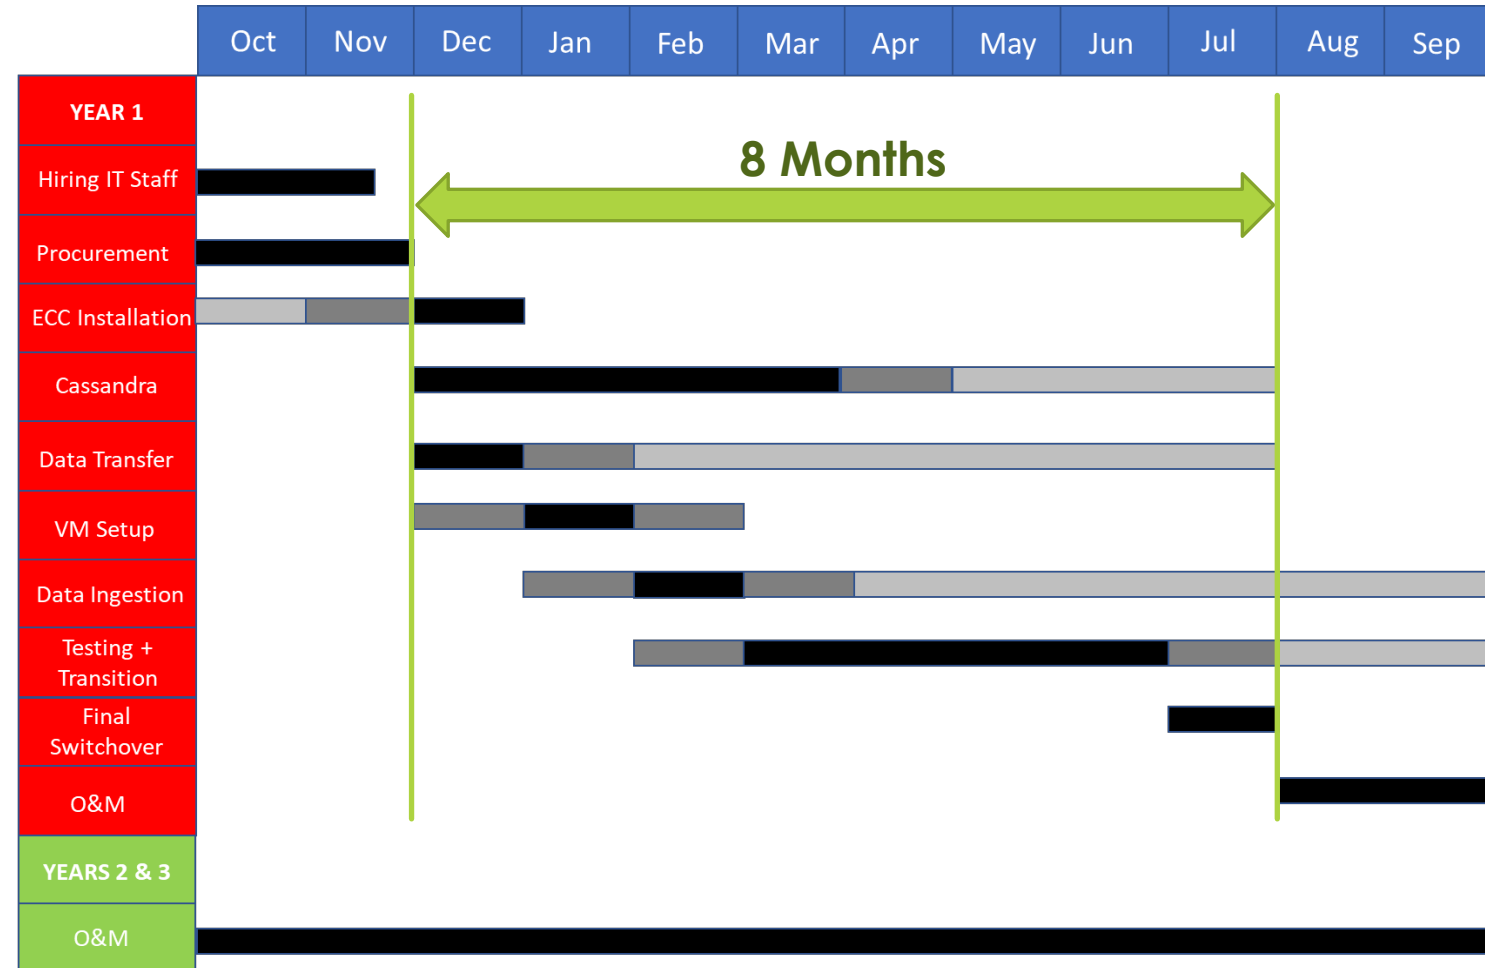
Relation to Other Large Facilities

- PI Koppers will bring to OOI his viewpoint that large facilities are paramount in the success of all modern-day ocean sciences (e.g. UNOLS, IODP, OOI, OBS, IRIS).
- He will help increase partnerships with other international observatory programs or with initiatives that would require future observatory science implementations.

Transition from Rutgers to OSU Data Center

- First two months:
 - Fast procurement by Dell and OSU.
 - Expedited hiring by OSU.
 - Install by Dell EMC to be finished by the beginning of December.
 - Zoom meetings with Rutgers+WHOI are critical to collaboratively prove the transition timeline and procedures.
- Procedures carried out will include:
 - Hardware testing by Dell during burn-in period, testing that will continue in normal daily monitoring routines.
 - Disaster Recovery testing of various backup and recovery scenarios.
 - Testing with WHOI and Rutgers.

Contracted Timeline



Improvements that we have focused on ...

- Transparency throughout.
- Versioning using Git integrated with vSphere and/or Ansible.
- Full benefit through hyperconverged virtualization.
- Increased cybersecurity.
- Single Isilon scale out storage solution.
- More capable Dev/Test environment.
- More capacity in Cassandra.

Central Management Platform

- Because of a simplified management platform, providing concise documentation, more meaningful monitoring, and enhanced metrics are readily within reach.
- Central management of Virtual Machines (VMs) and versioning of applications through e.g. Git will be easier and more efficiently managed.
- Again, because of the simplified architecture and more centralized management, access to all relevant VMware and data stores for administrators and programmers can be easily achieved and guaranteed.
- One simple scale out Isilon data file store, with integral 3x replication, and with Disaster Recovery (DR) duplication into cloud that also has 3x replication.

Improvements that we have focused on ...

- Transparency throughout.
- Versioning using Git integrated with vSphere and/or Ansible.
- Full benefit through hyperconverged virtualization.
- Increased cybersecurity.
- Single Isilon scale out storage solution.
- More capable Dev/Test environment.
- More capacity in Cassandra.

Hyperconverged Virtualization

- We have added significant compute resources to be able to clone VMs on the hyperconverged servers to add both more ability and capacity. Server nodes can be easily added to increase compute power.
- There now is ample VM capacity to launch more than one node of e.g. stream engine servers from the get-go or any of the other server for that matter.
- The Dell vSphere software that also is Ansible enabled will allow for a strong control on load balancing.
- In Year 1 we will focus on establishing VMs and use cloning of VMs to increase compute power; however, containerization (e.g. through Kubernetes) will be considered in Years 2 & 3 if indeed we need to expand compute onto servers in the cloud.

Improvements that we have focused on ...

- Transparency throughout.
- Versioning using Git integrated with vSphere and/or Ansible.
- Full benefit through hyperconverged virtualization.
- Increased cybersecurity.
- Single Isilon scale out storage solution.
- More capable Dev/Test environment.
- More capacity in Cassandra.

Cassandra Cluster Issues

- Strongly improved compute power will speed up replay significantly in populating the Cassandra cluster.
- This could be enhanced by launching multiple VMs to support replay in parallel queues.
- We have increased the number of nodes from 21 to 28 (▲33%) with an overall increase in storage (▲113%), the number of cores (▲166%), and total memory (▲300%).
- The OOI Cassandra database is highly concurrent as well as heavy-read. Therefore, the improved CPU power and 166% increase in the number of cores will be highly relevant as well as the large 300% increase in RAM, which results in less frequent memory flushes to disk.
- When writing is needed from memory to disk, the new system is extremely fast by using NVMe SSD only.
- If nodes need to be added, this is relatively inexpensive at roughly \$15k/node.

OOI Facilities Board Meeting
13-14 October 2020



Oregon State
University

OOI-CI DATA CENTER

Toward an Extensible Cyberinfrastructure for the Ocean Observatories Initiative

THANK YOU !!!

WE ARE HAPPY TO ANSWER QUESTIONS !!!