

Ceph for Tier-2 Storage at MSI

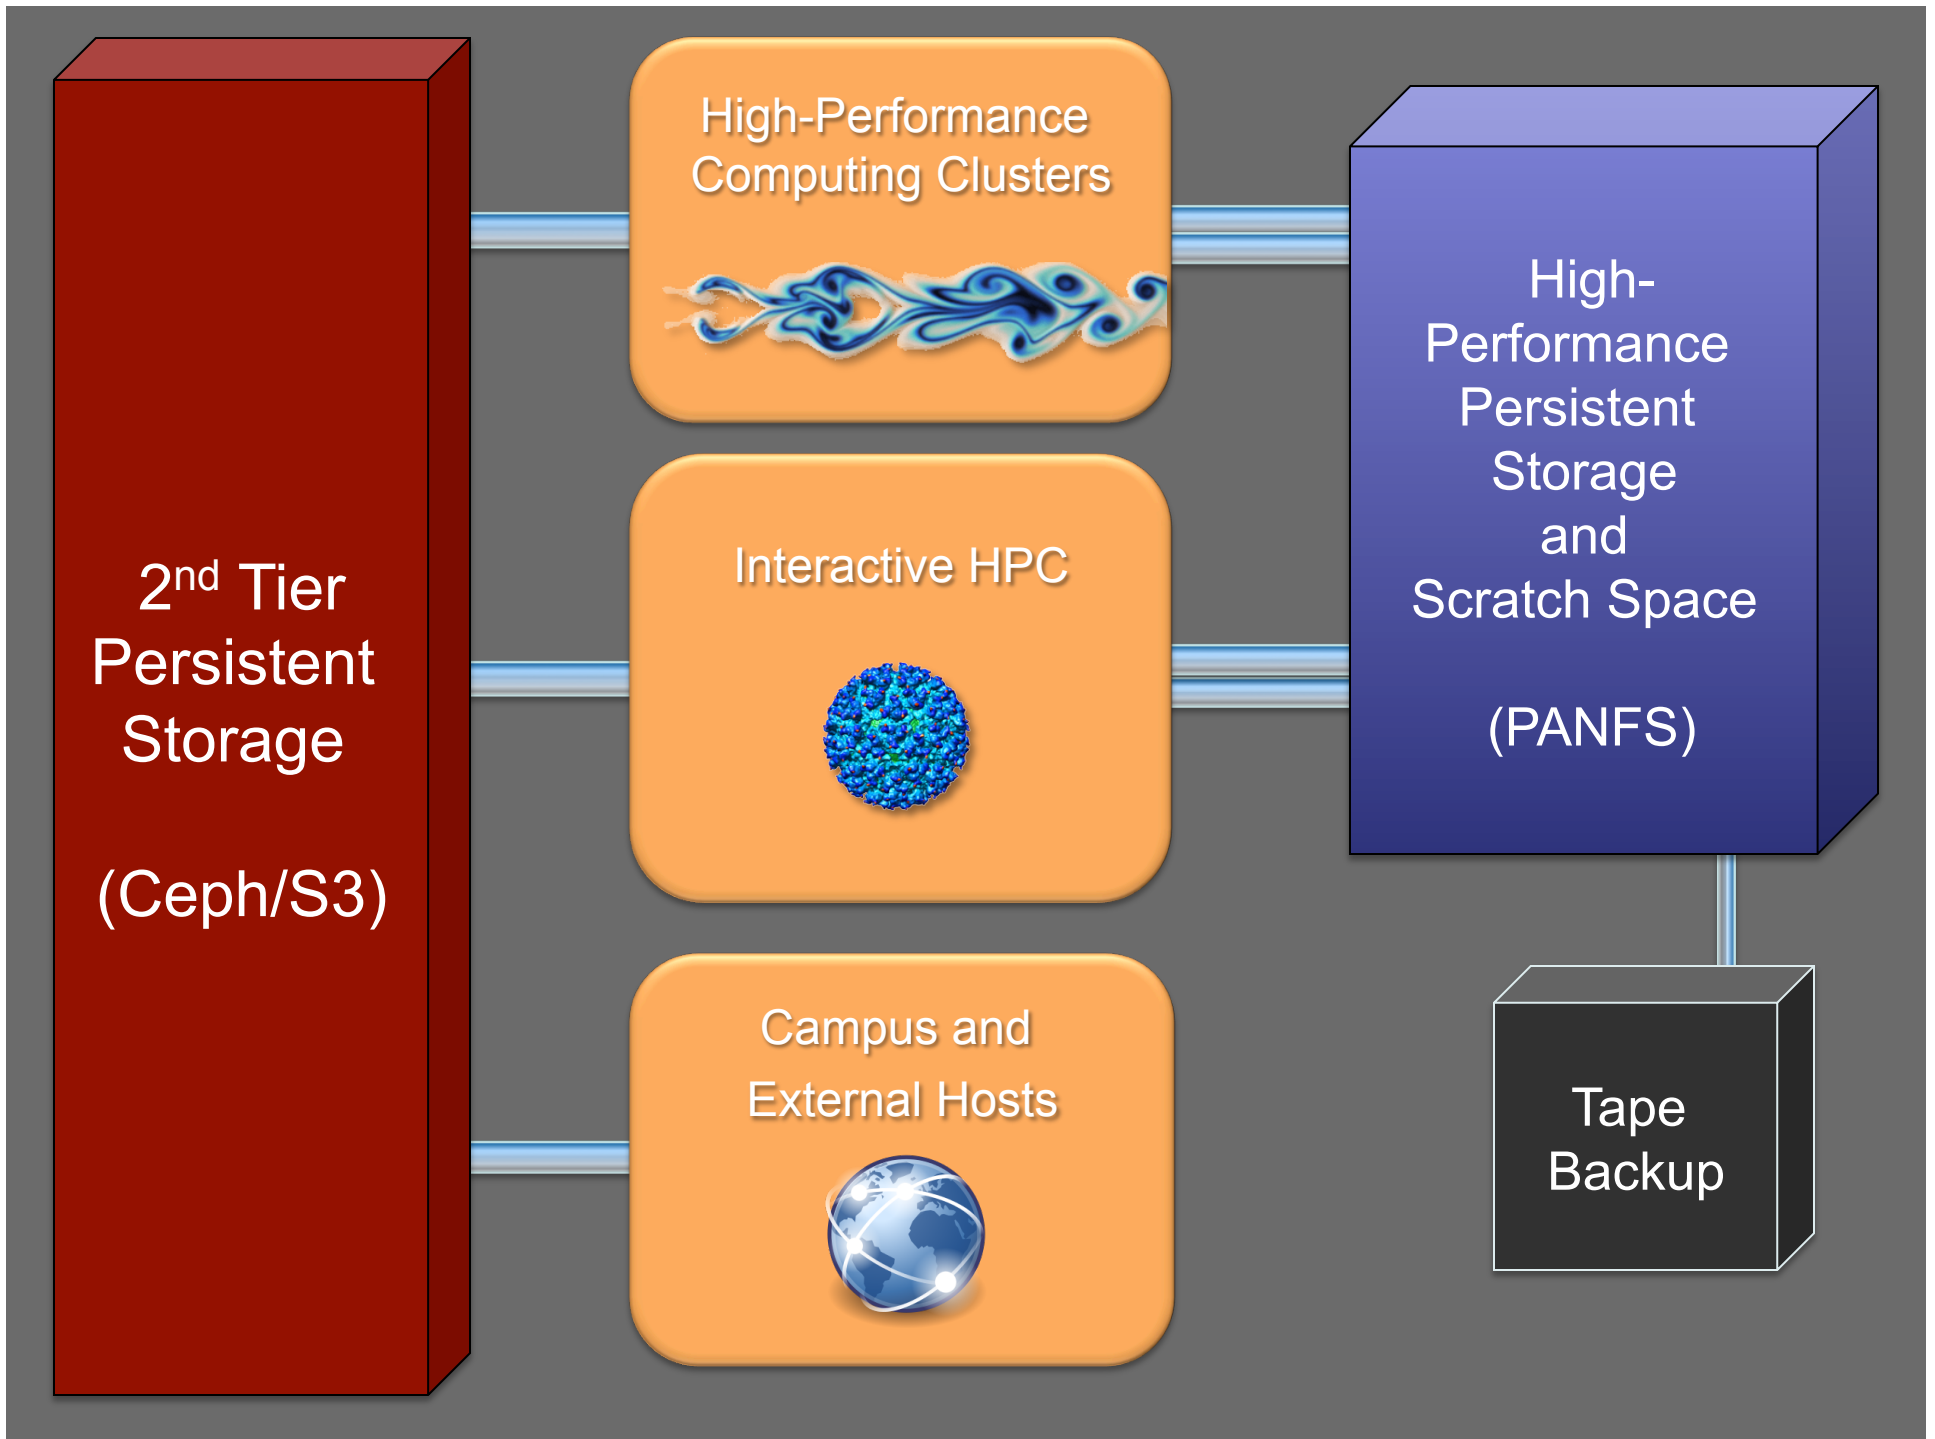
Benjamin Lynch, Ph.D.

© 2014 Regents of the University of Minnesota. All rights reserved.

Minnesota Supercomputing Institute
for Advanced Computational Research




UNIVERSITY OF MINNESOTA
Driven to DiscoverSM



High-Performance
Computing Clusters



Interactive HPC




Campus and
External Hosts



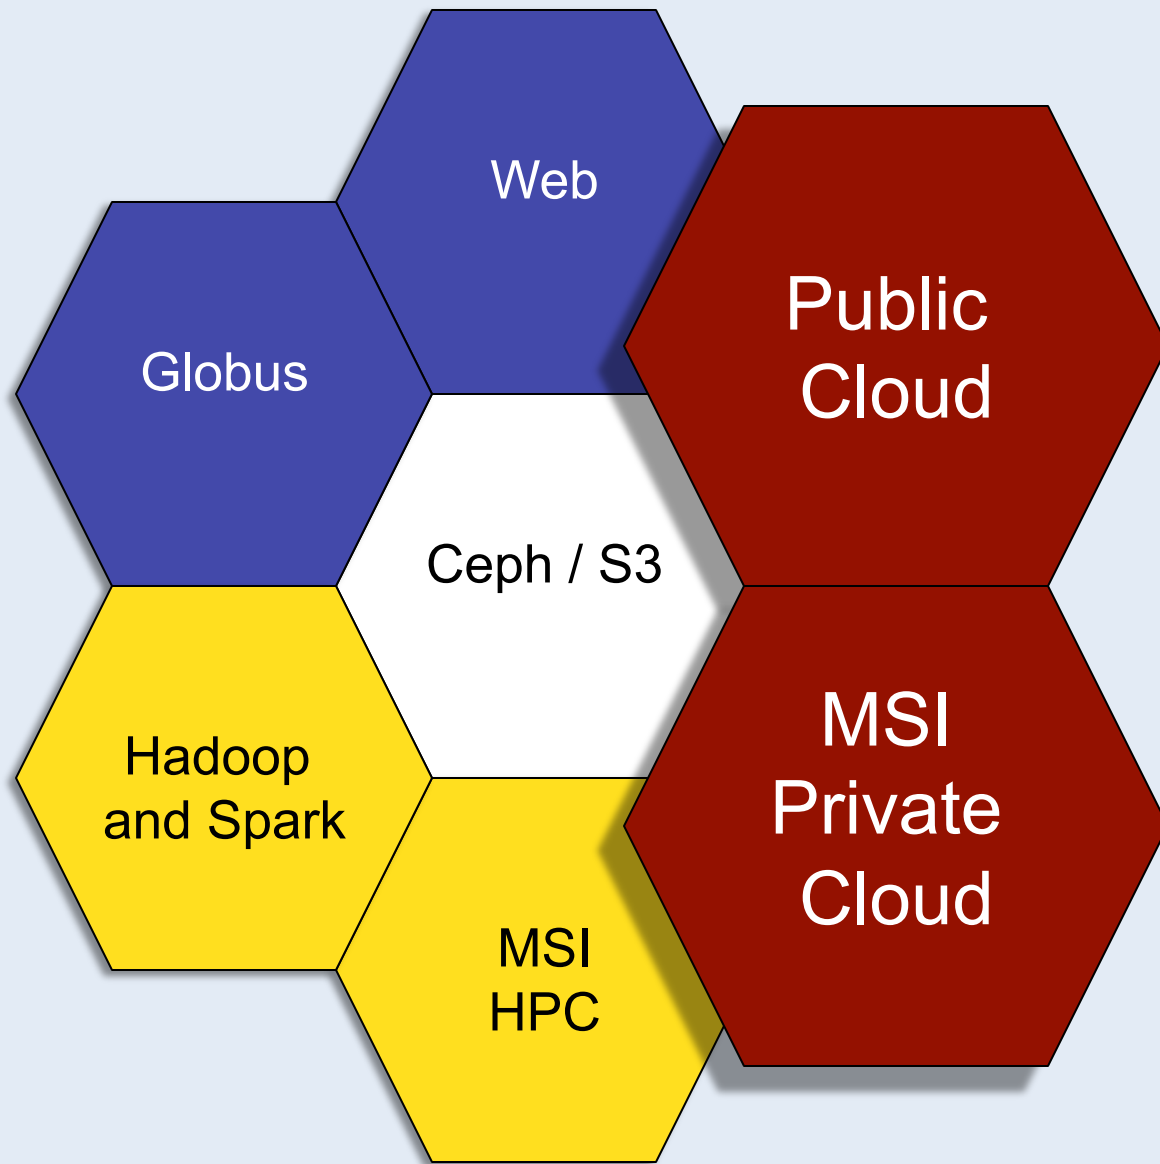
High-
Performance
Persistent
Storage
and
Scratch Space
(PANFS)

Tape
Backup

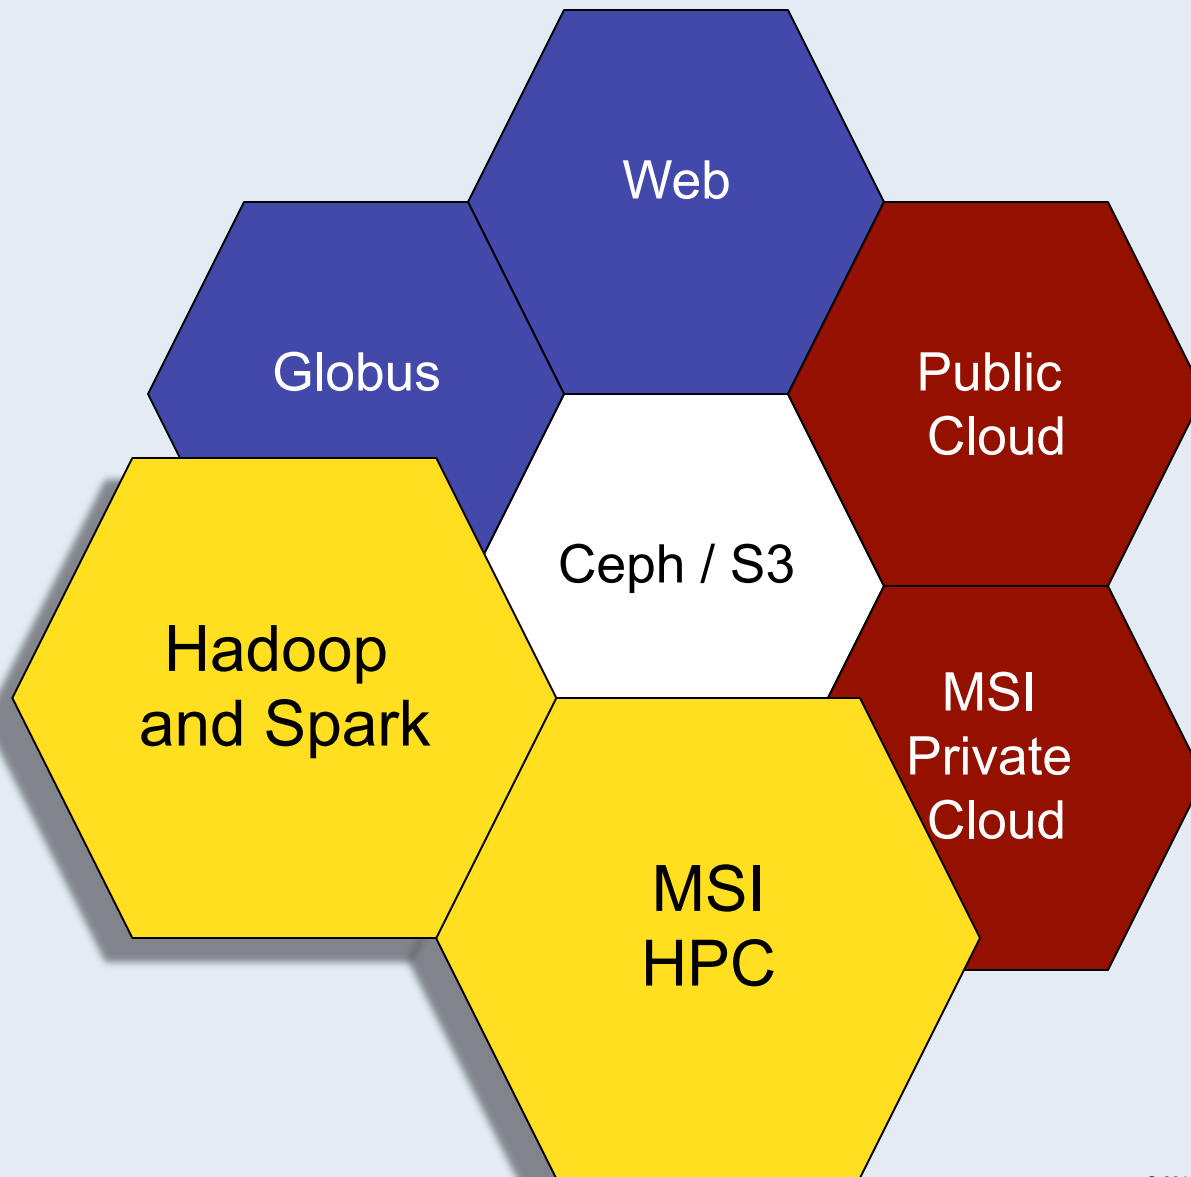
2nd Tier
Persistent
Storage
(Ceph/S3)



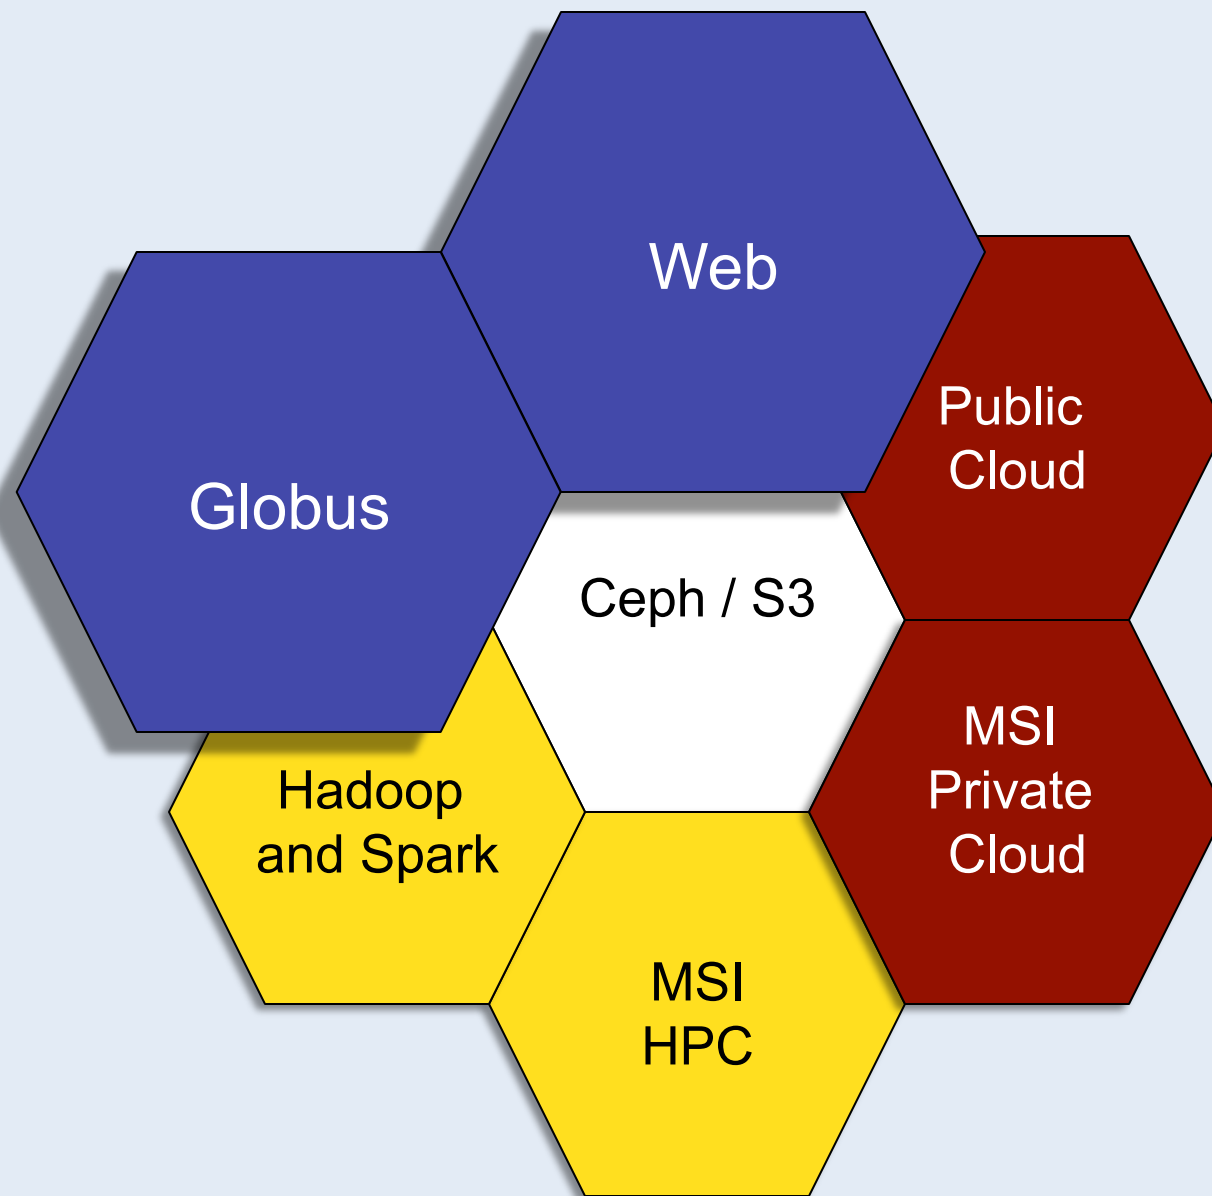
© 2014 Regents of the University of Minnesota. All rights reserved.



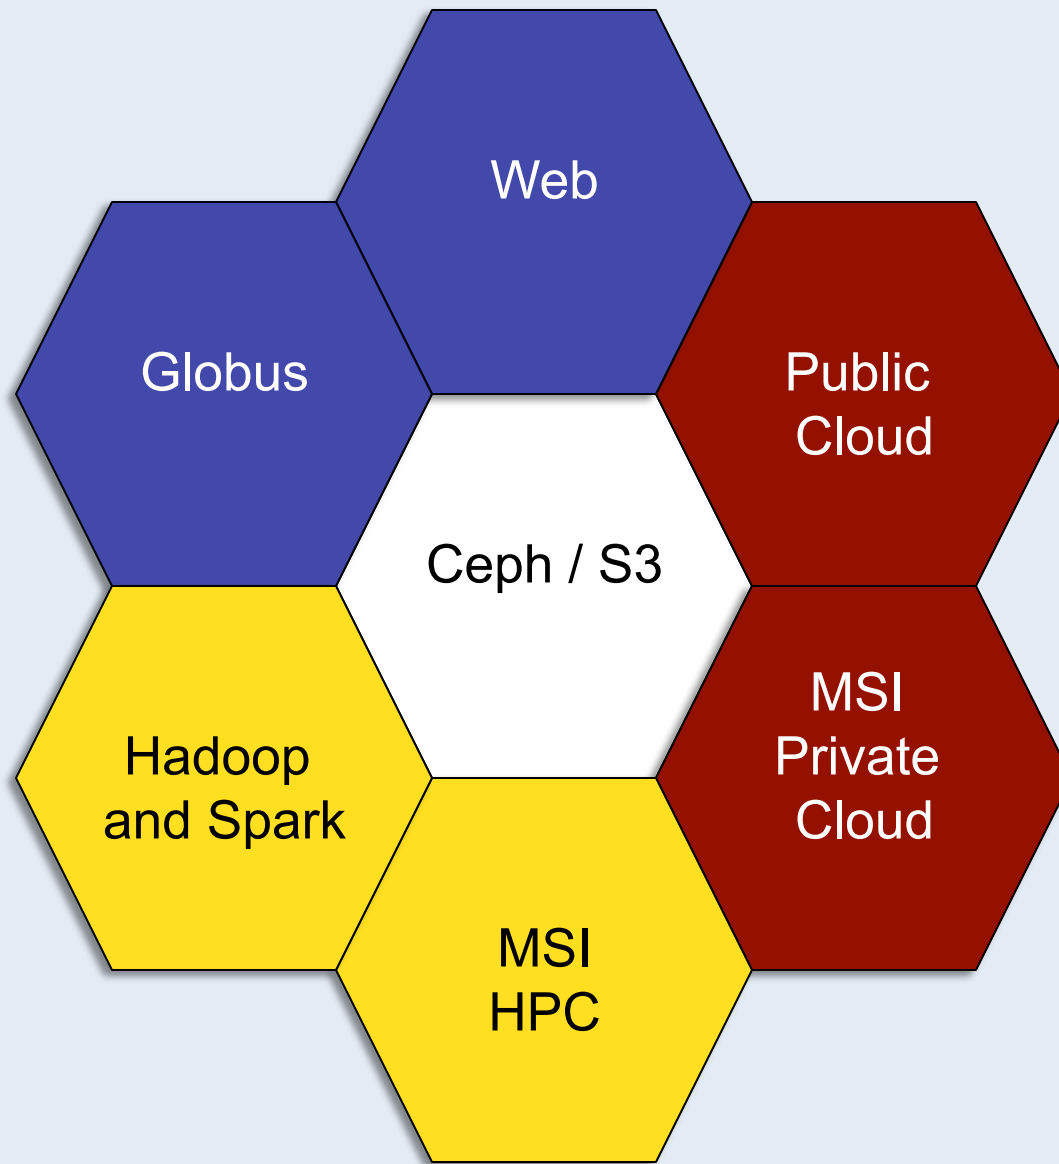
© 2014 Regents of the University of Minnesota. All rights reserved.



© 2014 Regents of the University of Minnesota. All rights reserved.

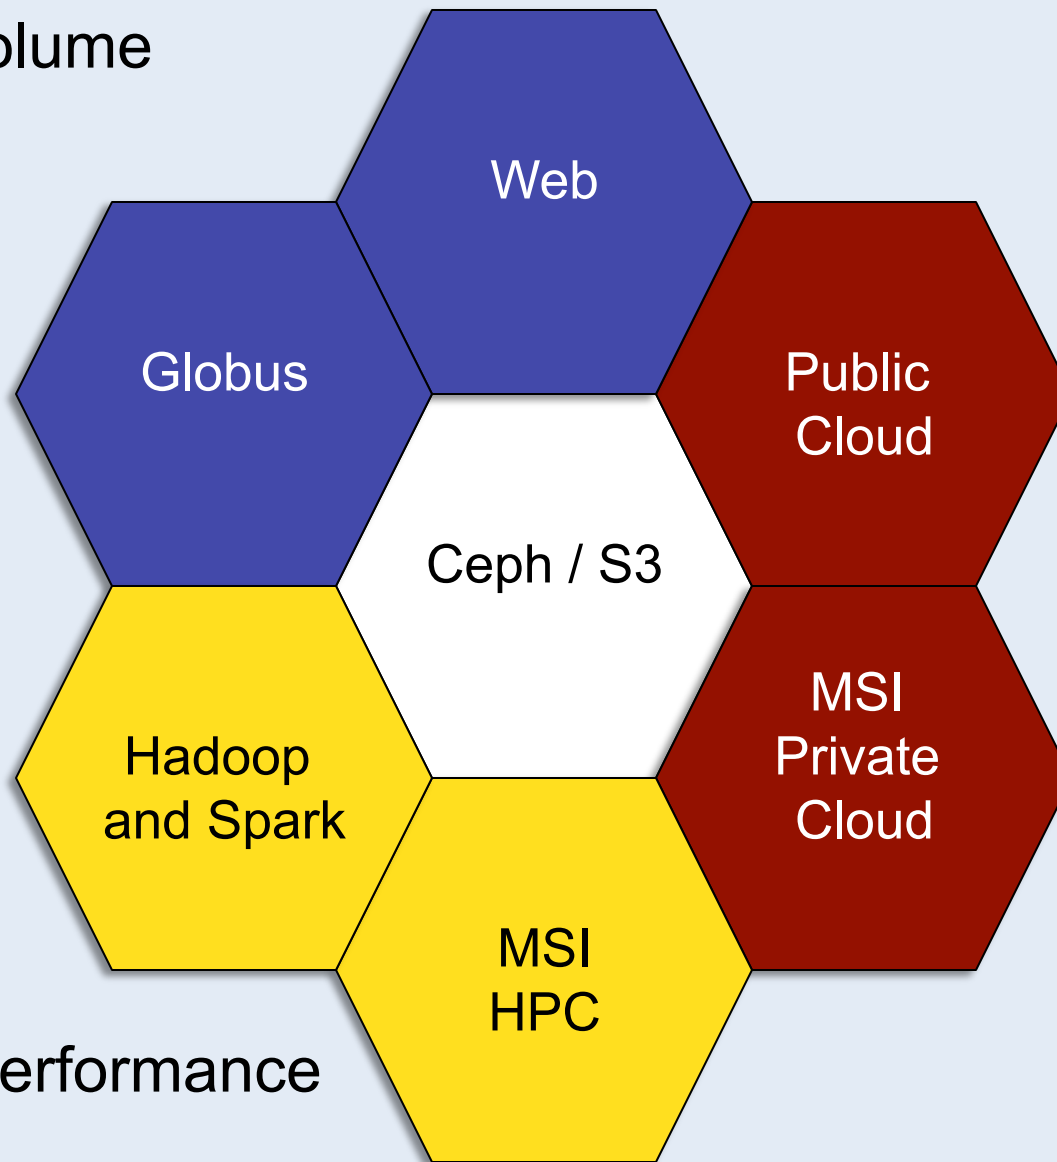


© 2014 Regents of the University of Minnesota. All rights reserved.



© 2014 Regents of the University of Minnesota. All rights reserved.

Scalable volume



High availability

Scalable performance

© 2014 Regents of the University of Minnesota. All rights reserved.

Why Ceph?

- Scalable
- Fault-tolerant
- Flexible configuration
- Flexible support

© 2014 Regents of the University of Minnesota. All rights reserved.



Our Ceph Cluster

- Cluster Specs
 - Ceph 0.94.5 (Hammer)
 - 3 PB raw disk
 - 4+2 erasure coding for most data
 - 4 load-balanced gateways (civetweb)
 - 10GigE network
 - 3 monitors
- Applications
 - Genomics workflows on public & private cloud compute resources
 - Data staging for HPC applications
 - Data staging for Hadoop / Spark cluster
 - Bulk store of stale data

© 2014 Regents of the University of Minnesota. All rights reserved.



Acknowledgements

- Luke Burns (System Administrator)
- Matt Mix (Network Administrator)

© 2014 Regents of the University of Minnesota. All rights reserved.

