



Forum **TERATEC** **23**

Unlock the future

31 MAY & 1 JUNE 2023 • At the Parc Floral, Paris

An event organised by

 **infoprodigital**





Hierarchical Storage Management for Exascale Applications in the IO-SEA Project

J. A. Grogan, K. O'Connor, C. Feller, B. Benek-Gursoy
Irish Centre for High End Computing (ICHEC)
james.grogan@ichec.ie



OLLSCOIL NA
GAILLIMHÉ
UNIVERSITY
OF GALWAY

Unlock the future

Overview

- Irish Centre for High End Computing (ICHEC)
- The IO-SEA Project
- Hierarchical Storage Management for Exascale Applications
- The IO-SEA HSM-Object Store Interface (Hestia)

Irish Centre for High End Computing



- Hosts **Kay** Supercomputer
- Upcoming EuroHPC JU Site – **CASPIr** Supercomputer
- **Activities:**
 - Quantum Computing
 - Machine Learning
 - Earth Observation
 - Performance Computing



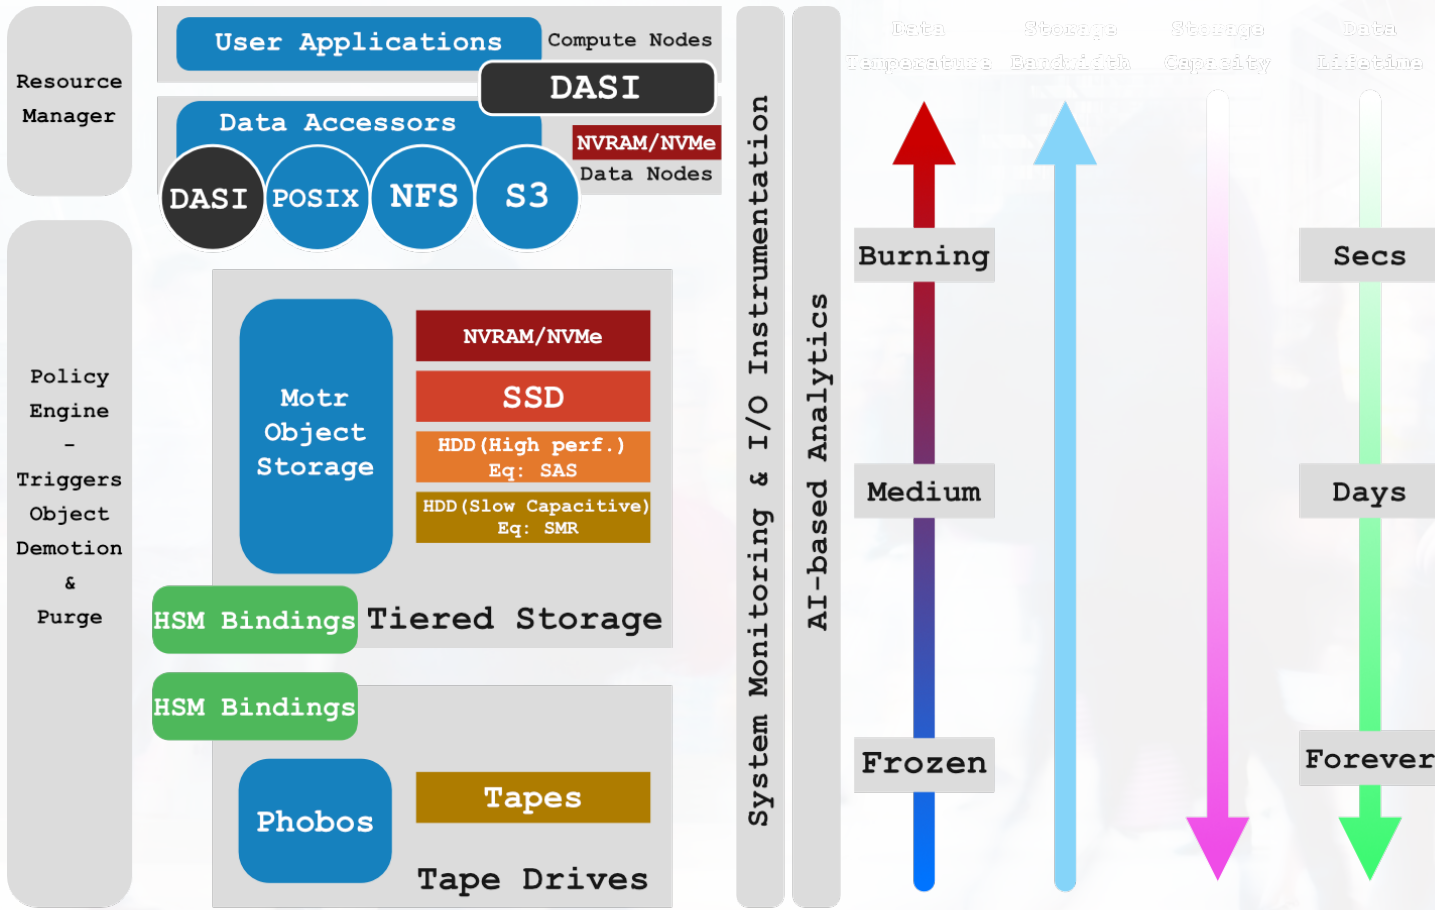
IO-SEA Project

**Unlock
the future**



IO-SEA Project

Storage I/O and Data Management for Exascale Architectures



IO-SEA Project - Partners

Storage I/O and Data Management for Exascale Architectures



IO-SEA Project - Foundations

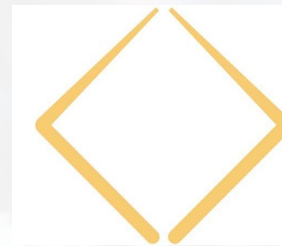
Storage I/O and Data Management for Exascale Architectures

Storage Abstraction
APIs

Ephemeral Data
Lifecycle Management

Advanced
Instrumentation

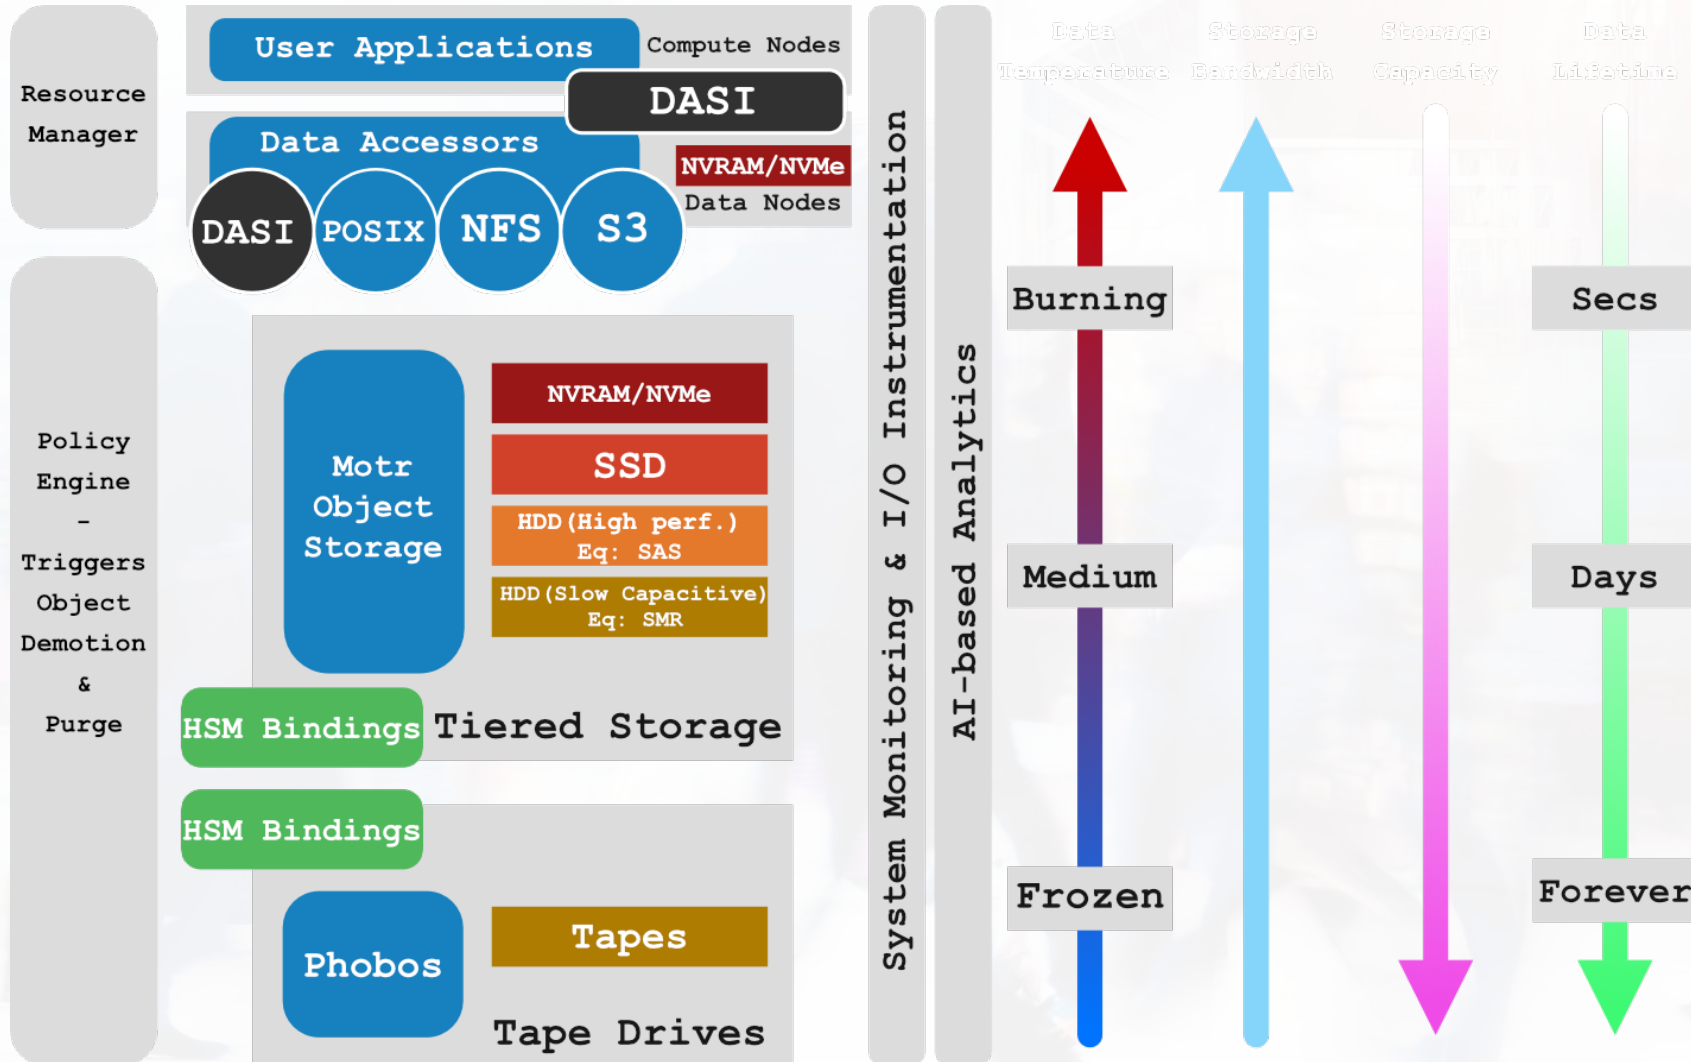
Multi-tiered Storage



MAESTRO
DATA ORCHESTRATION

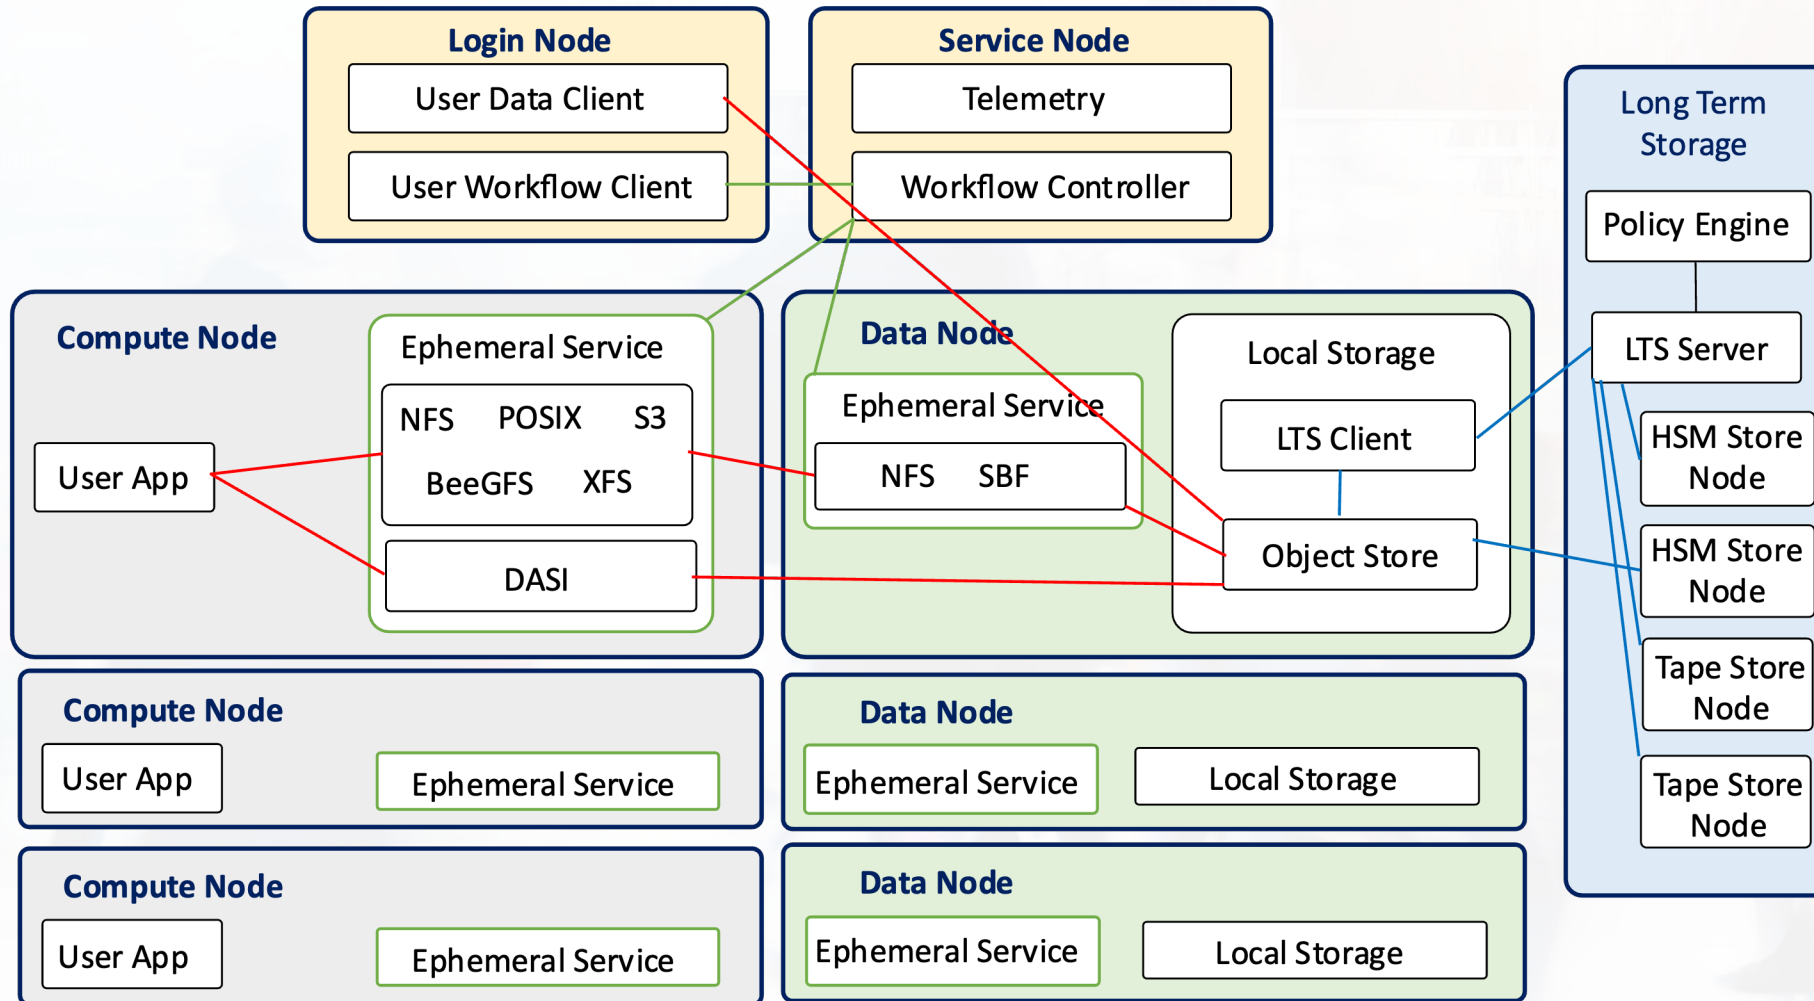
IO-SEA Project - Architecture

Storage I/O and Data Management for Exascale Architectures



IO-SEA Project - Architecture

Storage I/O and Data Management for Exascale Architectures



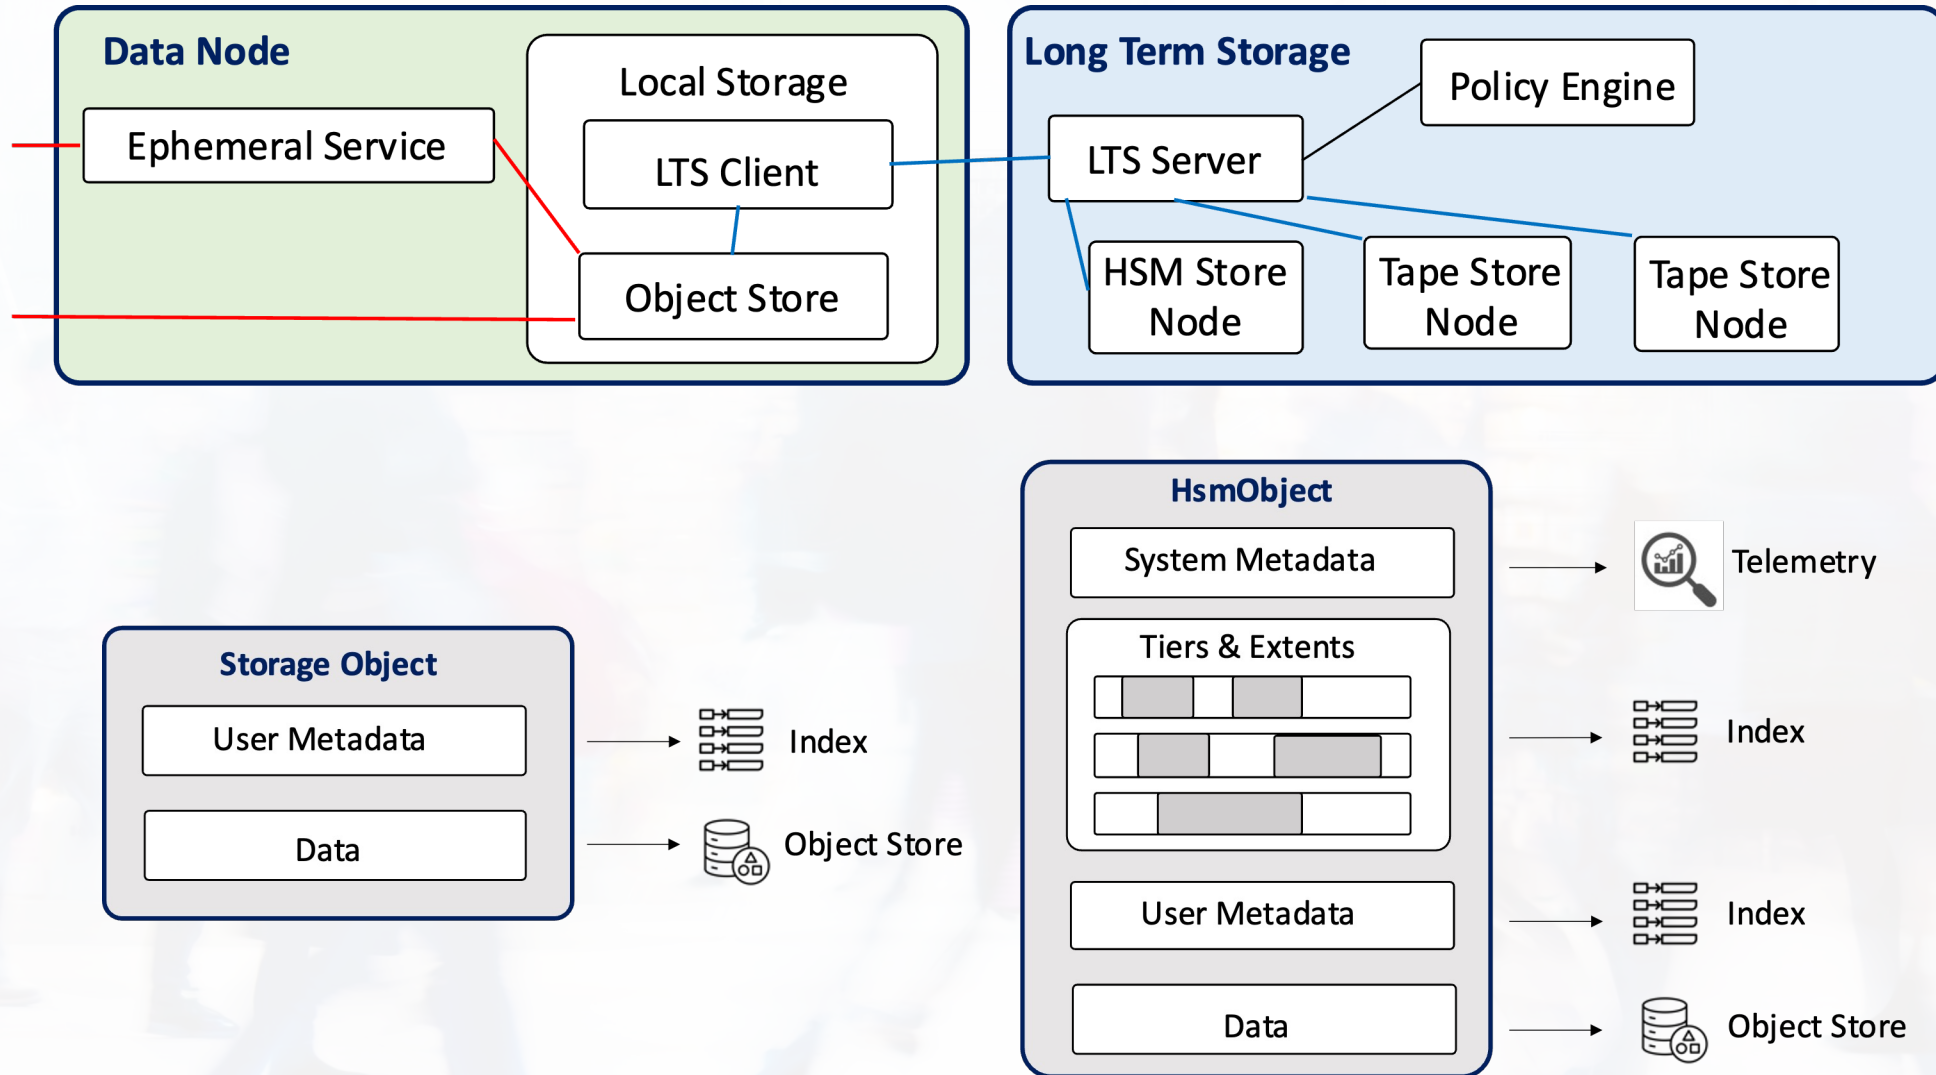


HSM for Exascale

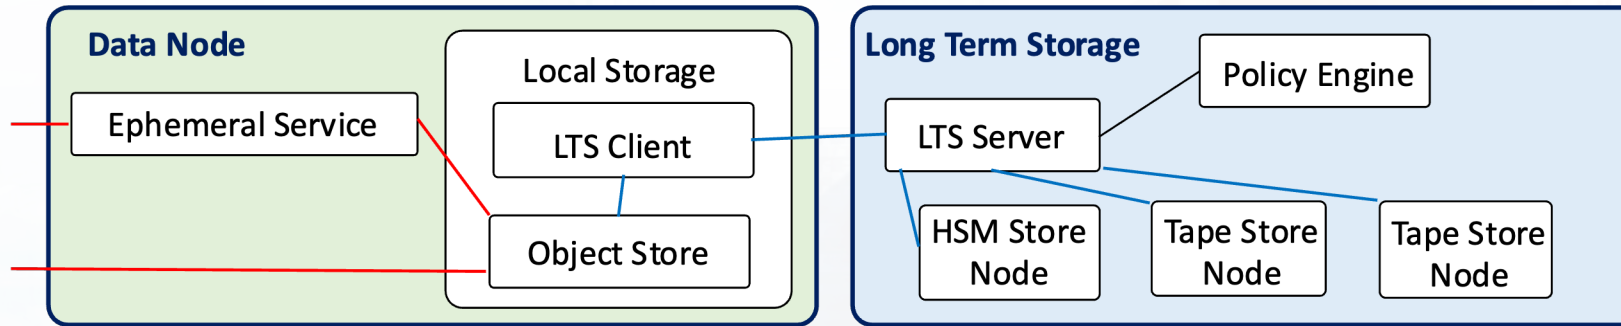
**Unlock
the future**



HSM for Exascale



HSM for Exascale

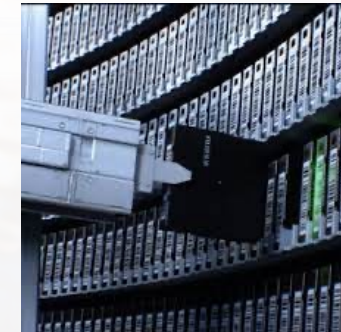


CORTX

Motr

Distributed Object Storage Prototype System

<https://github.com/Seagate/cortx-motr>

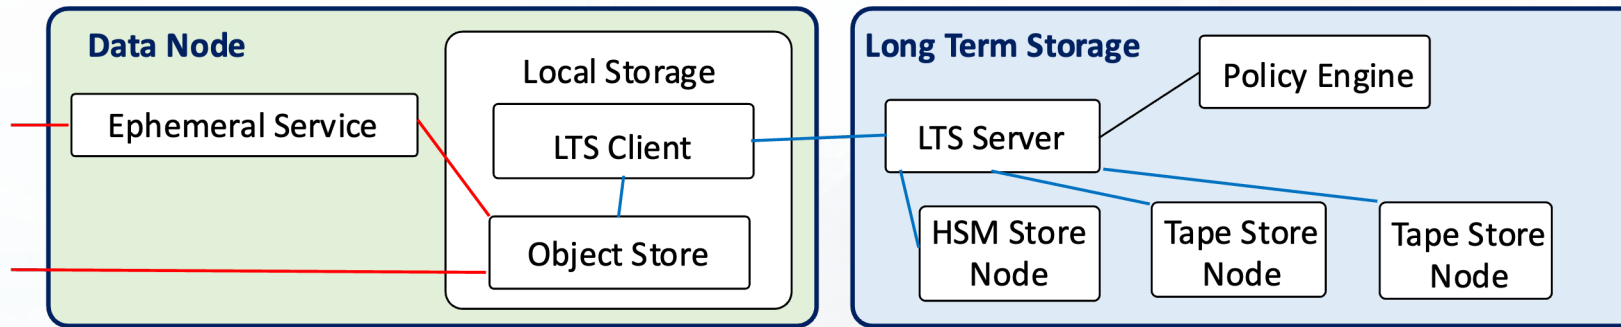


Phobos

Object Storage on Tapes

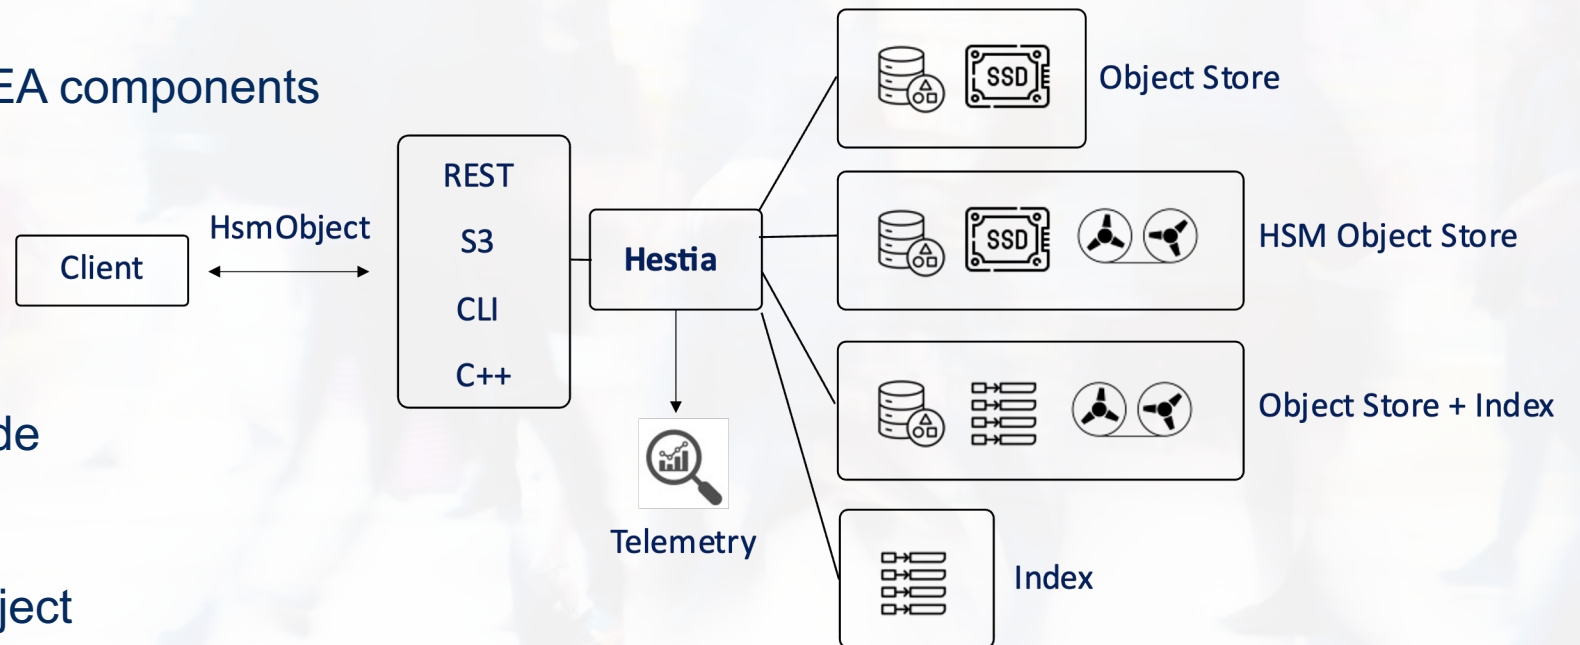
<https://github.com/cea-hpc/phobos>

HSM API for Exascale – Hestia Component



Goals:

- 1) Support interaction of IO-SEA components via a 'HSM API'
- 2) Minimal impact on:
 - Performance
 - IO-SEA component code
 - Introspection ability
- 3) Be extensible beyond the project





Hestia Component

**Unlock
the future**



Hestia Implementation

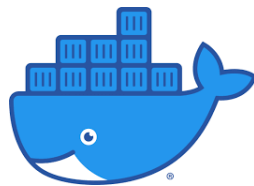
Modern C++ Library – CMake Build



Standard Dependencies and Formats

- Curl, LibS3
- Interchangeable WebServer
- Yaml/JSON
- Spdlog/CLI11

CI with containerized builds

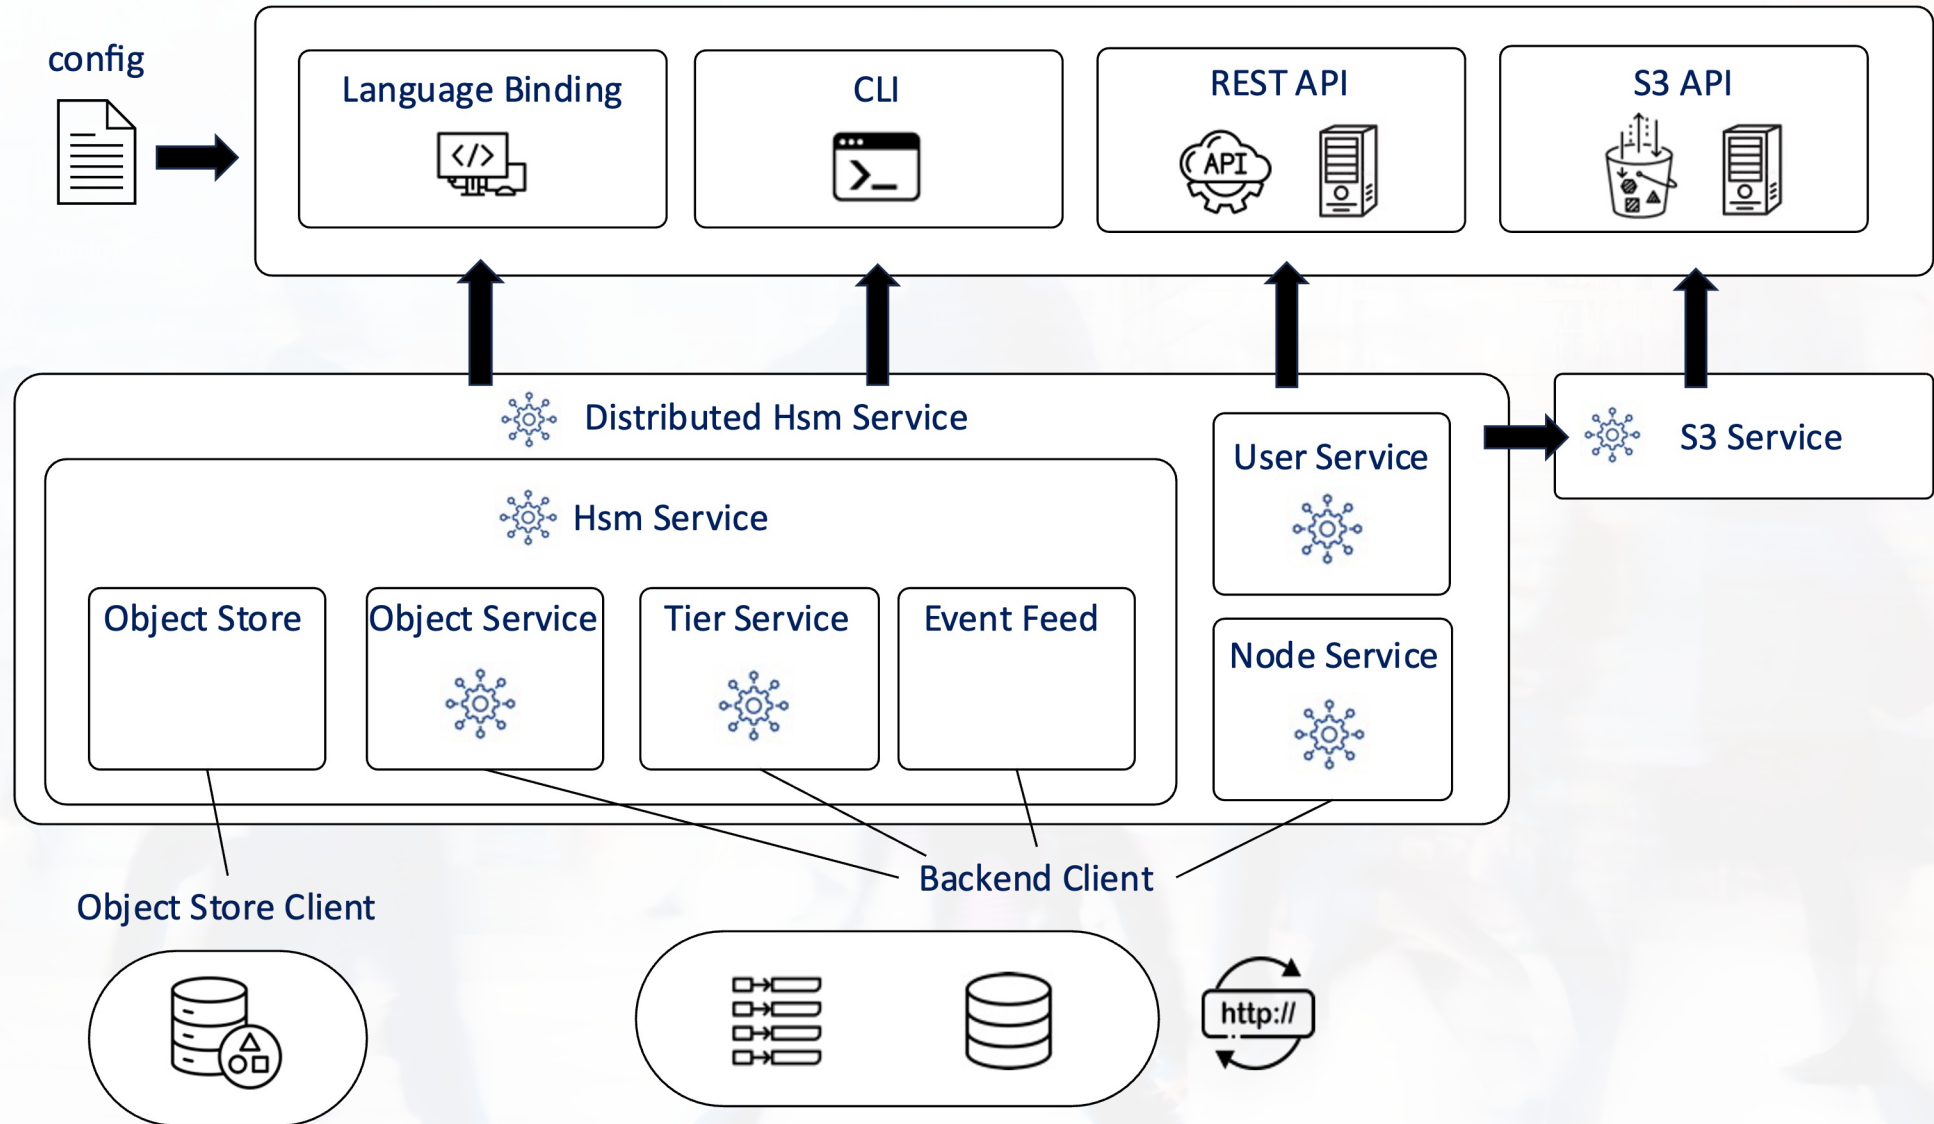


Open Source – MIT License

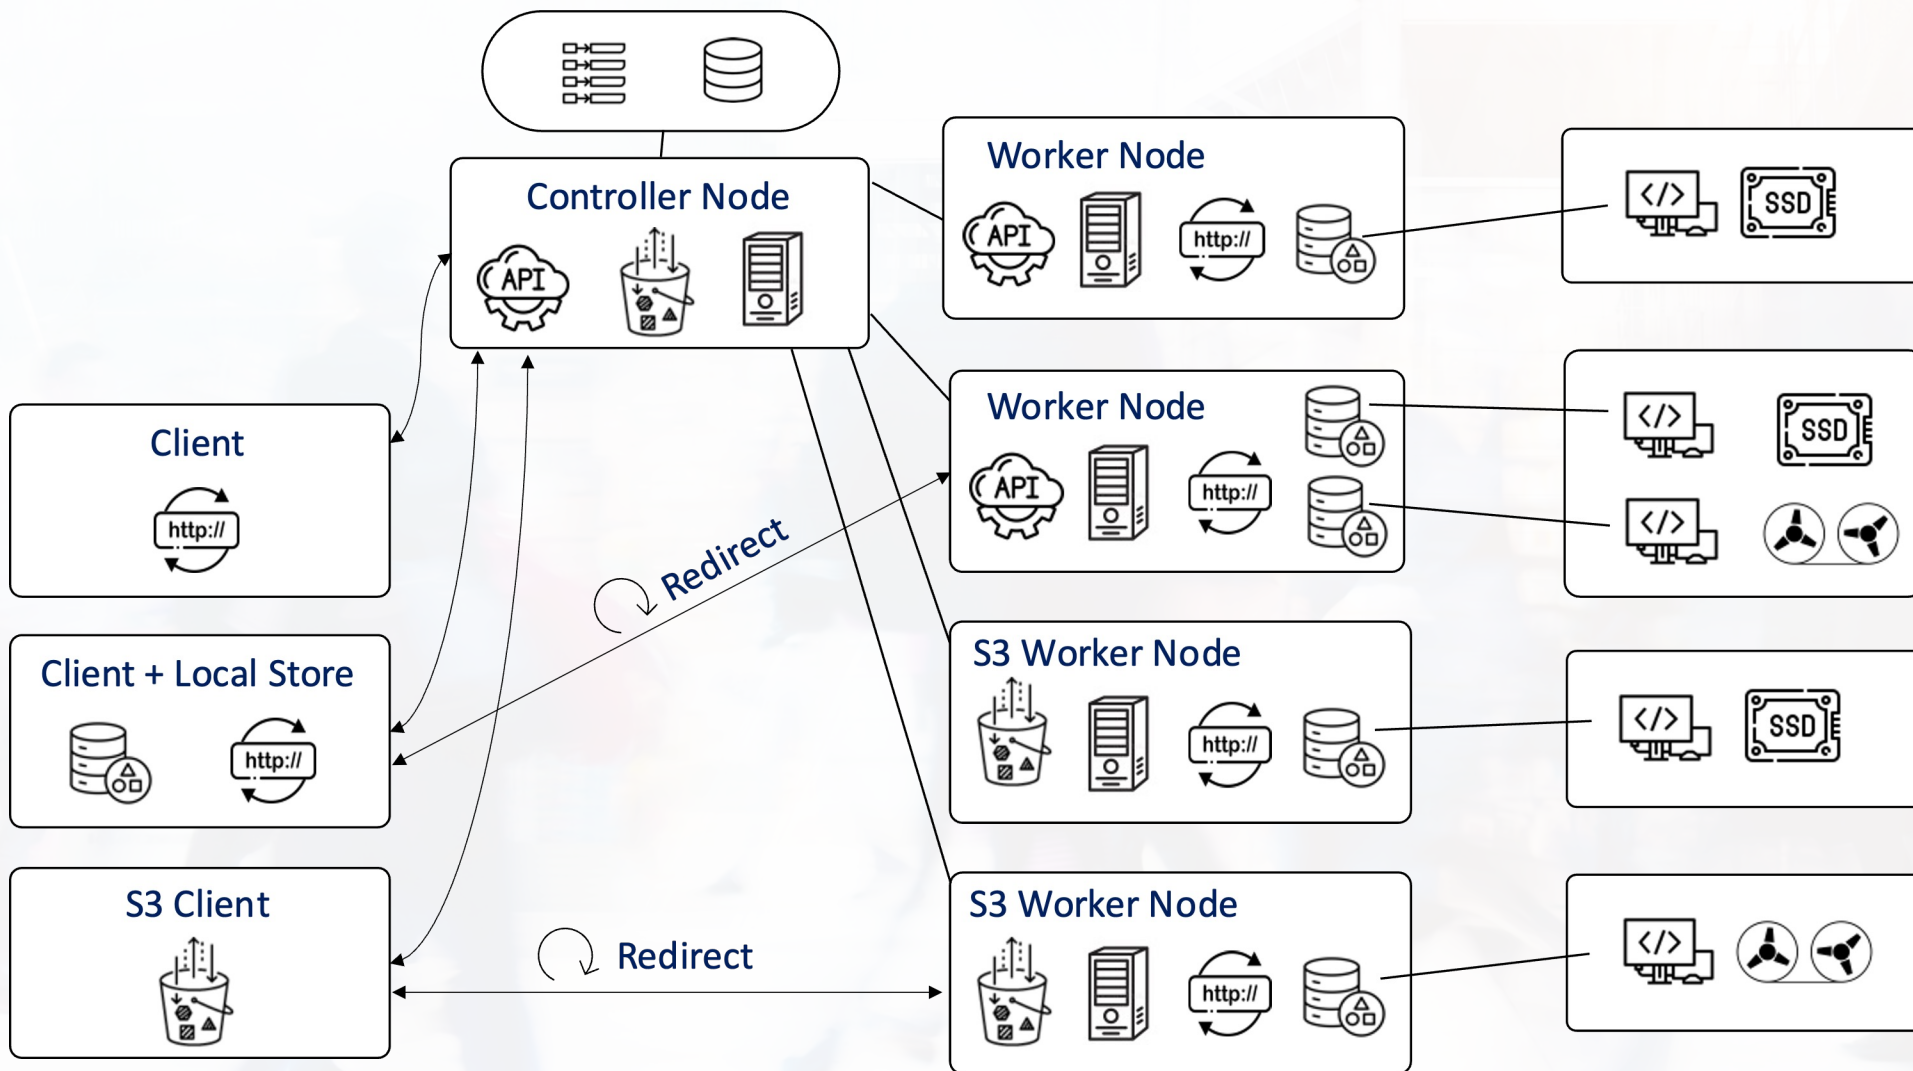


<https://git.ichec.ie/io-sea-internal/hestia>

Hestia Components



Hestia System





Applications

**Unlock
the future**




Applications – Data Mover

Lustre to Remote Tape



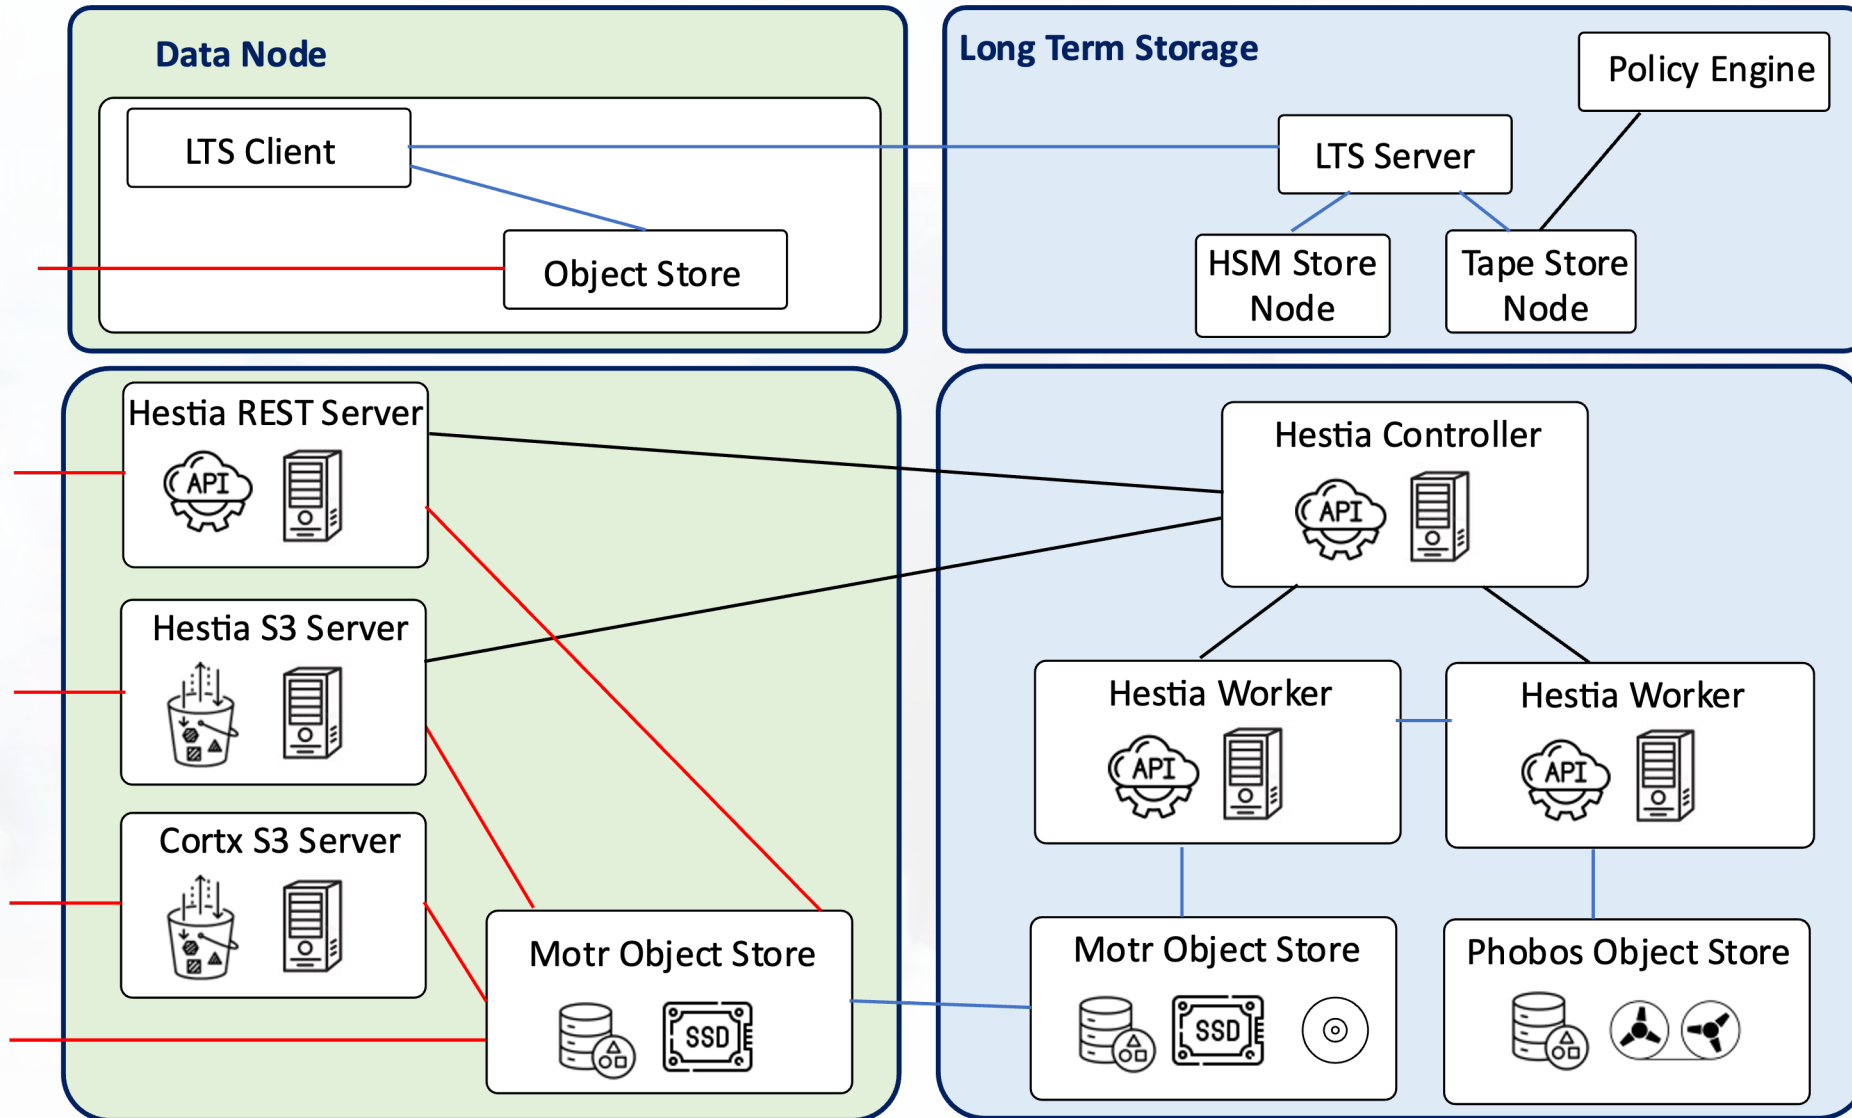
Lustre to Local HSM Object Store



<https://git.ichec.ie/performance/storage/estuary>

<https://github.com/ComputeCanada/lustre-obj-copytool>

Applications – IO-SEA



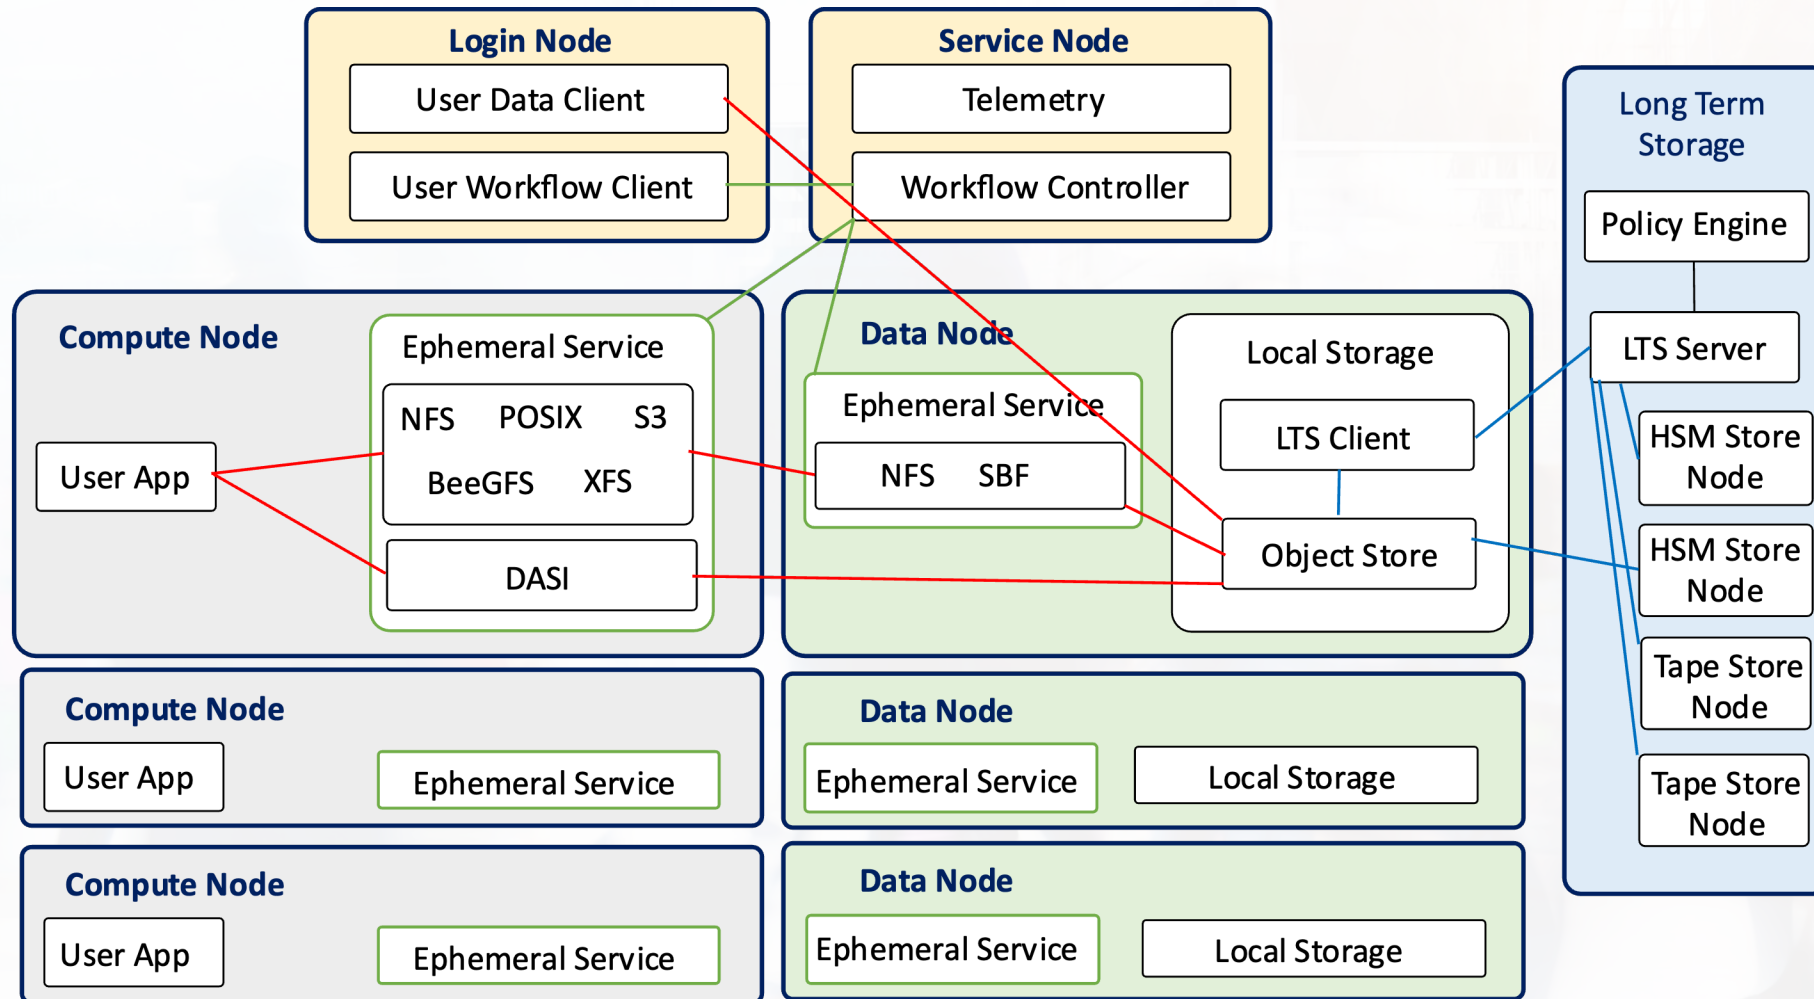


Roadmap

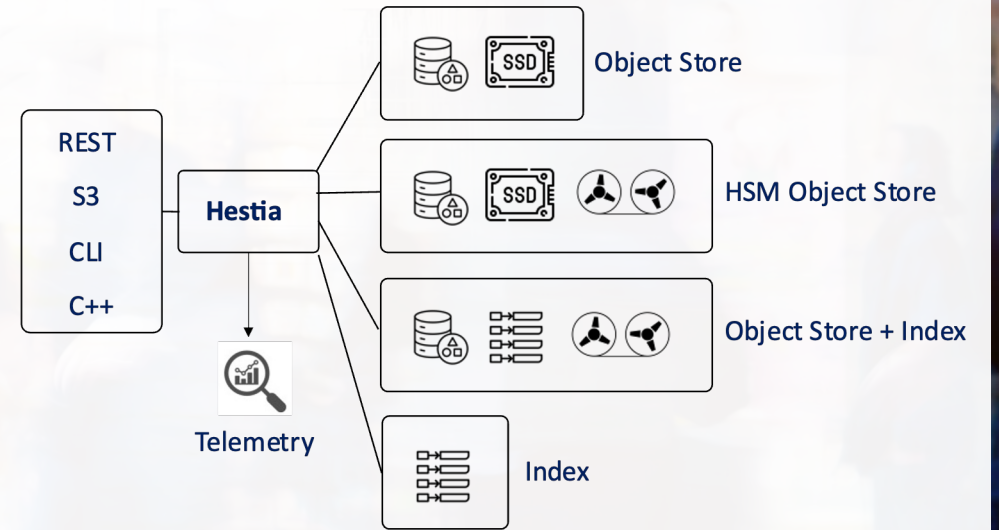

**Unlock
the future**



Roadmap – IO-SEA



Roadmap - Hestia



Resources

Software

IO-SEA: <https://github.com/io-sea>

Hestia: <https://git.ichec.ie/io-sea-internal/hestia>

DASI: <https://github.com/ecmwf-projects/dasi>

Motr: <https://github.com/Seagate/cortx-motr>

Phobos: <https://github.com/cea-hpc/phobos>

Contact

ICHEC: <https://www.ichec.ie> – james.grogan@ichec.ie

IO-SEA: <https://iosea-project.eu/contact/>