Towards Total Knowledge of I/O at NERSC through Holistic Monitoring

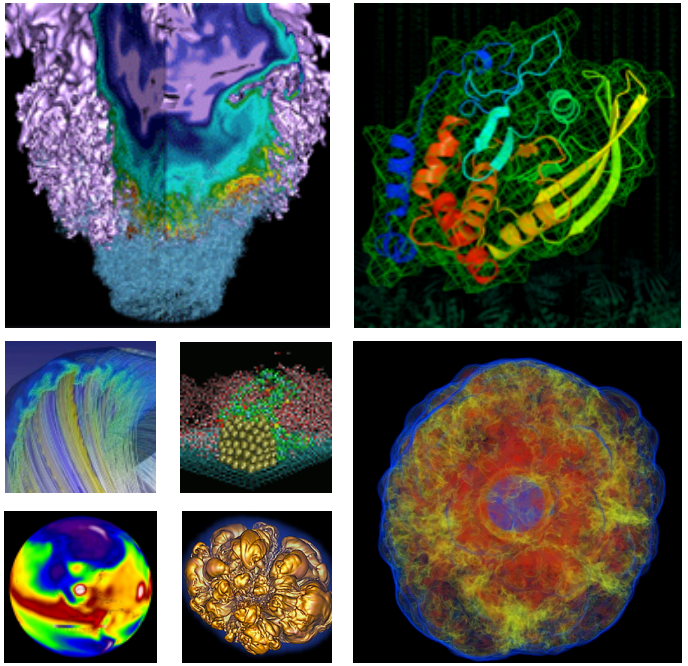

Glenn K. Lockwood, Ph.D.
Advanced Technologies Group

November 16, 2017

Total Knowledge of I/O with holistic analysis

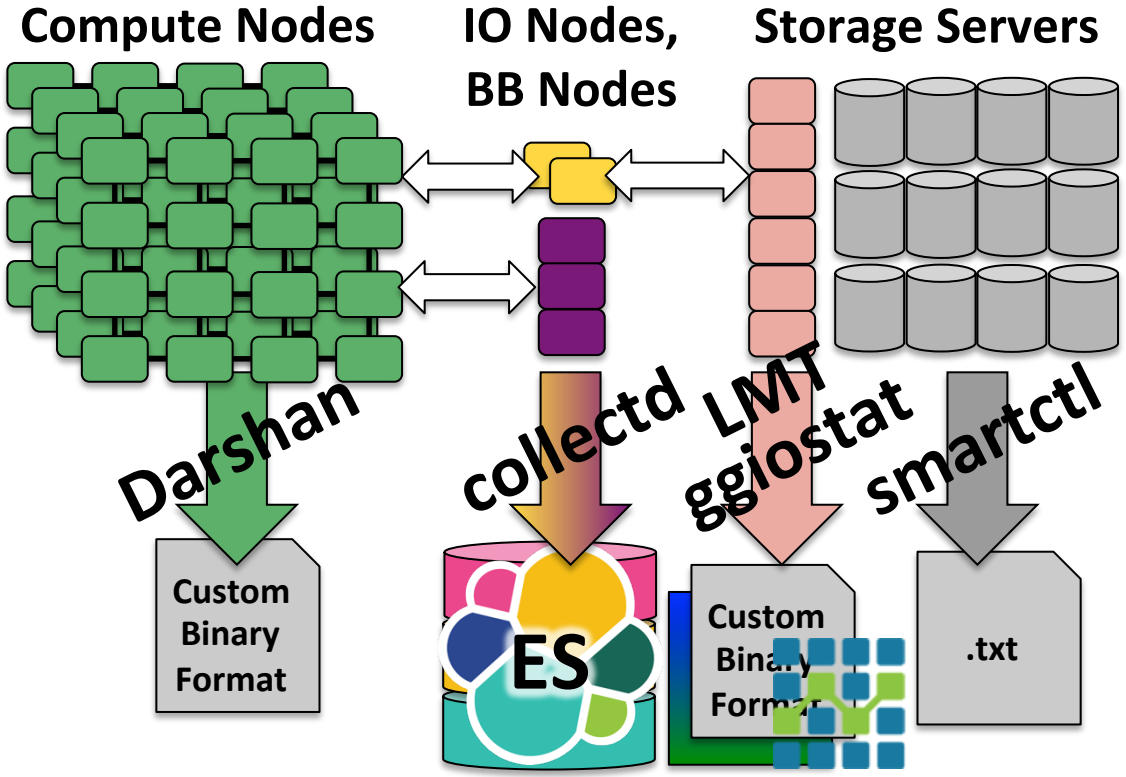

- Our vision: augment expert knowledge using existing tools
- Index and normalize all available data
- Provide a holistic view through a single pane (UMAMI)

Total Knowledge of I/O with holistic analysis

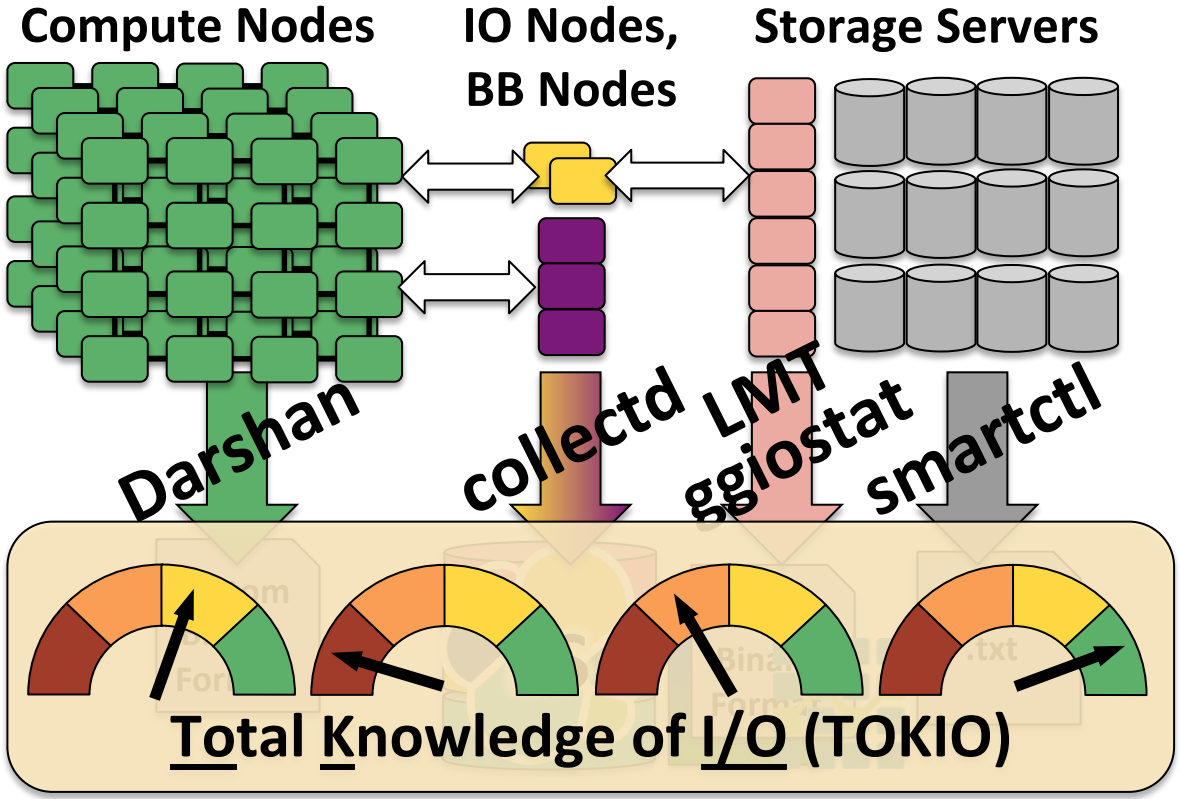

- Our vision: augment expert knowledge using existing tools
- Index and normalize all available data
- Provide a holistic view through a single pane (UMAMI)

UMAMI: variation due to contending bandwidth

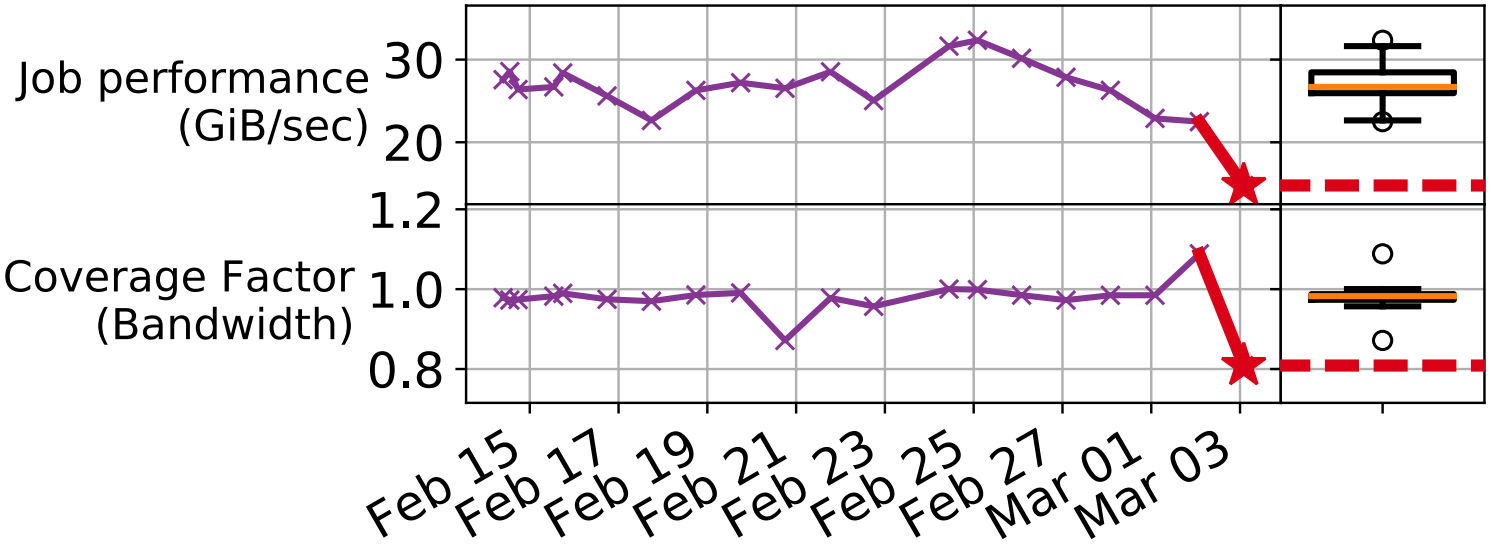

Group job performance by I/O motif (similar transaction size, file/process ratio, client count, etc)

UMAMI: variation due to contending bandwidth

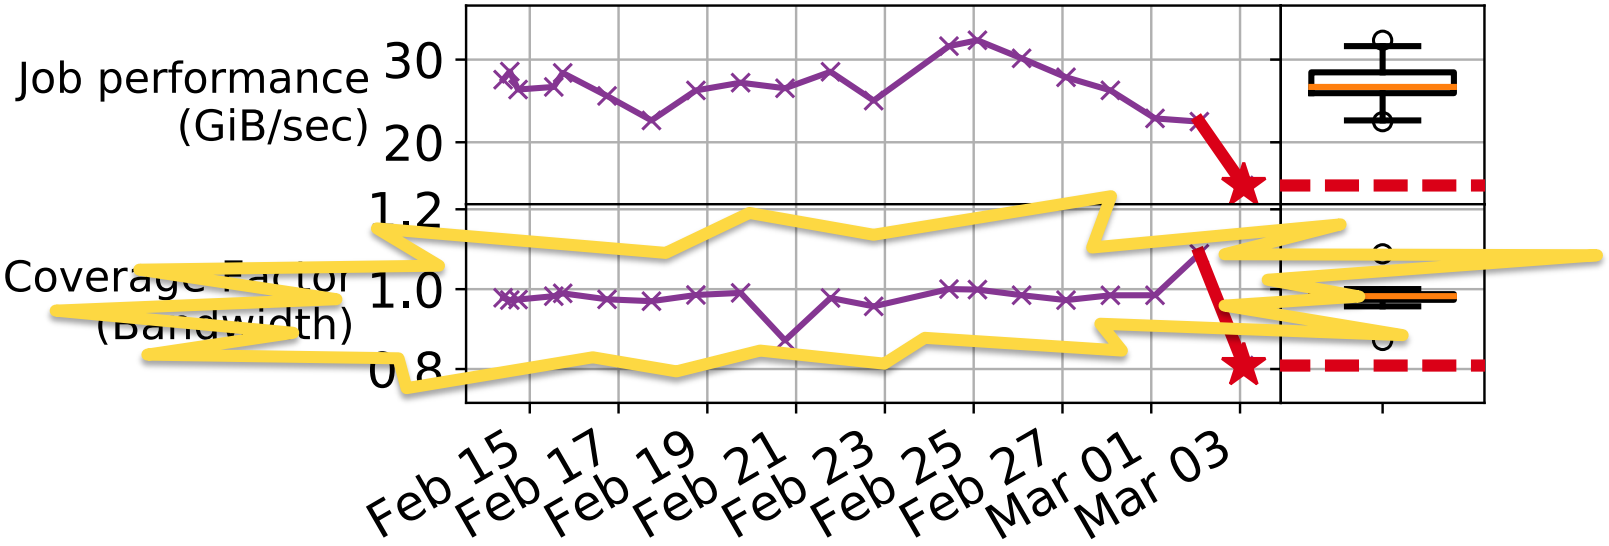

Most jobs get exclusive access to Lustre bandwidth ($CF_{bw} \approx 1.0$)

UMAMI: variation due to contending bandwidth

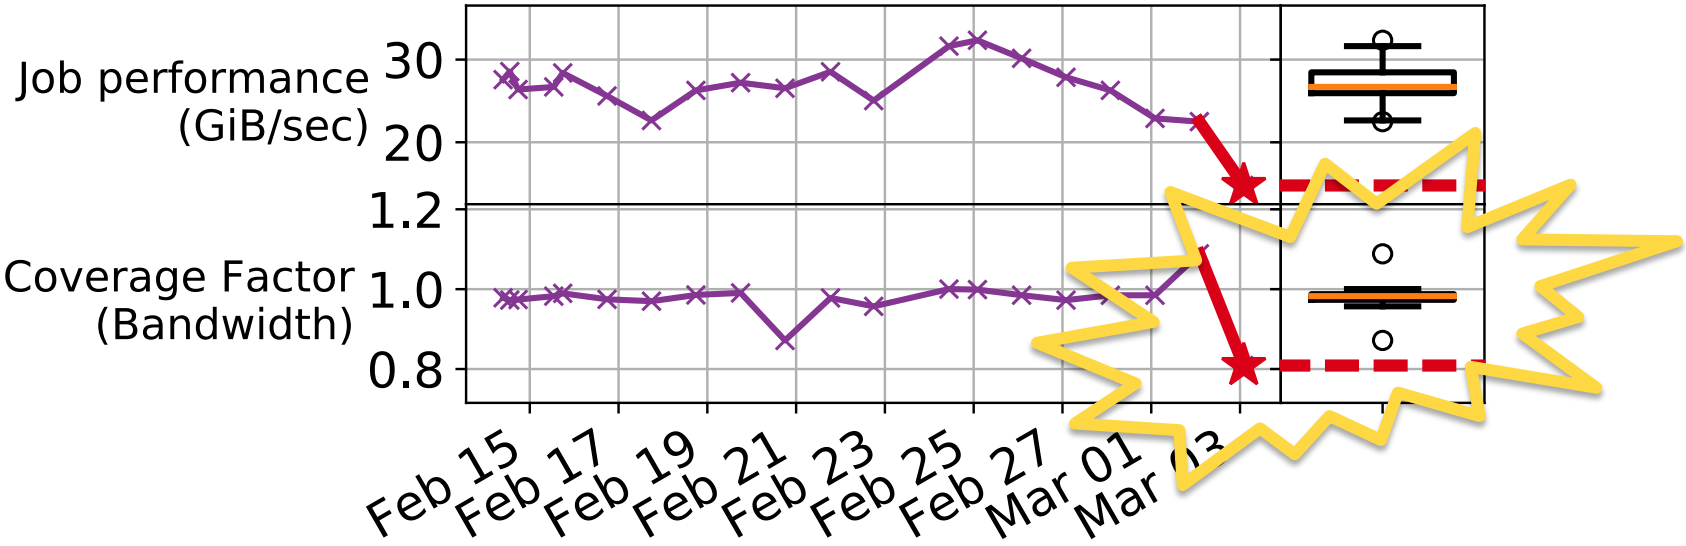

Performance variation caused by bandwidth contention

Variation due to metadata contention (GPFS)

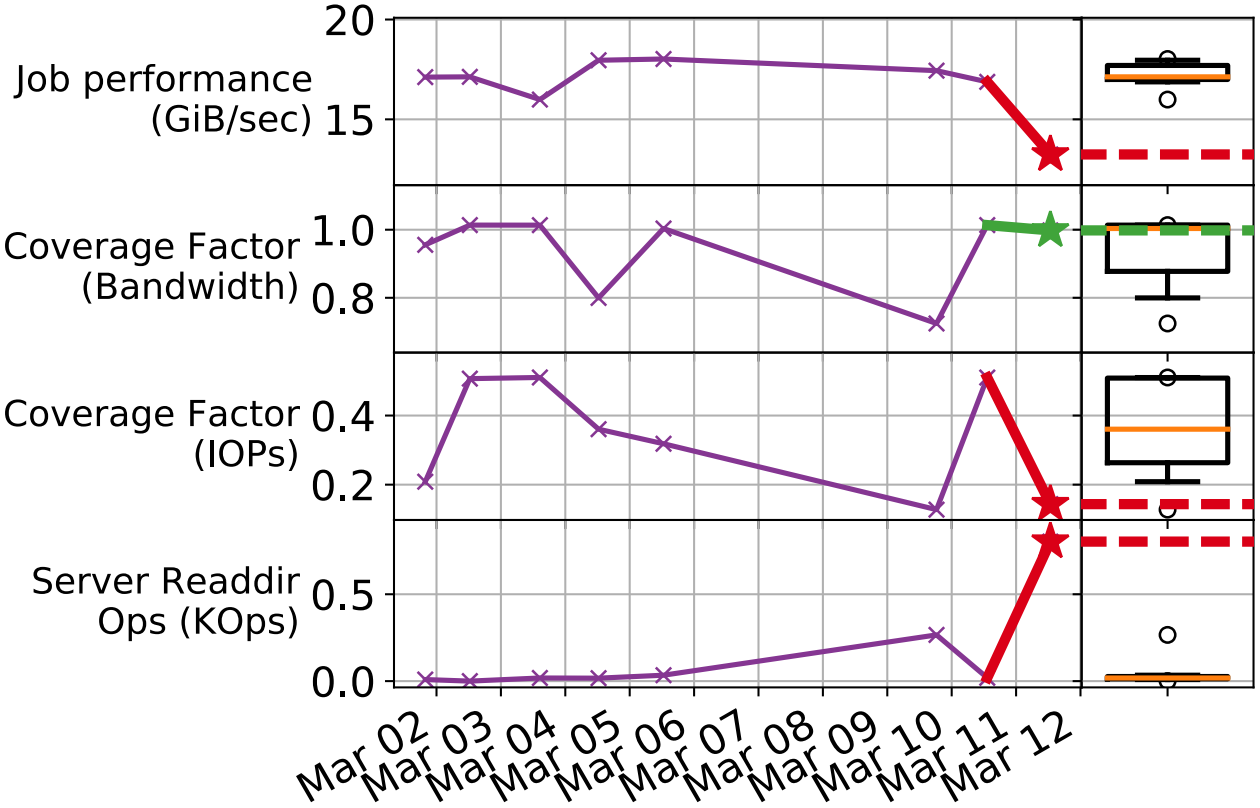

- Bandwidth was uncontended
- IOPS contended by high readdir rates
- Effected by ALCF's GPFS architecture

Variation due to extremely full file system (Lustre)

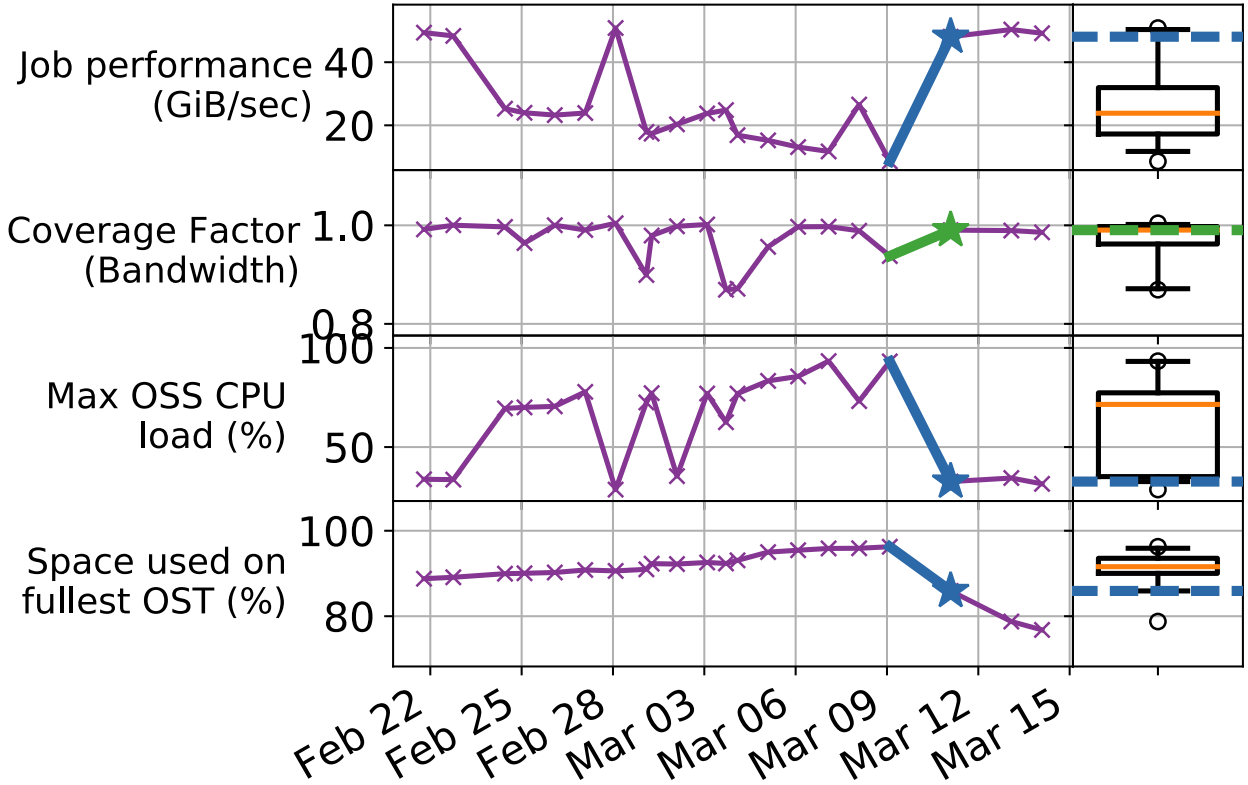

- Moderate negative correlation: Perf vs. OSS CPU load
- Strong negative correlation: fs fullness
- Result of Lustre block allocation at >90% fullness

TOKIO Project—come join the party!

- **Implemented in the pytokio Python package: https://github.com/nersc/pytokio/**
  - Jupyter notebooks: demonstrate useful analyses
  - CLI tools: interact with component-level data
  - Unit tests (and integration tests, smoke tests, etc): basic usage examples and sample input data sets
- **1000% open-source (BSD)**
  - pytokio is open source *and* open development
  - REST API allows researchers to take data with them
- **Supported by DOE SC (DE-AC02-05CH11231 and DE-AC02-06CH11357; *A Framework for Holistic I/O Workload Characterization*; program manager: Dr. Lucy Nowell)**

https://doi.org/10.1145/3149393.3149395

https://github.com/nersc/pytokio