Argonne NATIONAL LABORATORY

Simulation-time data analysis and I/O acceleration at extreme scale with GLEAN

Venkatram Vishwanath, Mark Hereld and Michael E. Papka

Argonne National Laboratory

venkatv@mcs.anl.gov

U.S. DEPARTMENT OF ENERGY

Simulation-time Analysis Opportunities on the Argonne Leadership Computing Facility

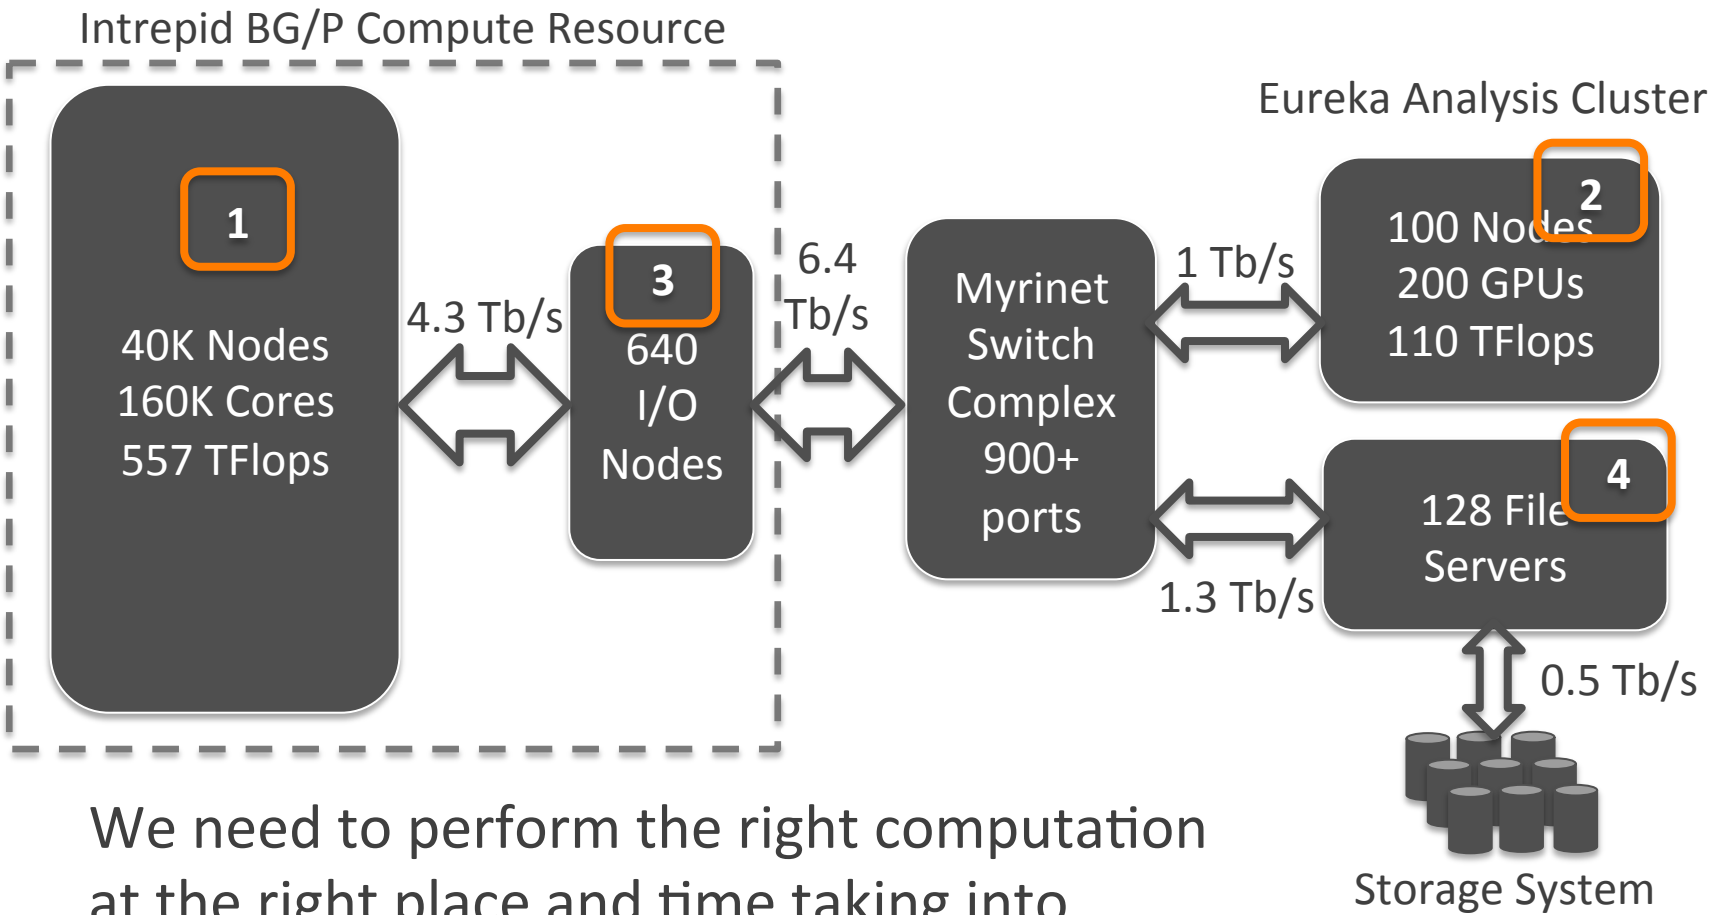

We need to perform the right computation at the right place and time taking into account the characteristics of the simulation, resources and analysis

Our approach - GLEAN

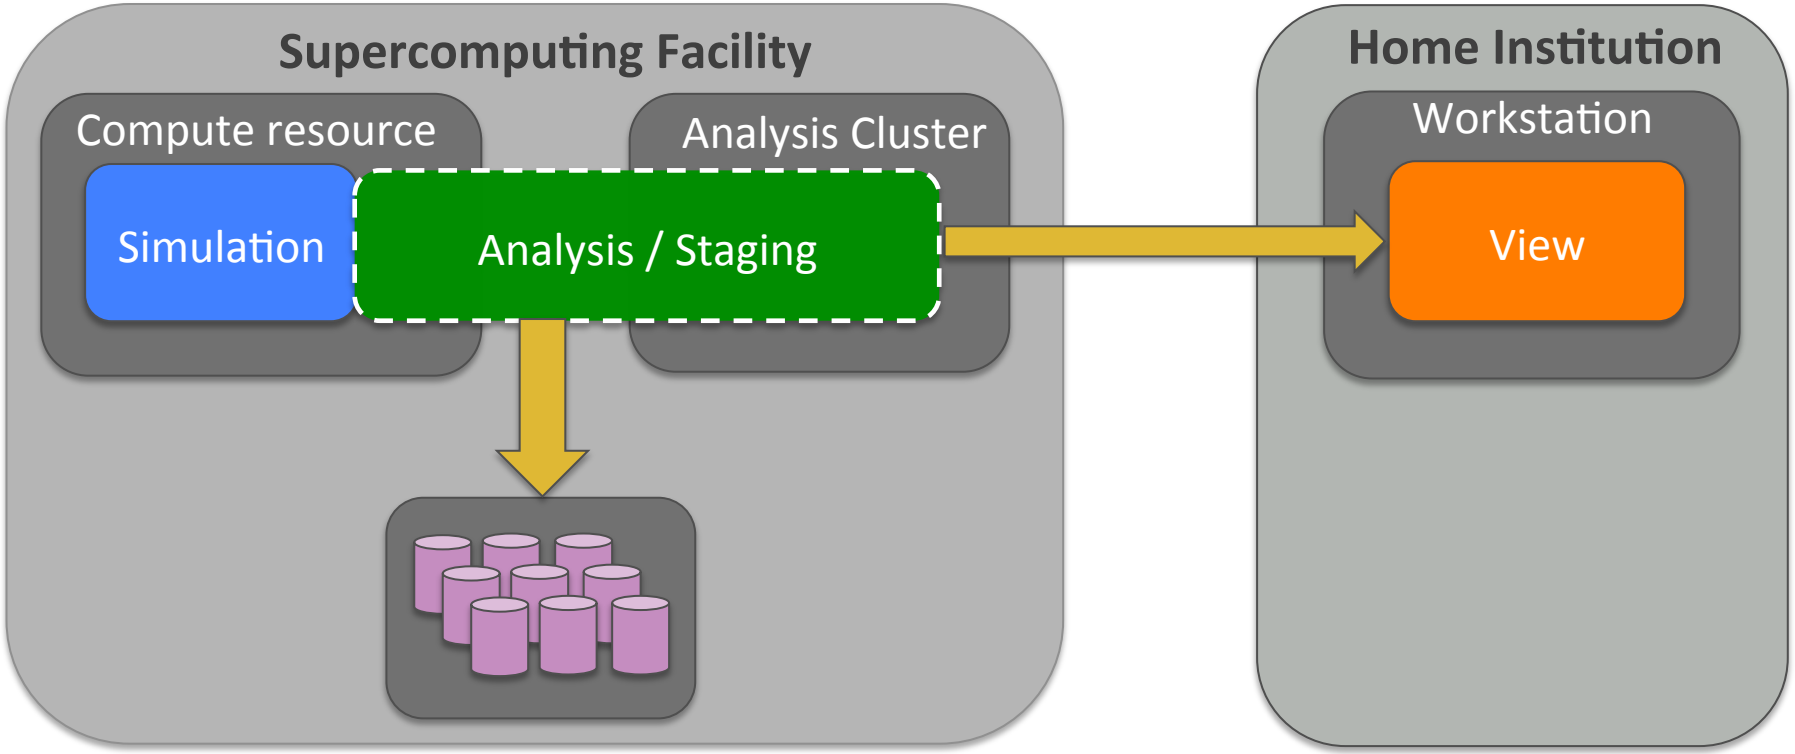

GLEAN is a flexible and extensible framework for simulation-time data analysis and I/O acceleration taking into account application, analytics and system characteristics to perform **the right analysis at the right place and time**.

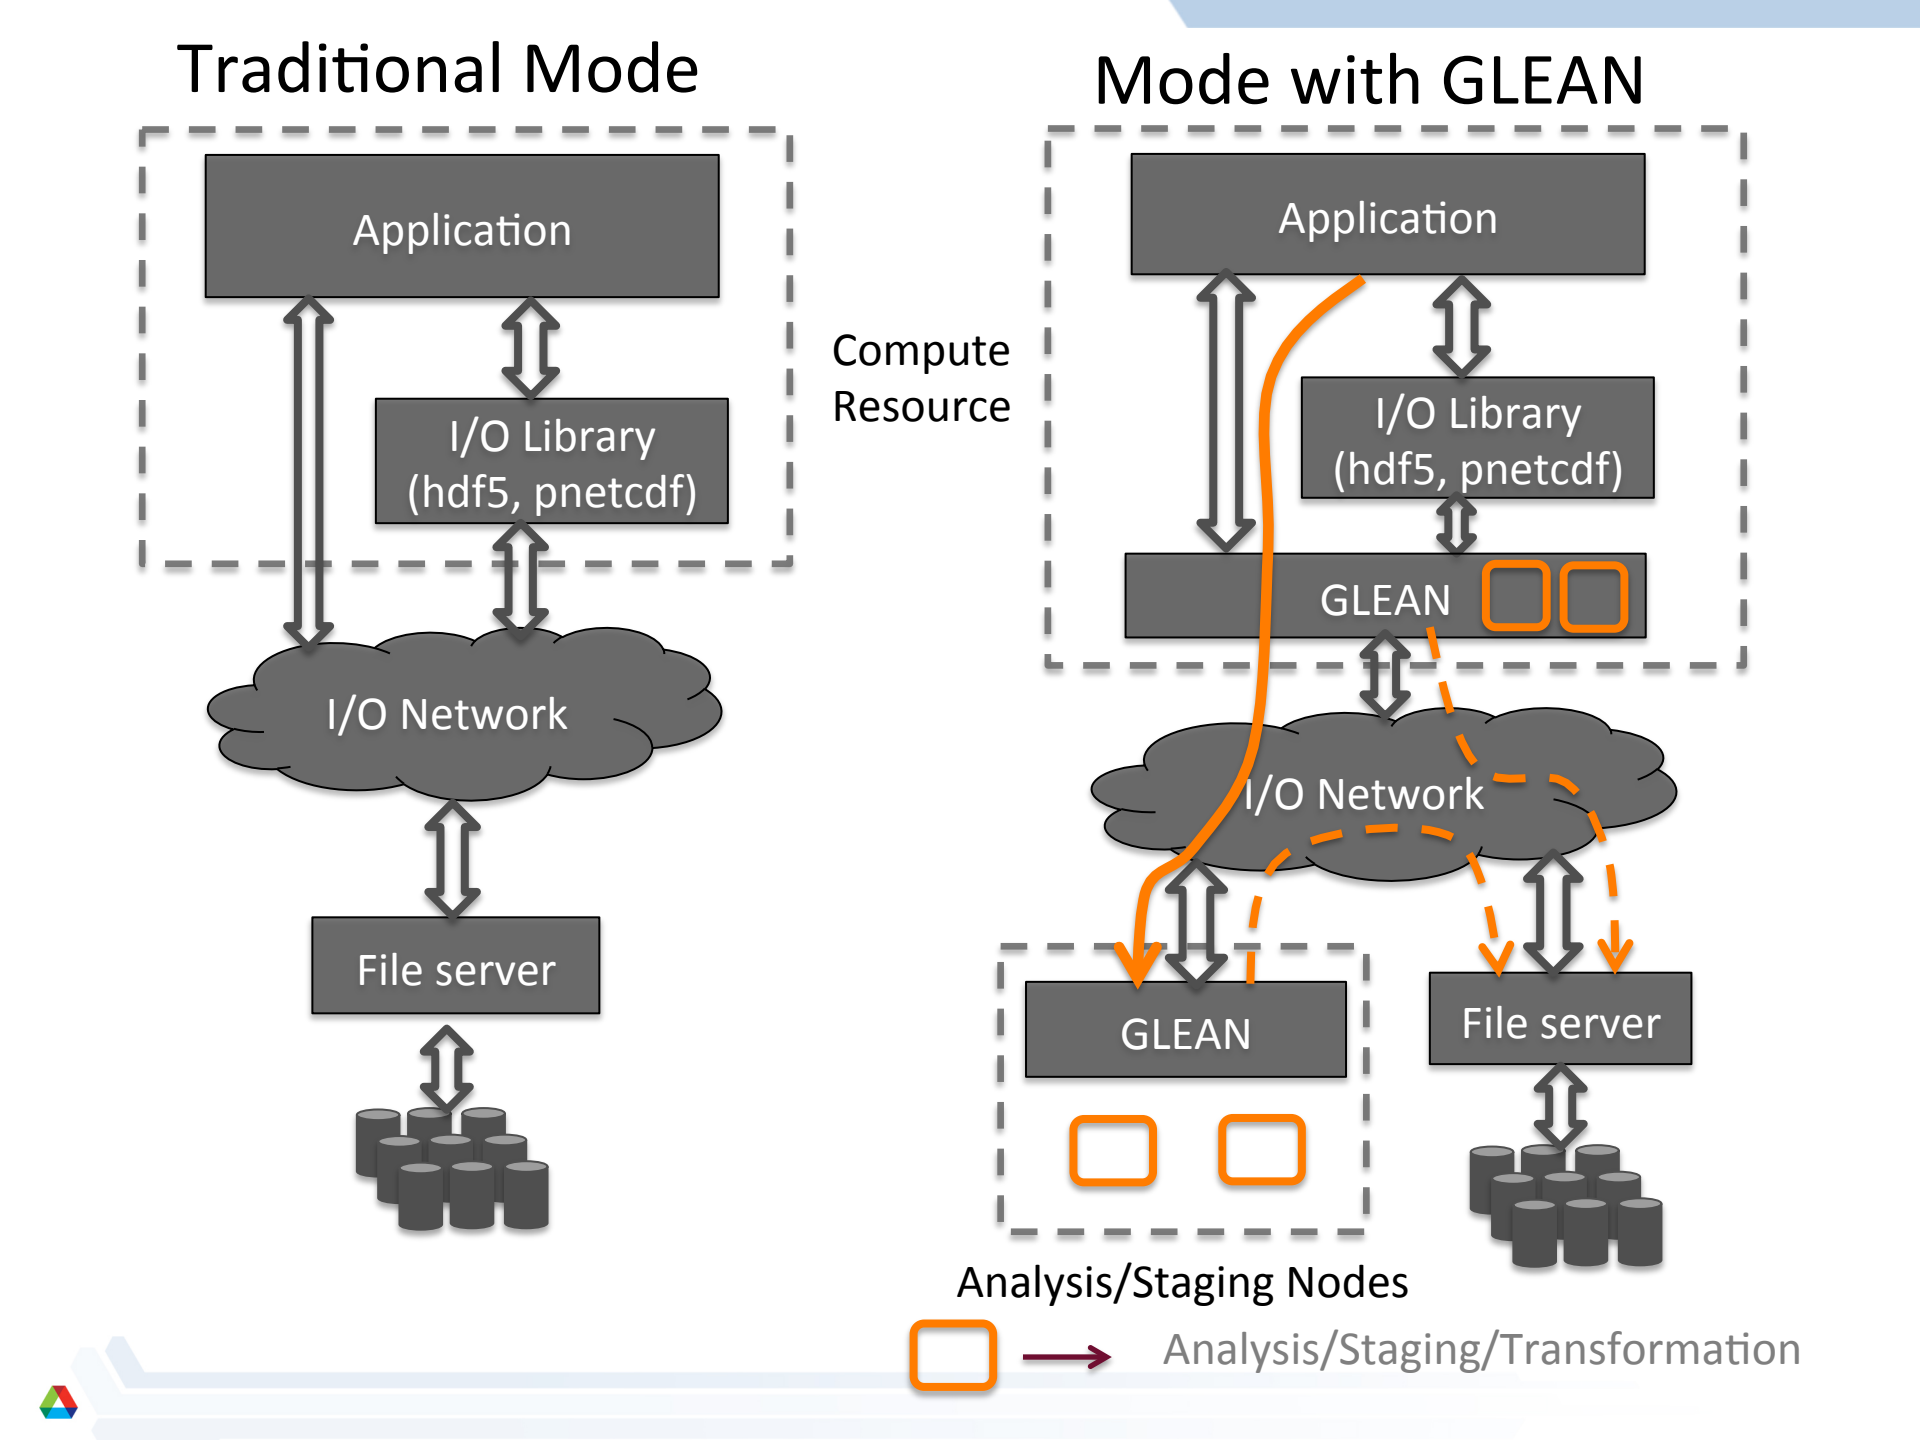
Traditional Mode
Application
I/O Library
(hdf5, pnetcdf)
I/O Network
File server
Compute
Resource
Mode with GLEAN
Application
I/O Library
(hdf5, pnetcdf)
GLEAN
I/O Network
GLEAN
File server
Analysis/Staging Nodes
Analysis/Staging/Transformation

Key features of GLEAN

- Exploit the underlying network topology to speed data movement
- Leverage data semantics of applications
- Provide non-intrusive integration with existing applications
- Enable simulation-time data analysis, transformation and reduction by providing a flexible and extensible API
- Provide asynchronous data I/O via staging nodes
- Provide transparent integration with native application data formats.

Strong scaling performance to write 1GiB

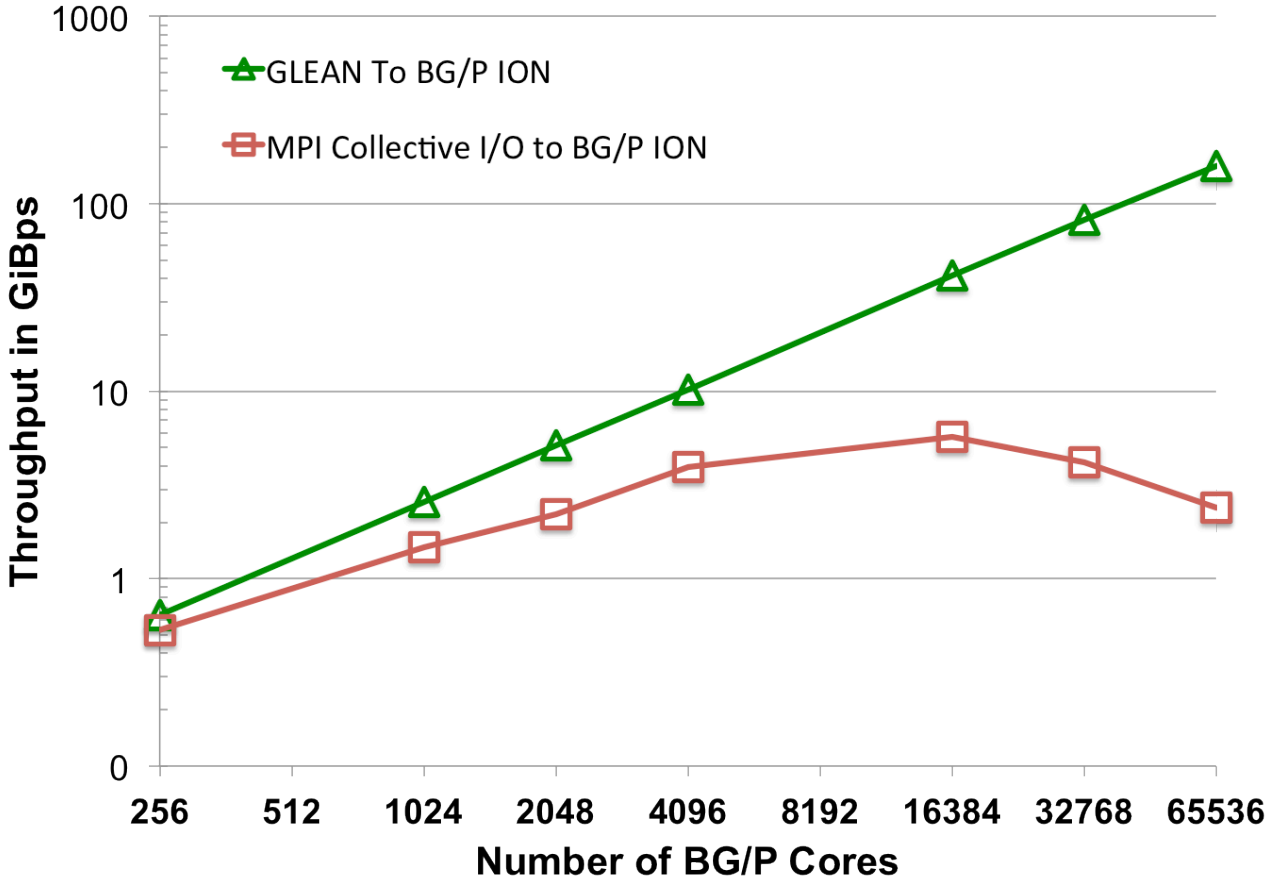

By leveraging the topology of BG/P, we can achieve both weak scaling as well as strong scaling for data movement

"Topology-aware data movement and staging for I/O acceleration for IBM Blue Gene/P supercomputing applications", V. Vishwanath et. al. (To appear SC 2011)

Performance for FLASH checkpoints

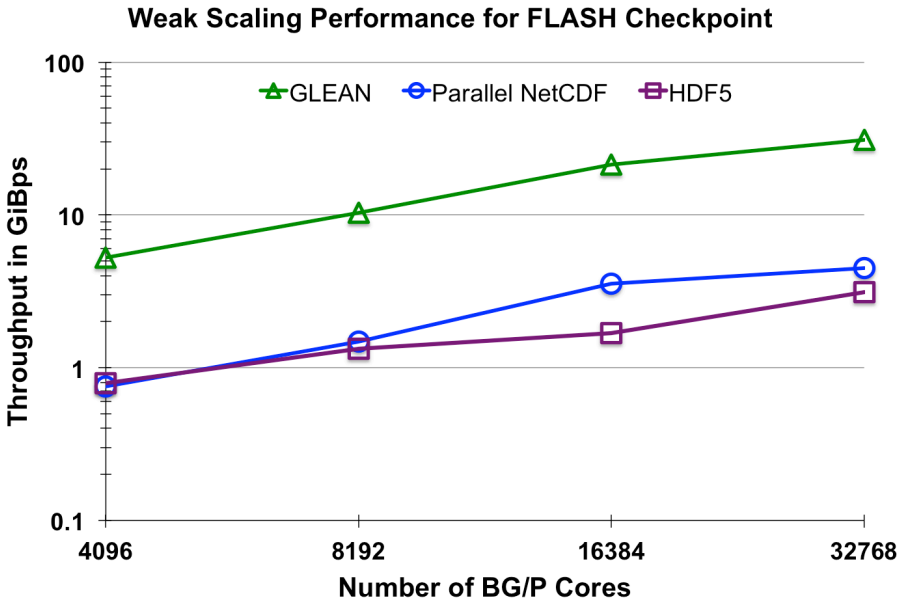

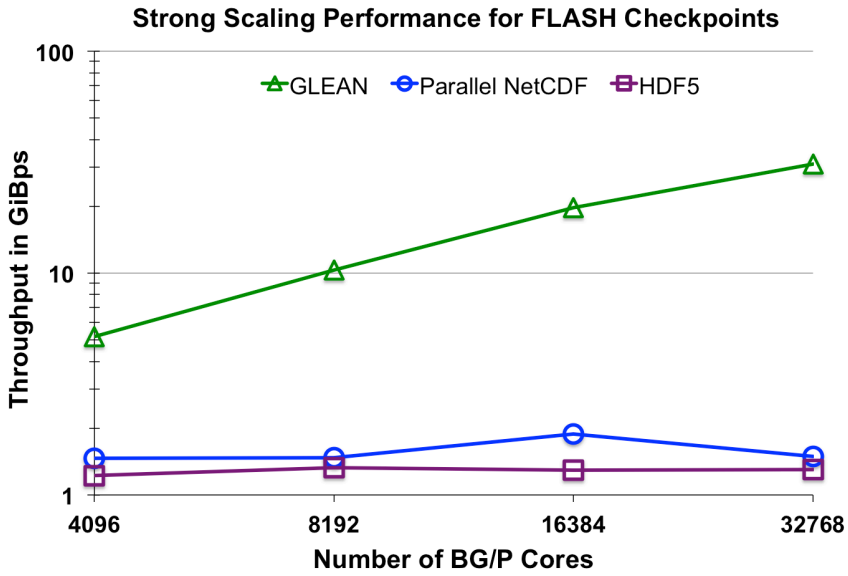

- For weak scaling at 32,768 cores, GLEAN sustains 31 GiBps and achieves an observed speedup of **10-fold** over pnetcdf and hdf5
- For strong scaling at 32,768 cores, GLEAN sustains 27 GiBps and achieves an observed speedup of **15-fold** over pnetcdf and hdf5
- 16.3 GiBps to Storage at 32K cores

in situ analysis of FLASH using GLEAN

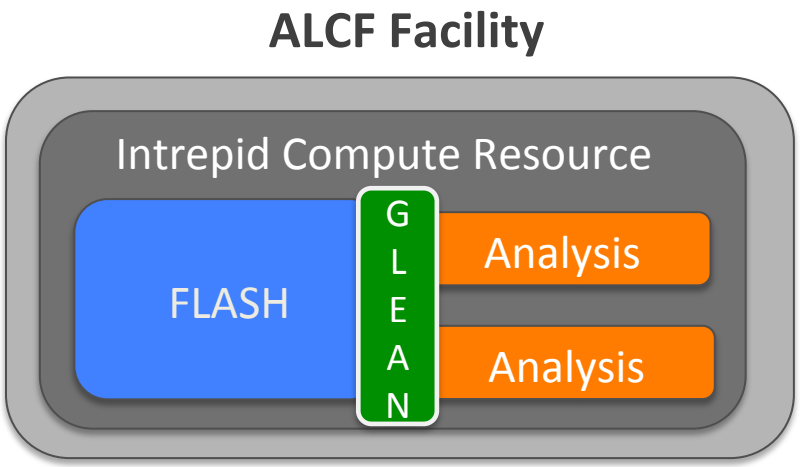

- Fractal Dimension illustrates the degree of turbulence in a particular time step as well as within a sub-region of the domain
- Analysis using GLEAN **required no code changes** to FLASH

- in situ analysis to compute fractal dimension for 5 variables of a FLASH simulation on 2048 BG/P processors

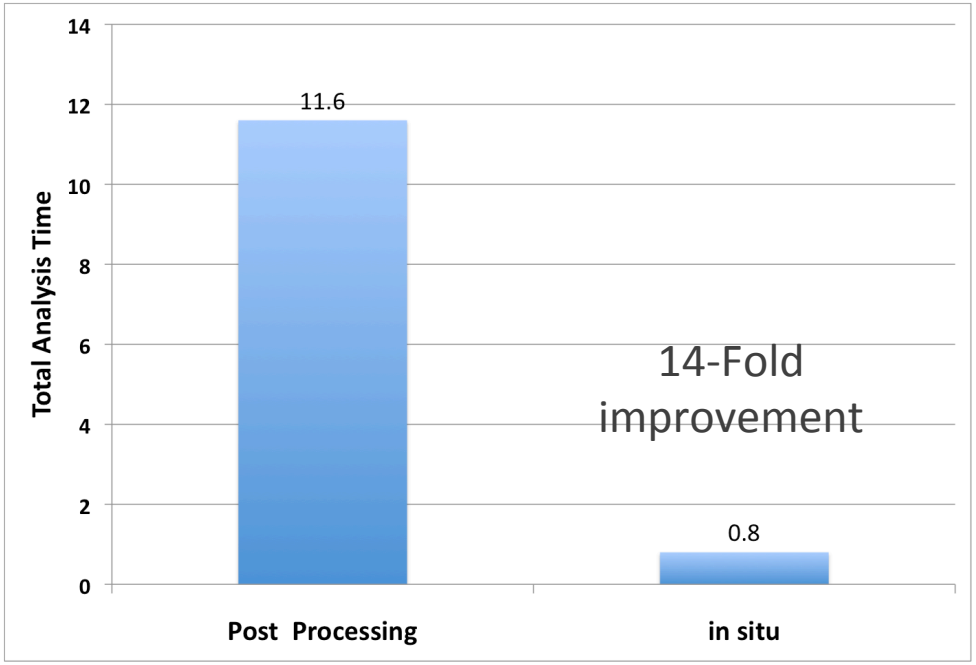

Simulation-time analysis of PHASTA on 160K Intrepid BG/P cores

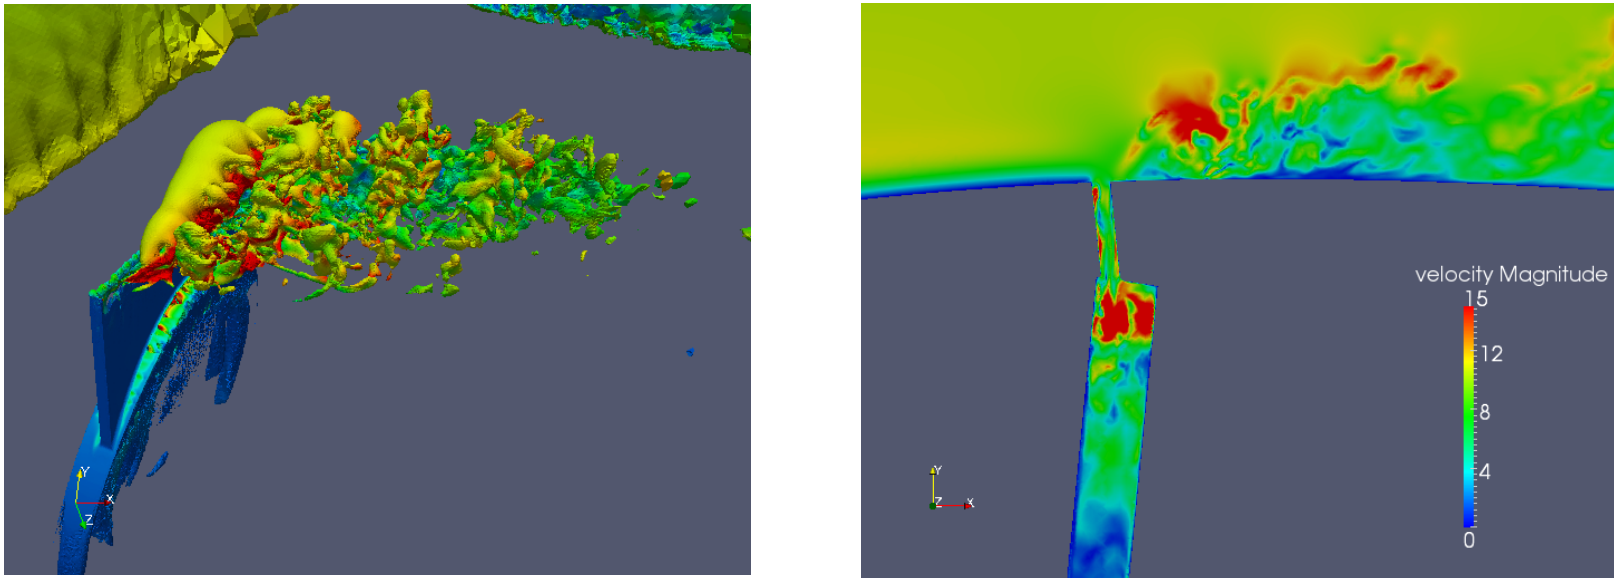

Isosurface of vertical velocity colored by velocity and cut plane through the synthetic jet (both on 3.3 Billion element mesh). *Image Courtesy: Ken Jansen*

- Visualization of a PHASTA simulation running on **160K cores** of Intrepid using ParaView on 100 Eureka nodes **enabled by GLEAN**

- GLEAN achieves **48 GiBps** sustained throughput for data movement enabling simulation-time analysis

GLEAN- Enabling simulation-time data analysis and I/O acceleration

- A **flexible** and **extensible** data analysis framework taking into account application, analytics and system characteristics to perform **the right analysis at the right place and time**
- Provides I/O acceleration by asynchronous data staging
- Scaled to **entire ALCF infrastructure (**160K BG/P cores + 100 Eureka Nodes)

Infrastructure	Simulation	Analysis
Co-analysis	PHASTA	Visualization using Paraview
Staging	FLASH, S3D	I/O Acceleration
In situ	FLASH	Fractal Dimension, Surface Area, Histograms
In flight	MADBench2	Histogram

- Leverages **data models** of applications including adaptive mesh refinement and unstructured meshes
- Generic design to enable deployment on any platform

Acknowledgements

- DOE Office of Advanced Scientific Computing Research
- ANL Director's Fellow Award
- Argonne Leadership Computing (ALCF) Resources
- ANL - Mike Papka, Mark Hereld, Joseph Insley, Eric Olson, Aaron Knoll, Tom Uram, Rob Ross, Tom Peterka, Rob Latham, Phil Carns, Kevin Harms, Kamil Iskra, Vitali Morozov, Susan Coughlan, Ray Loy, and the ALCF team
- FLASH Center – Chris Daley, George Jordan, Anshu Dubey, John Norris, Randy Hudson and Don Lamb
- Kitware - Pat Marion and Berk Geveci
- PHASTA – Ken Jansen, Michel Rasquin and the PHASTA team

venkatv@mcs.anl.gov

Summary

- Exploiting topology, data semantics and asynchronous data staging is critical as we scale to future systems
- GLEAN is a flexible and extensible framework for data analysis and I/O acceleration taking into account application, analytics and system characteristics
- Demonstrated GLEAN successfully with DOE INCITE and ESP applications for simulation-time data analysis and I/O acceleration at scale on leadership computing systems

venkatv@mcs.anl.gov