

End-to-End Online Performance Data Capture and Analysis of Scientific Workflows

G. Papadimitriou, C. Wang, K. Vahi, R. Ferreira da Silva, A. Mandal, Z. Liu, R. Mayani, M. Rynge, M. Kiran, V. Lynch, R. Kettimuthu, E. Deelman, J. Vetter, I. Foster

13th Workflows in Support of Large-Scale Science (WORKS) – Supercomputing 18 November 11, 2018

Funded by the US Department of Energy under Grant #DE-SC0012636M



Panorama 360: Project Overview

- Leverage the Pegasus WMS to structure, execute and monitor workflow execution
 - Characterize performance: instrument data capture, summarize data, and publish results
- Create an open access common repository for storing end-to-end workflow performance and resource data captured using a variety of tools

*Open for external contributors

- Apply and develop ML techniques for workflow performance analysis and infrastructure troubleshooting
- Record findings, distill best practices, and share and refine them with other program teams





Data Sources: Application and Infrastructure

- Pegasus Kickstart wraps the execution of jobs and provides execution statistics (duration, I/O, memory usage)
- **Pegasus Kickstart Online** collects resource usage traces with frequency as low as 1 second in real-time
- Darshan collects file access statistics (eg. POSIX, MPI-IO) during the execution
- Globus collects transfer statistics and general information about the transfer (throughput, file transfer errors etc.)
- **TSTAT** captures low level network statistics



DARSHAN HPC I/O Characterization Tool





Data Sources: Problems

- They are scattered across multiple locations (Eg. execution site, cloud service, pegasus logs)
- They don't contain metadata about the workflow, and it's very hard to locate and match them in the future
- Captured data don't have a common format
 - Pegasus Kickstart logs are in xml format
 - Pegasus Kickstart online logs are in json format
 - Globus logs are in json format
 - Darshan logs are in binary format
 - Tstat logs are in csv format

This is messy !!!





Data Collection: End-to-End Workflow Execution Monitoring

- Pegasus apart from planning and running the workflow, orchestrates the data collection
- A message queueing system is used, to decouple the publishers from the datastore
- Flexible search and visualization engines are used to explore the data







Data Collection: Architecture Overview







Data Collection: Tools enhancements and new tools

- pegasus-monitord: extended with JSON output format, the ability to pickup job related monitoring messages, and publish capability
- pegasus-transfer: extended to support Globus transfers and the ability to publish statistics in json format

pegasus-darshan: wrapper to darshan-parser, that pushes darshan logs in JSON

format to pegasus-monitord

Achieved Common Data Format and Automated Collection !!!





Visualization: Detailed Workflow and Job Characteristics

Workflow Dashboard Panorama 360 Workflow Performance Explorer







Visualization: Time Series Data of Workflow Performance







Repository: Open access data



https://data.panorama.isi.edu

https://kibana.panorama.isi.edu



https://panorama360.github.io



Repository: Organization

ElasticSearch Index	Description
panorama_transfer	Globus logs
panorama_kickstart	Pegasus-Kickstart online traces
panorama_stampede	Workflow Events and Darshan logs





Repository: Paving the way for in depth analysis

- End to end characterization of scientific workflows
 - Understanding the requirements of future systems
- Online workflow analysis
 - Network behavior
 - Filesystem access patterns
 - Compute resource usage
- Training ML models on historical data and capturing anomalies as they happen







 GitHub: <u>https://github.com/Panorama360</u>

 Website: <u>https://panorama360.github.io</u>



George Papadimitriou

Computer Science PhD Student University of Southern California

email: georgpap@isi.edu



School of Engineering Department of Computer Science

https://panorama360.github.io/