Pegasus Users Group

MEETING

End to End Workflow Monitoring

George Papadimitriou
USC PhD Student in Computer Science
ISI Graduate Research Assistant

2/25/2021

U.S. DEPARTMENT OF ENERGY | Office of Science

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

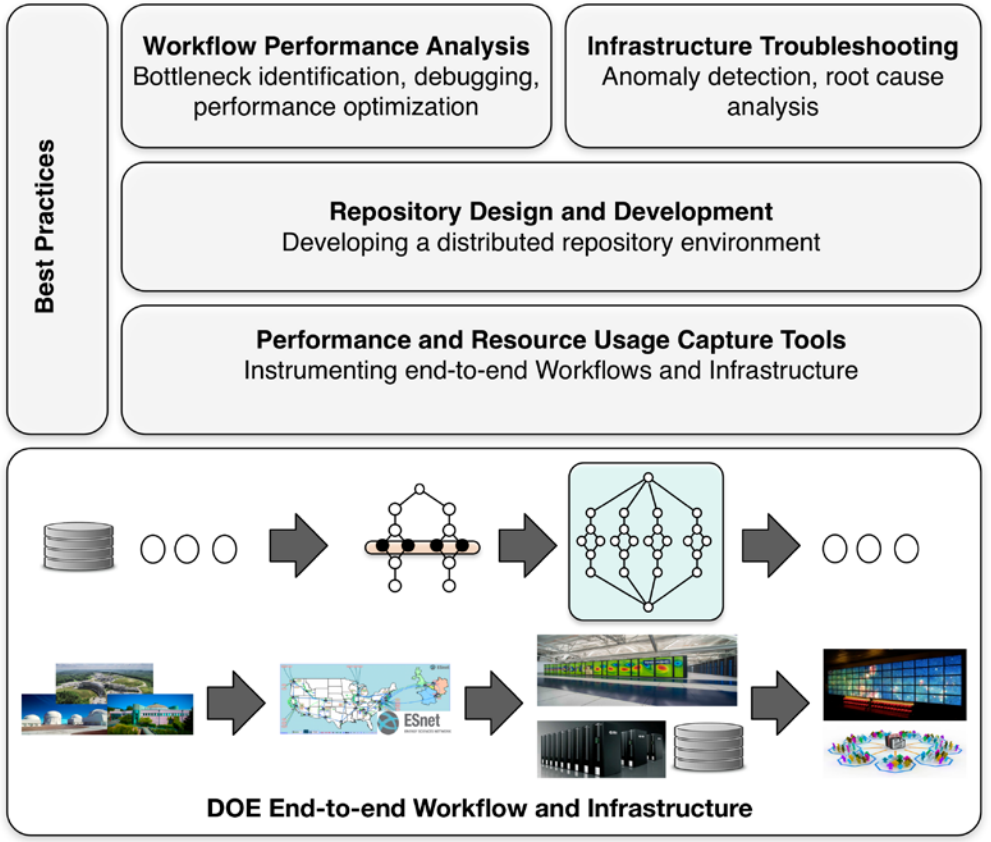

Panorama 360 Overview.

Data Sources: Application and Infrastructure

- **Pegasus Stampede** events regarding the workflow and its status
- **Pegasus Kickstart Online** collects resource usage traces with frequency as low as 1 second in real-time
- **Pegasus Transfer** collects transfer statistics and general information about the transfer (throughput, file transfer errors etc.) and can query transfer services (e.g., Globus)
- **Darshan** collects file access statistics (eg. POSIX, MPI-IO) during the execution

DARSHAN
HPC I/O Characterization Tool

globus

Data Collection: Architecture Overview

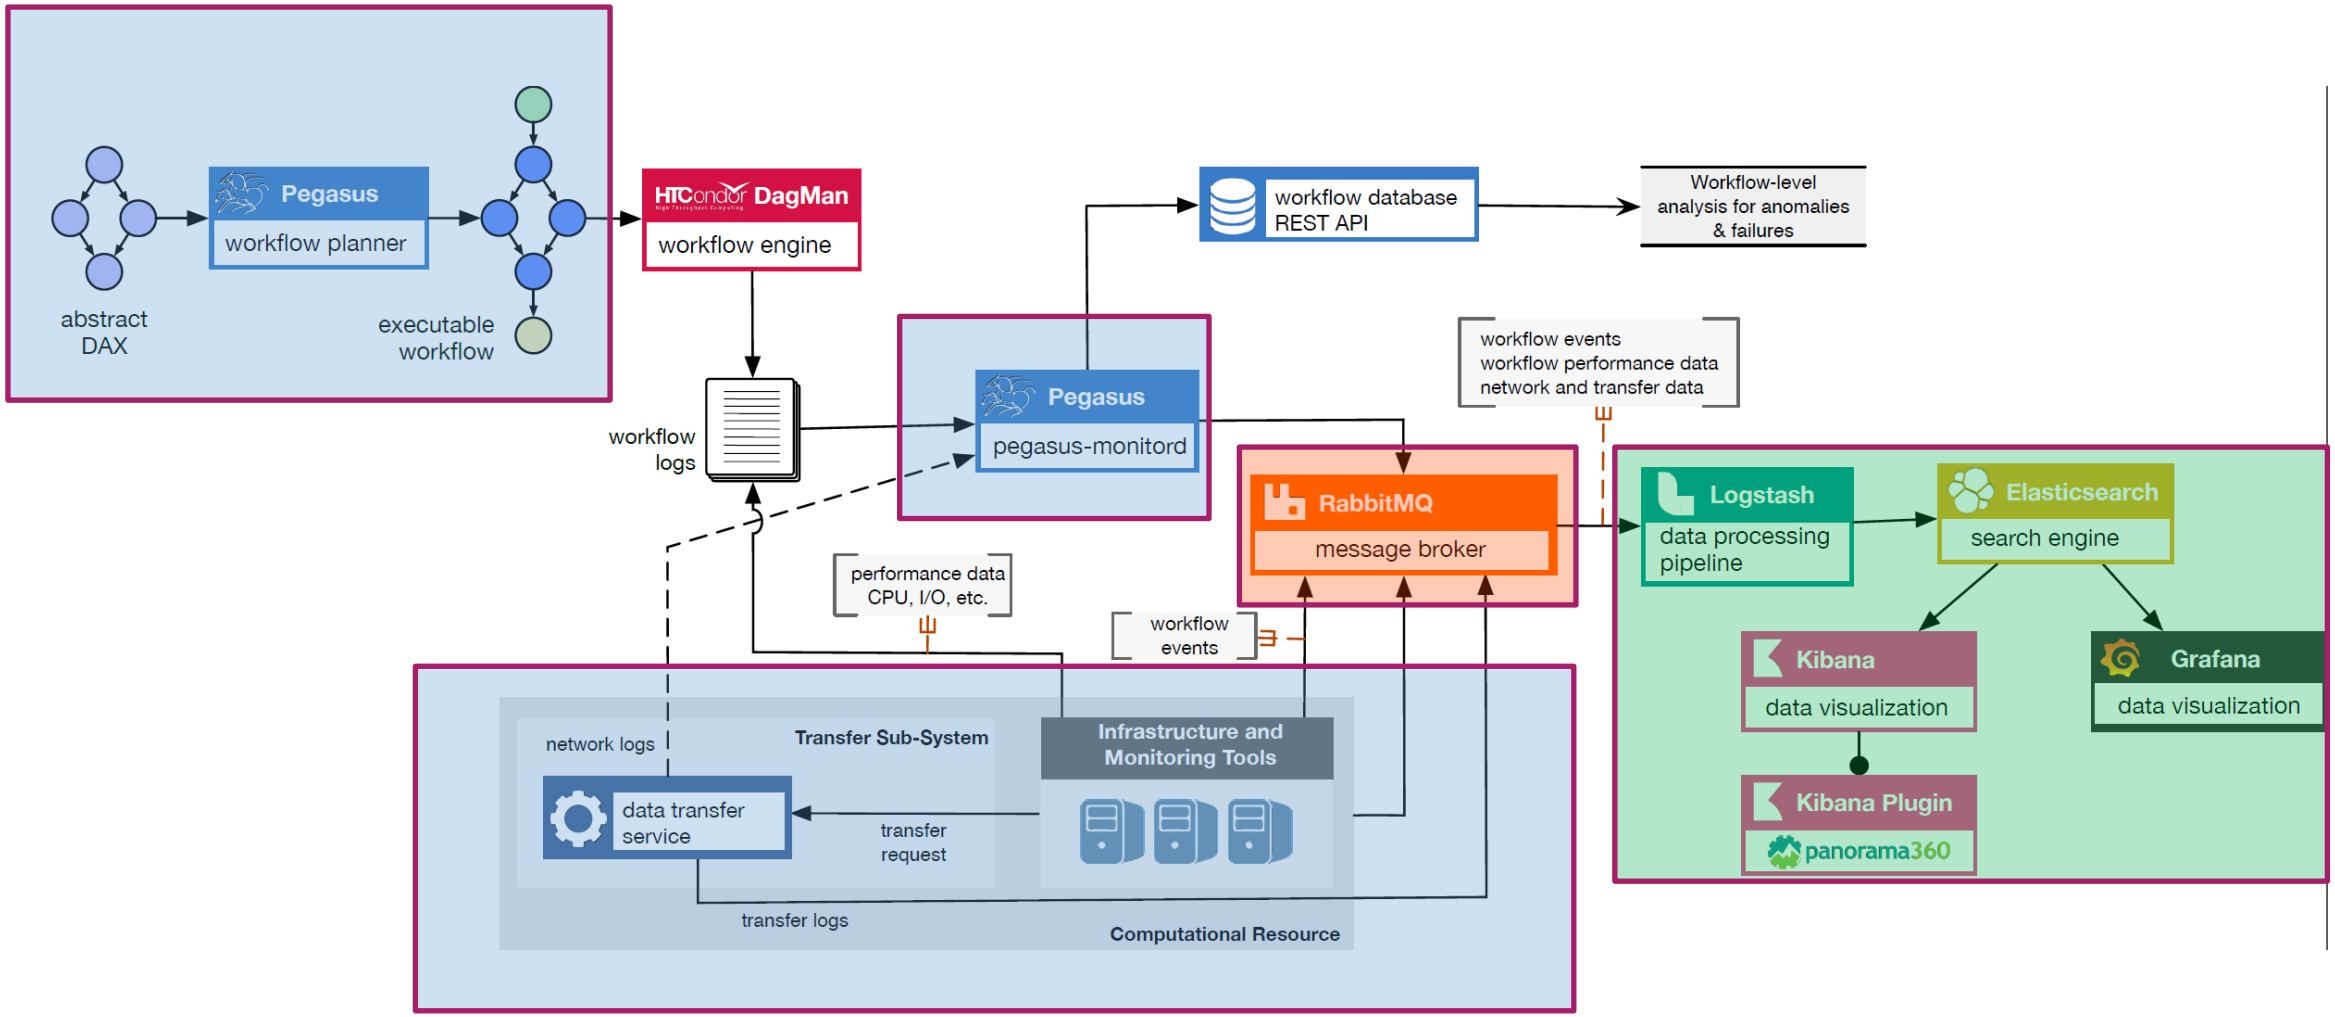

Visualization: Detailed Workflow and Job Characteristics

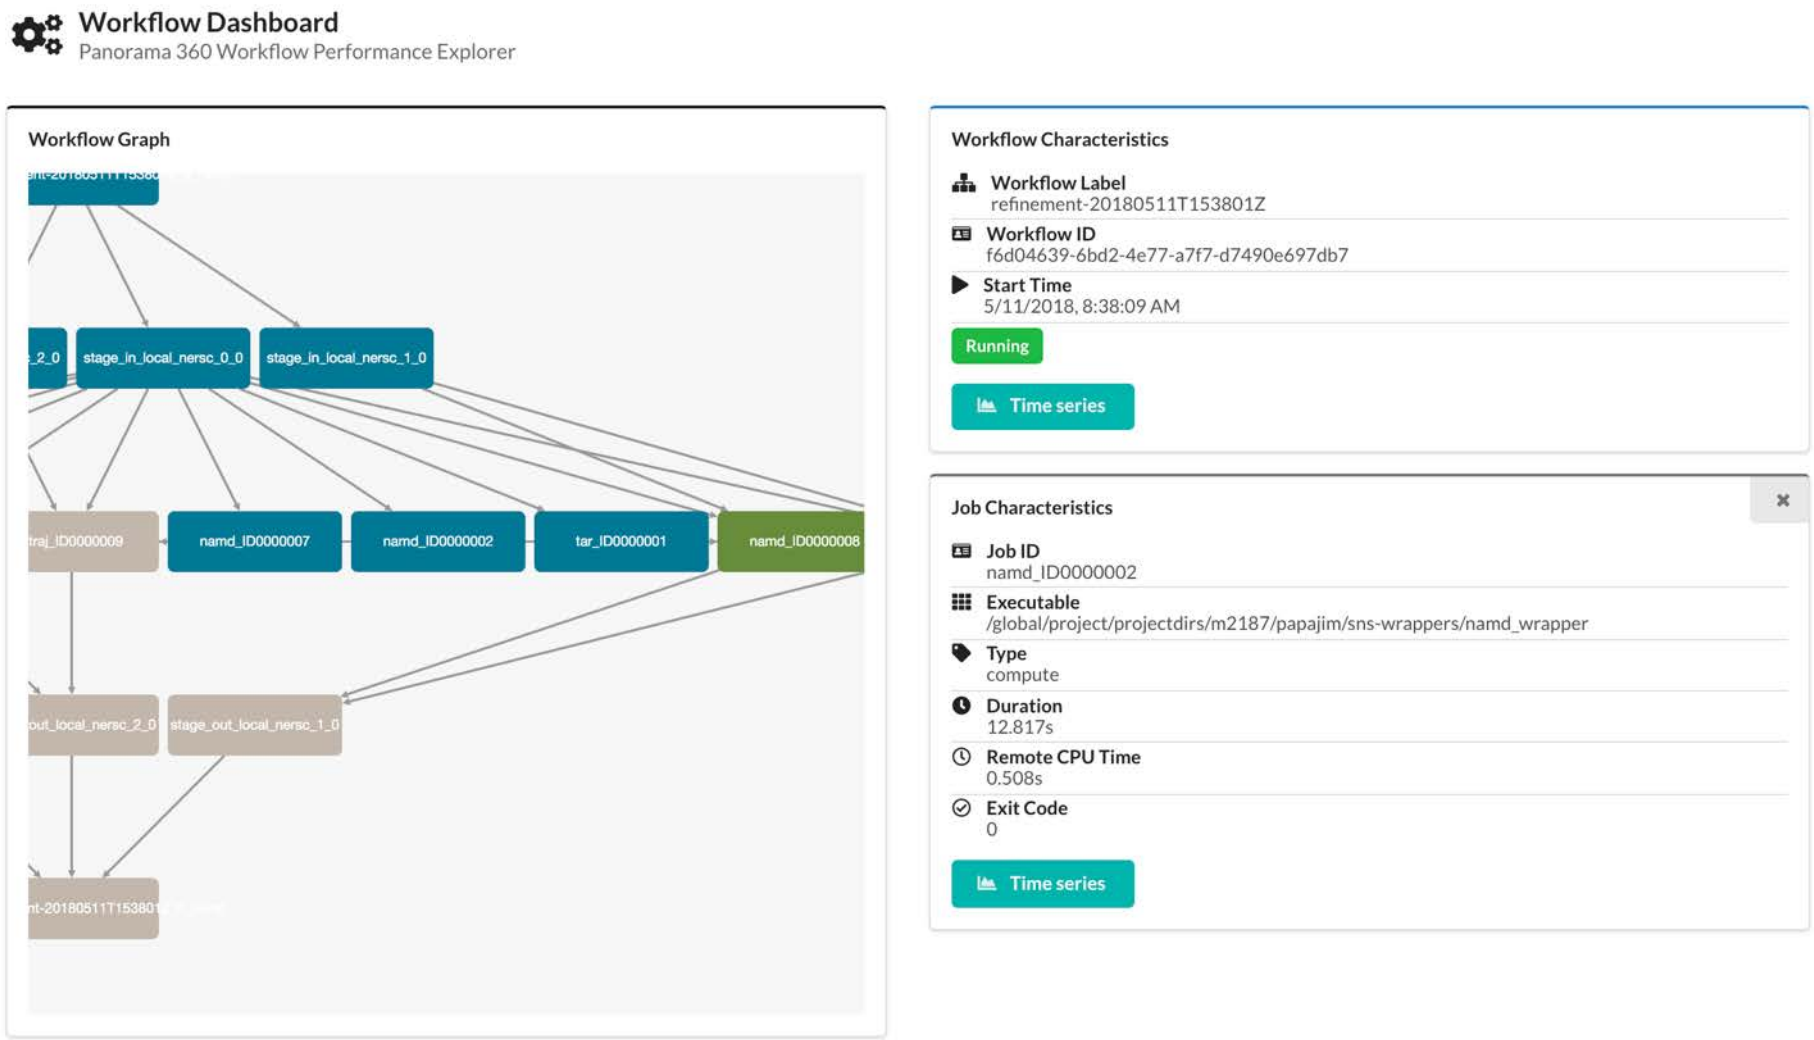

Visualization: Time Series Data of Workflow Performance

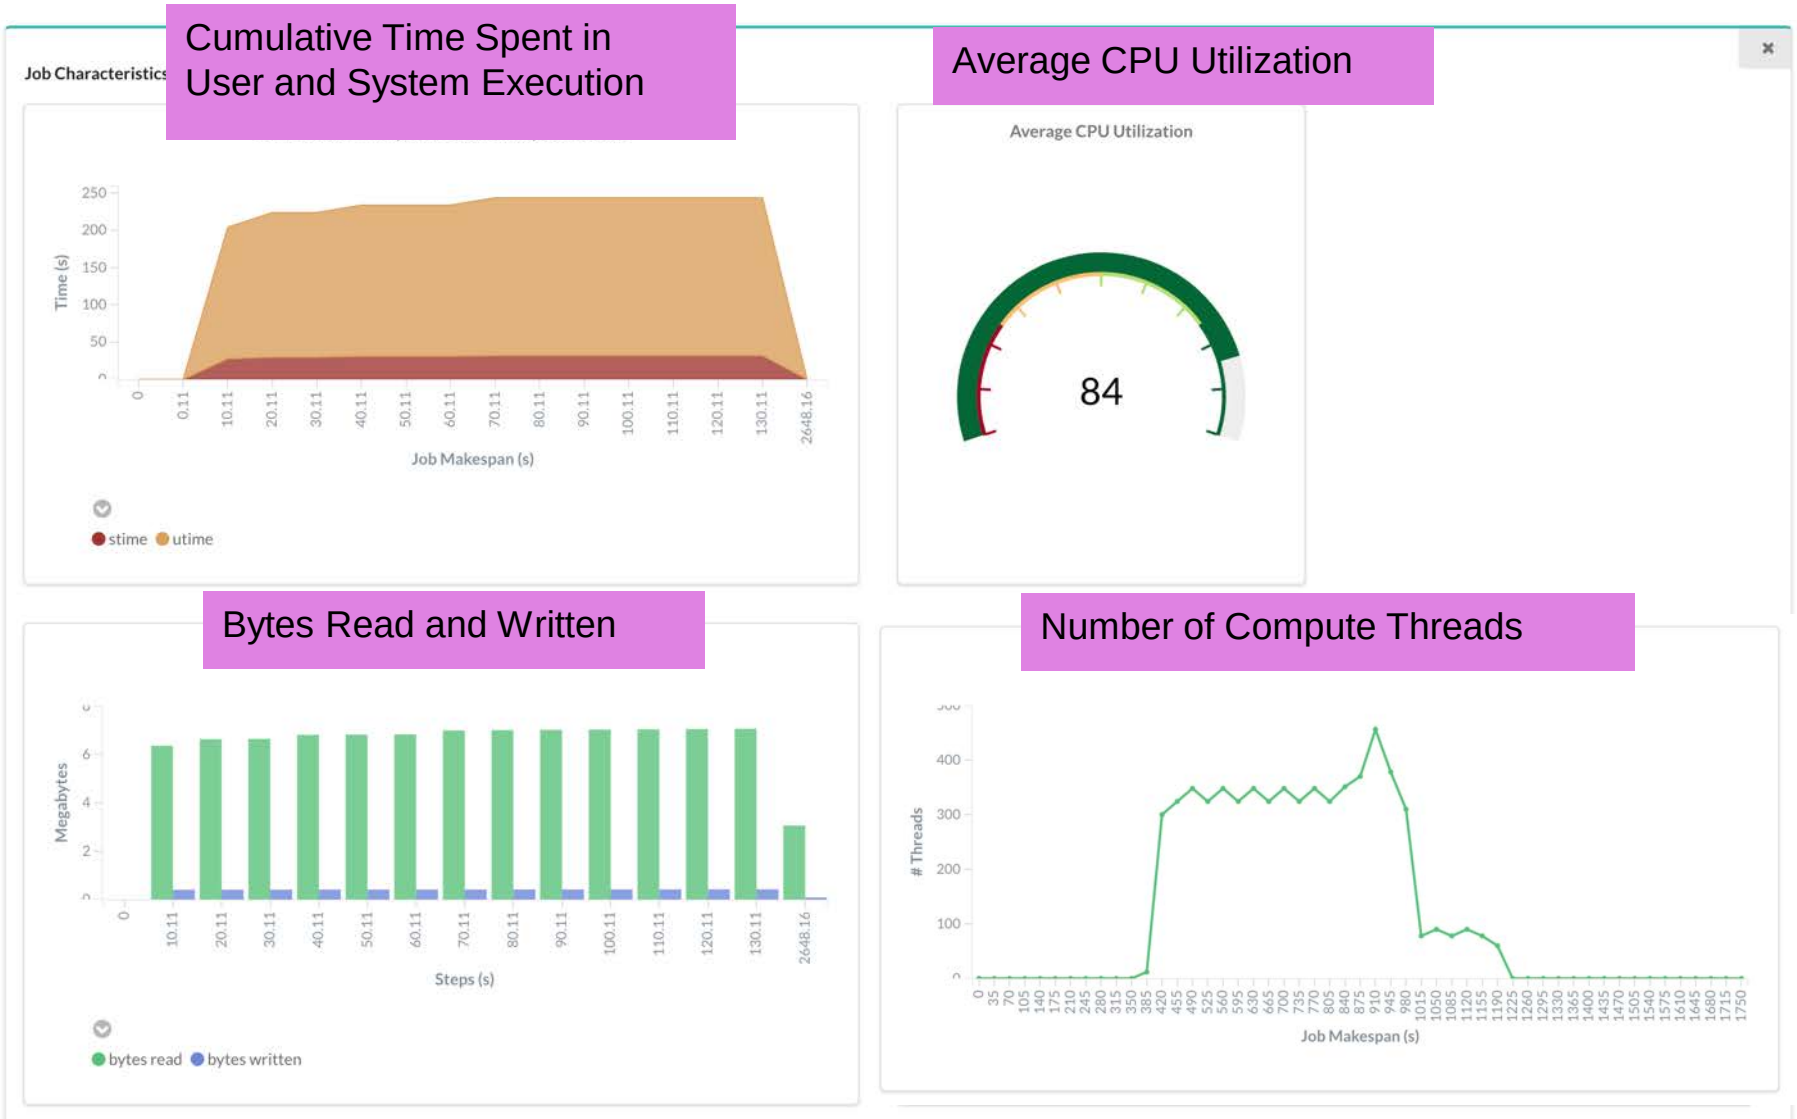

Visualization: Workflow and Infrastructure Utilization Overview

Snapshot of all the currently active workflow ensembles

Snapshot of the data movement across sites

Repository: Open access data

https://data.panorama.isi.edu

https://kibana.panorama.isi.edu

Repository: Organization

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

How to Deploy: Prerequisites

- HTCondor 8.6+:
  - https://research.cs.wisc.edu/htcondor/downloads/
- Pegasus Panorama:
  - Compile from source: https://github.com/pegasus-isi/pegasus/tree/panorama
  - Pre-compiled binaries: http://download.pegasus.isi.edu/pegasus/5.1.0panorama/
- Docker 17.02+:
  - https://docs.docker.com/install/
- Docker Compose:
  - https://docs.docker.com/compose/

How to Deploy: Monitoring Backend (RabbitMQ, ELK Stack)

- On a host that has Docker and Docker Compose installed, clone https://github.com/Panorama360/data-collection-arch
- Change to the cloned directory and execute the following command:

docker-compose up –d

- **Example:**

```
georgpap@iris:~/GitHub/panorama360/data-collection-arch$ docker-compose up -d
Creating network "panorama_net" with driver "bridge"
Creating panorama-rabbitmq      ... done
Creating panorama-elasticsearch ... done
Creating panorama-kibana        ... done
Creating panorama-logstash      ... done
georgpap@iris:~/GitHub/panorama360/data-collection-arch$
```

How to Deploy: Enabling Stampede Events

- In order to get pegasus-monitord to publish **all** of its events to an AMQP endpoint in JSON format, 3 properties must be specified in the workflow's properties file (eg. "pegasus.properties").
  - pegasus.monitord.encoding = json
  - pegasus.catalog.workflow.amqp.url = amqp://[username:password]@<hostname>[:port]/<exchange_name>
  - pegasus.catalog.workflow.amqp.events = stampede.*

- **Example:**

```
19  # Monitord Events
20  pegasus.monitord.encoding=json
21  pegasus.catalog.workflow.amqp.url=amqp://panorama:panorama@amqp.isi.edu:5672/panorama/monitoring
22  pegasus.catalog.workflow.amqp.events = stampede.*
```

- More about stampede events: https://pegasus.isi.edu/documentation/stampede_wf_events.php

What's next?

- Support GPU monitoring (online)
- Generate detailed QoS reports by harvesting the collected logs
- Integrate analysis models for root cause analysis

Learn More

- GitHub: https://github.com/Panorama360
- Website: https://panorama360.github.io
- Extended version of this talk:
  - https://scitech.isi.edu/presentations/2019/pegasus-office-hours-monitoring.pdf
  - https://www.youtube.com/watch?v=fir9ZCxK9Gg

G. Papadimitriou, C. Wang, K. Vahi, R. Ferreira da Silva, A. Mandal, L. Zhengchun, R. Mayani, M. Rynge, M. Kiran, V. E. Lynch, R. Kettimuthu, E. Deelman, J S. Vetter, and I. Foster, "End-to-End Online Performance Data Capture and Analysis for Scientific Workflows," *Future Generation Computer Systems*, vol. 117, pp. 387-400, 2021., (Funding Acknowledgments: DOE DE-SC0012636)

Thank You!

panorama360

George Papadimitriou
Computer Science PhD Student
University of Southern California
email: georgpap@isi.edu