

# **Center-wide Pervasive Performance Monitoring**

# Goals

- Gather application-level, system-level, and infrastructure performance data, correlate and analyze it, and discover new performance problems
- Improve (scientific) productivity for concurrent workloads using shared resources
- Get all the performance data in a single data store (or federated databases) that enable development of new analysis techniques
  - with the goal of sharing these techniques among centers

# Problems

- What APIs for data collection are available at the various layers?
- How do we collect the data without duplicating too much existing functionality or data transfers?
- We (the computing centers) know our infrastructure and software, but are not well informed on others'
- How to convince management and users to allow always-on performance monitoring?
  - policy issues related to security, use of fixed computer-time allocations
  - needs to include methods for turning up the monitoring resolution
- Need to leverage experience data analysts
- Centers reluctant to deploy traditional or BigData databases

# Wish List

- common analysis despite disparate monitoring/data collection
- getting application developer buy-in for continuous performance data collection and evaluation (e.g., TAUdb)
- high-level, low overhead profile reports for every application/job run
- identify sources of performance variability
- monitor power/energy on application level
- define necessary breadcrumbs for connecting layers of performance data
- cross-center analysis to guide deployment on new sites
- resilience data for failure analysis and prediction

# Action Plan

- Develop a taxonomy of existing monitoring and capabilities/deficiencies
- Identify data we don't yet know how to collect
- Identify stakeholders and ask them what it is they want to know that they can't currently obtain, and existing things they wish to improve
- Define metrics and what data is necessary to compute them