Quantitatively Understanding Workflow Performance using prov:Bundle to Associate Traditional Workflow Provenance to System Metrics

[a]Eric Stephan, [a]Todd Elsethagen, [c]Kerstin Kleese van Dam

[a]Bibi Raju, [b]Alok Singh, [b]Ilkay Altintas, [a]Darren Kerbyson

[a]Pacific Northwest National Laboratory, [b] San Diego Supercomputer Center, University of California, San Diego, [c]Brookhaven National Laboratory

Author Email: eric.stephan@pnnl.gov

As high performance scientific compute infrastructures continue to grow in capability and complexity, so do the workflow applications that they serve. The US Department of Energy Integrate end-to-end Performance Prediction and Diagnosis for Extreme Scientific Workflows (IPPD) project is currently investigating an integrated approach to the diagnosis and prediction of scientific workflow performance in extreme-scale computing environments. To gain a quantitative understanding of workflow performance, we need to capture workflow properties and activities, as well as the characteristics and behavior of the underpinning execution environment. The IPPD provenance management solution (ProvEn) (an W3C PROV implementation) was specifically developed to capture this information. ProvEn offers a scalable and load-balanced hybrid database solution for both the traditional semantic workflow provenance and associated system metrics streams representing the environmental factors. The metrics information is hereby stored in an Open Source time-series database, whereas the traditional provenance information is stored in a scalable triple store. During workflow execution, workflow activities are captured using an extension of the PROV-O called the Workflow Performance Provenance (WFPP) ontology. The `prov:Bundle` provides a means to describe how streaming system metrics collected in the time-series database correspond to workflow activities in the triple store.

Introduction

For many domains, the W3C PROV [LEB13] provenance vocabulary offers a complete solution to trace process history; track data origin, changes, and data use; and can describe the role of agents interacting with data and processes. In some situations however, if provenance is disclosed in another form (referred to as the provenance native source) it can still be used to augment traditional PROV graphs, also known as *provenance of provenance*. The W3C PROV accommodates provenance of provenance as

means to describe provenance native sources using the `prov:Bundle` class [LEB13], and as explained in the Linking Across Provenance Bundles technical note [MON13].

In high performance computing applications, workflow history is captured in PROV while observed system metrics are captured in a separate time-series database, both the workflow history and environmental factors influencing workflow behavior is essential evaluating complex, distributed workflows and to clearly identify the sources of commonly observed workflow performance variability. A deeper understanding of these issues will help to optimize workflow design a priori and can help inform workflow optimization at run time.

Performance Metrics in a Provenance Context

The IPPD project [KER14] created the Workflow Performance Provenance (WFPP) [WFP15] vocabulary extending PROV-O to track disclosed workflow activities in the context of system behavior. Typical metrics include memory access, processor use, system load etc. To better understand factors influencing workflow performance, WFPP correlates the disclosed provenance of running workflow processes with instances of the class `proven:NativeSource` (a child of the class `prov:Bundle`). Each instance of `proven:NativeSource` provides properties identifying the time-series database IP address, port, measurement-name, measurement-field, type, unit of measure, measurement rate.

Provenance and Metrics Synchronization

Provenance and metrics are collected using a hybrid triple store and time-series database solution: Provenance Environment (ProvEn) [STEP13, W3C13]. During workflow construction scientists specify WFPP JSON-LD provenance messages [SPO14] that is used to identify workflow activities recorded by running workflow and desired system metrics that will be observed to provide insights for performance.

At runtime the workflow calls the ProvEn Client API (PAPI) along with the disclosing workflow activities stored on the ProvEn database server. The PAPI records `wfpp:Performance_Metric` and any related metrics streamed to a `proven:NativeSource` time-series database.

Provenance and Metrics Usage

Because ProvEn offers hybrid database solution, it can: form compound queries by combining semantic statements and time-series measurements aligned by time, represent time-oriented WFPP semantic statements as time-series metrics in the time-series database (Open source InfluxDB [DIX15]) to leverage the time-series query language, and represent time-series metrics in the ProvEn semantic store to annotate significant events.

To date, ProvEn has been used to reproduce thermal modeling experiments that were initially performed manually. PAPI was used to record each step of the stress test workflow (start and stop times of each workflow stage) while ProvEn initiated requests to collect metrics on observed power and thermal measurements. The disclosed workflow task start/end times serve as time markers for each stress test so that analysis of the observed power and thermal response can be observed with precision.

References

[LEB13] Lebo, Timothy, Satya Sahoo, Deborah McGuinness, Khalid Belhajjame, James Cheney, David Corsar, Daniel Garijo, Stian Soiland-Reyes, Stephan Zednik, and Jun Zhao. "Prov-o: The prov ontology." *W3C Recommendation*30 (2013).

[MON13] Moreau, Luc, and Timothy Lebo. "Linking Across Provenance Bundles." (2013).

[KLE15] van Dam, Kerstin Kleese, Ilkay Altintas, Eric Stephan, Todd Elsethagen, Bibi Raju, and Sriram Krishnamoorthy. "Enabling Structured Exploration of Workflow Performance Variability in Extreme-scale Environments."

[KER14] Energy Integrate end-to-end Performance Prediction and Diagnosis for Extreme Scientific Workflows, http://hpc.pnl.gov/IPPD/

[W3CI13] PROV Implementation Report, https://www.w3.org/TR/prov-implementations/#implementations

[STE13] Stephan, Eric G., Todd O. Elsethagen, Kerstin Kleese van Dam, and Laura Riihimaki. *What comes first, the OWL or the bean*. Technical Report 790738, figshare, 2013. http://dx. doi. org/10.6084/m9. figshare. 790738.

[DIX15] P. Dix, J. Shahid, and T. Persen. InfluxDB. http://influxdb.com/. [10. Aug. 2015].

[SPO14] Sporny, Manu, Dave Longley, Gregg Kellogg, Markus Lanthaler, and Niklas Lindström. "JSON-LD 1.0." *W3C Recommendation (January 16, 2014)*(2014).