

Discoverable Metadata for System Monitoring Data

S. Leak, A. Greiner, A. Gentile, J. Brant

Large systems such as NERSC Cori [1] produce logs from many subsystems - filesystems, networks, compute and service nodes, the batch system. Support staff collect telemetry data relevant to their responsibilities - environmental metrics, utilization snapshots, performance counters from various subsystems, software usage statistics, benchmark job timings. Helpdesk systems collect user requests and fault reports. System monitoring itself takes on the characteristics of a Big Data problem - large volumes of data collected at high velocities, stored in a variety of formats and locations, and for some sources such as user-reported faults, having imprecise veracity.

Insights about fault propagation, drivers and constraints for system utilization, root causes of performance bottlenecks and so on, should be possible via analysis of this data. Efforts to do so are however hampered by the scale and disparity of the data - combining data sources for analysis requires a bespoke approach customised for the analysis target.

The complexity is exacerbated by the fact that data is collected by different teams each for a specific purpose, with only that team's use-case in mind. Making the data useful for another team with a different purpose requires effort so often happens only in response to specific requests. Consequently, data collections often duplicate or overlap others, and are not known to the analyst who might gain from them.

The opportunity this data presents is also amplified by the fact that different centers have similar systems or use common subsystems - the differences and similarities between systems exhibiting eg a common failure, may help in tracing its root cause.

To address these opportunities and challenges, we are working on a metadata-based contribution. We envision a system where a user can query an RDF-graph for logs and data collections relevant to a given topic or event. People collecting data can publish information about what they are collecting without needing to prepare the data itself for wider consumption. This sidesteps many of the challenges around data publication - for example, a dataset that potentially contains sensitive information need not be published, but the knowledge that system manager Alice is collecting the dataset would allow analyst Bob to contact Alice about accessing just the part relevant to Bob's work - a much more tractable problem.

To this end we have defined an RDF vocabulary for system and log data and a schema for annotating system logs, and are working on tools simplify publishing and querying this metadata. The benefit of this approach is that it is machine-readable and network-friendly. In this WIP report we plan to describe the vocabulary and schema and some of the successes and challenges we have found.

[1] *Cori*, 2019, <https://www.nersc.gov/systems/cori/>. Accessed 11 Nov 2019.