



Darshan I/O Runtime Monitoring

Integrating DARSHAN into Continuous HPC Monitoring and Analysis

Presenter: Ann Gentile

Research Collaboration with Northeastern University and SNL HPC Monitoring and Analytics Team

Supercomputing 2022 – Analyzing Parallel I/O BoF



Motivation

I/O performance degradation is one of the main culprits of application behavior variation

Darshan currently collects detailed I/O information about an application run, however it aggregates data over the course of the run and reports out at the end

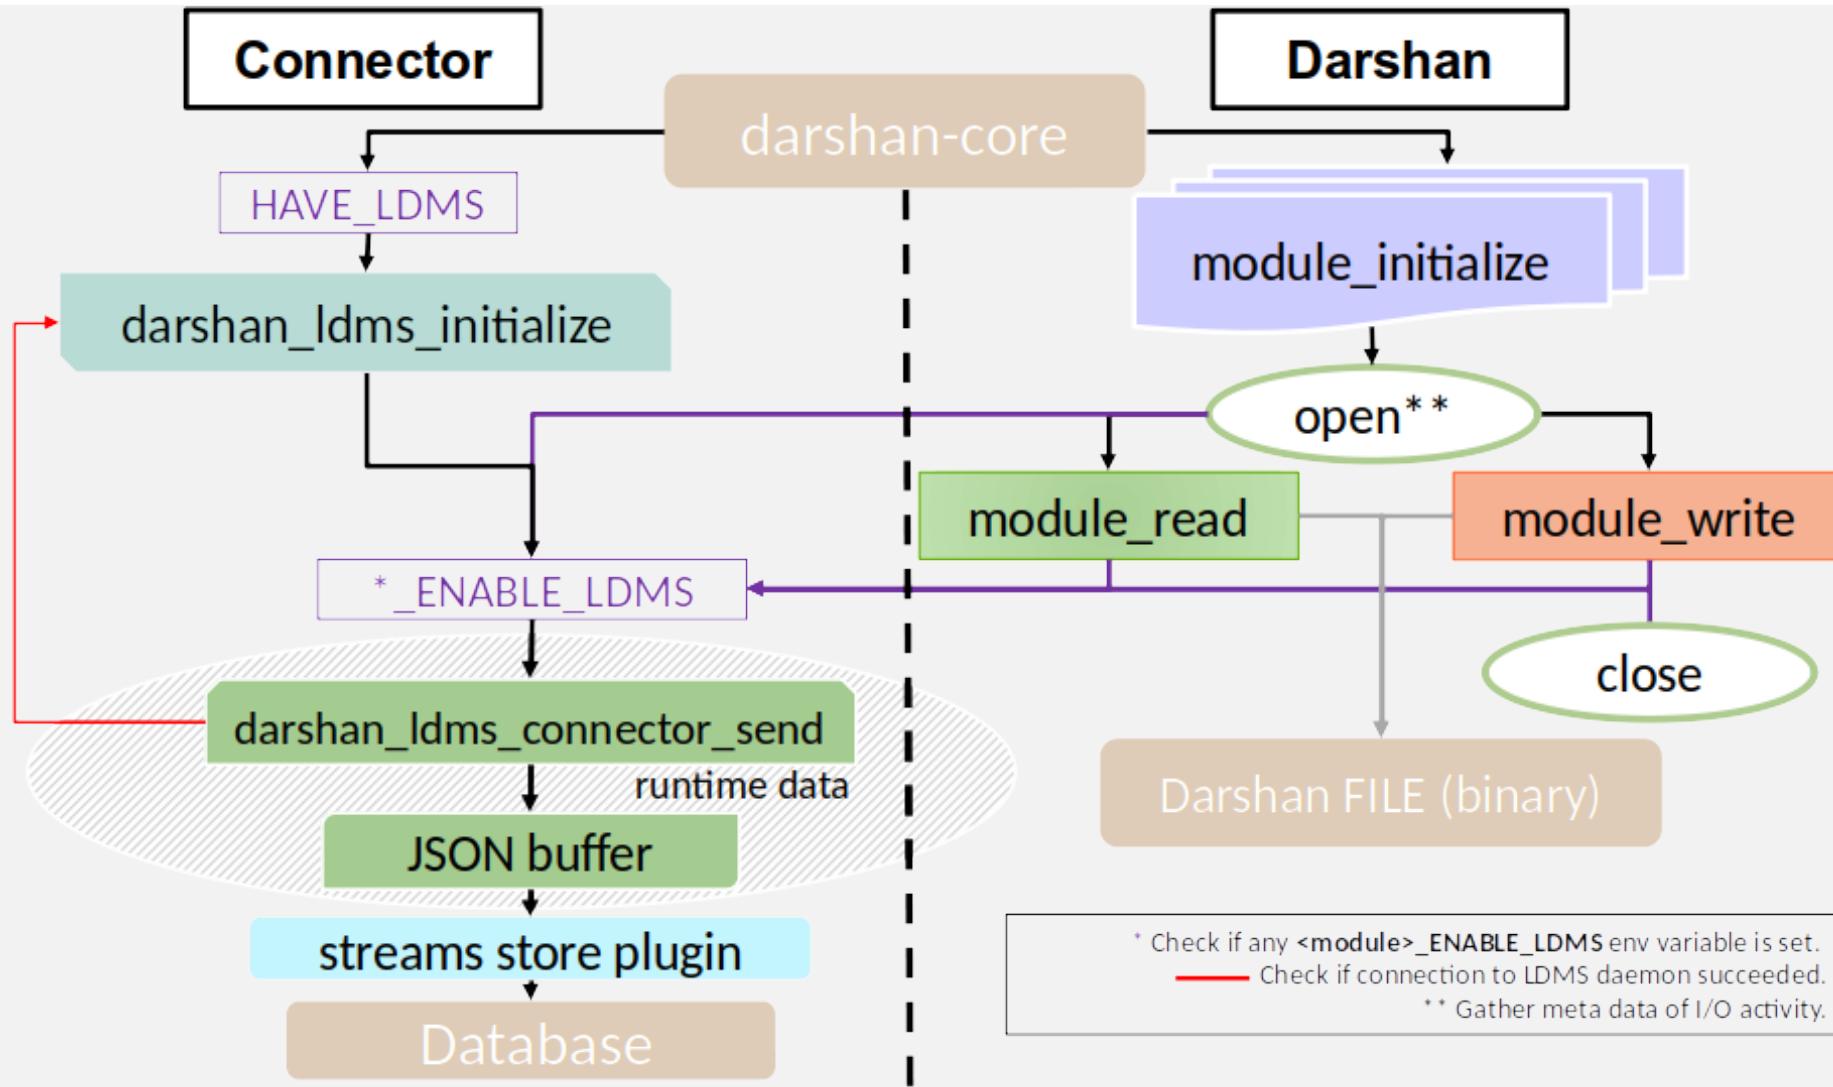
Darshan Extended Trace Plugin can acquire run time data but it is deliberately constrained in buffer/memory, and therefore data, size, and also reports at the end

Timestamped time-series of I/O data available at run time is needed to find errors, correlate with other events, and take actionable response

Benefits of Darshan+LDMS Interoperability

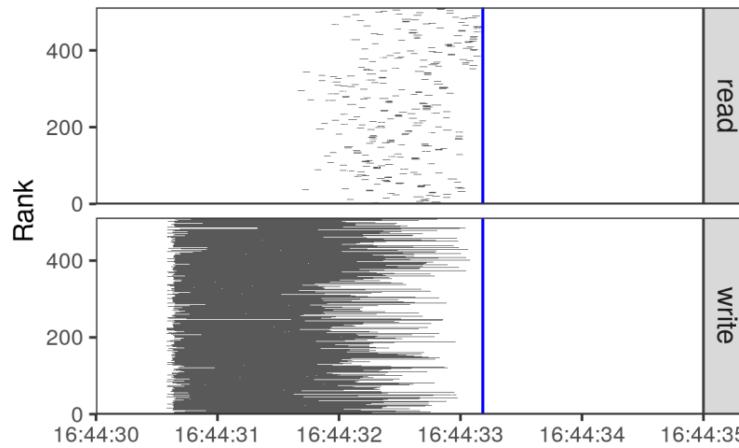
- Captures I/O events at runtime to generate time-series histories of I/O behavior
- Provides absolute time-series timestamps that can be used to evaluate I/O performance in many levels of the I/O subsystem and correlate to system monitoring data / logs
- Captures read/write/close/open/flushes
- Captures POSIX, MPI-IO and STUDIO
- Distinguishes STDIN, STDOUT, and STDERR for POSIX opens

Darshan – LDMS Block Diagram

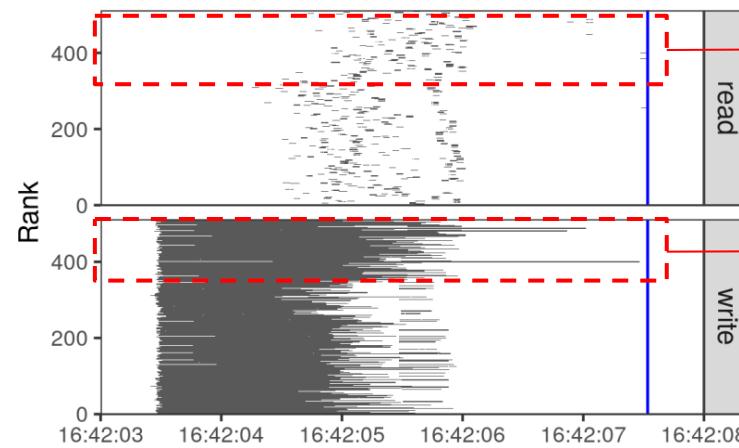


Use Case 1: Detect Runtime Bottlenecks

Fastest HACC-I/O run

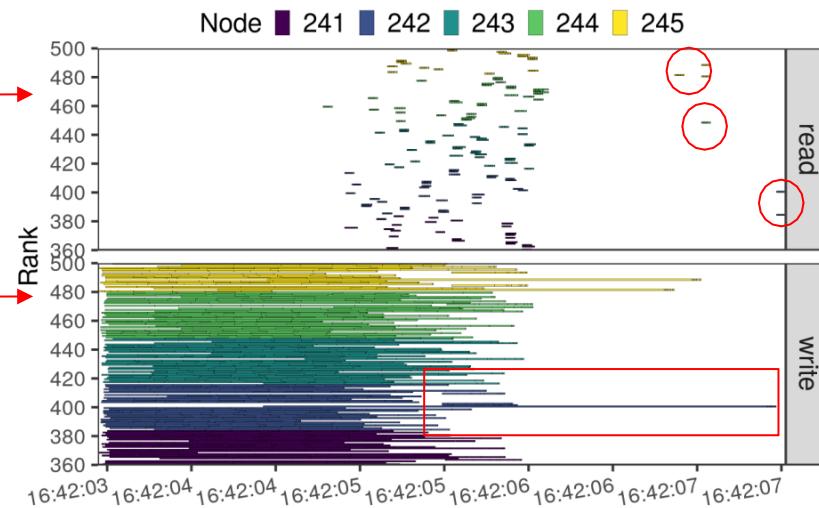


Slowest HACC-I/O run



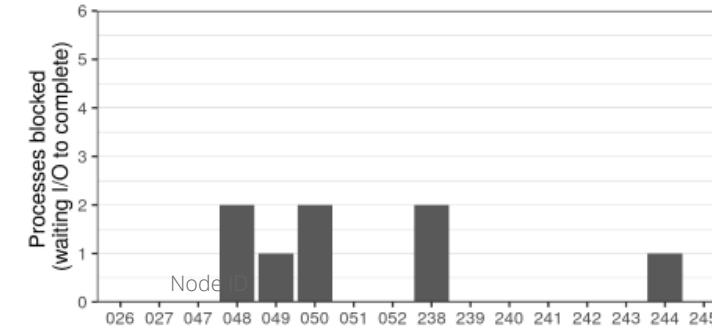
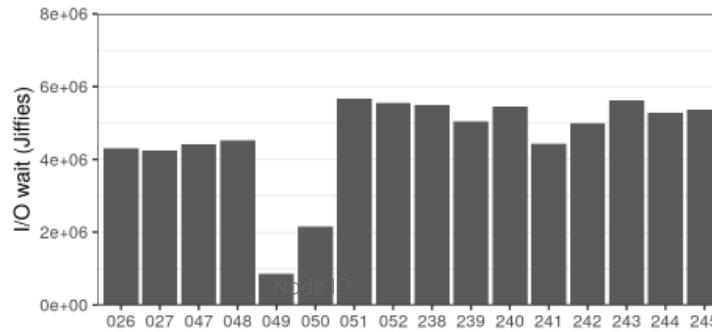
We can identify bottlenecks in the execution and use the **absolute timestamp** to correlate with the system health logs.

```
:245  read  488 10.00000 2022-09-08 12:42:07 0.002477
:245  read  480 10.00000 2022-09-08 12:42:07 0.002501
:242  write  400 20.00000 2022-09-08 12:42:07 2.988597
:242  write  400 20.00000 2022-09-08 12:42:07 0.005700
```



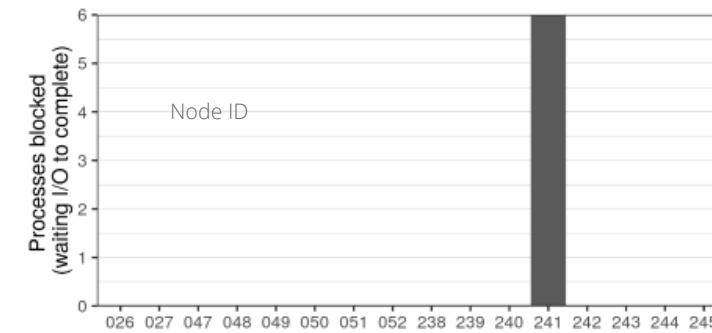
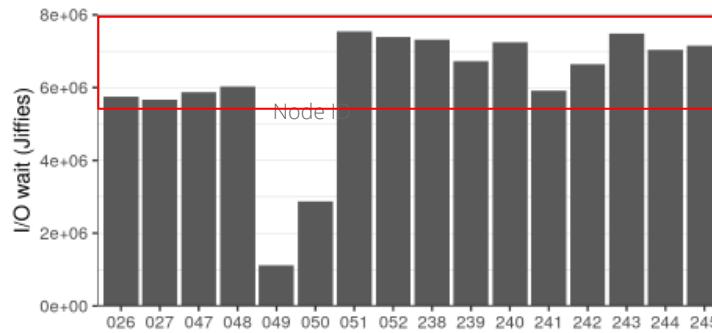
Use Case 2: Correlate with System Data

Fastest HACC-I/O run



/proc/stat iowait data

Slowest HACC-I/O run



Slowest case had 6 processes blocked in the same node (241).

This is one of the slowest nodes.

And 2e+06 Jiffies more in I/O wait than the fastest case.

Other Analysis Opportunities

LDMS also ingests application progress information using a similar architecture (e.g., kokkos, caliper):

- Understand performance impact wrt the science and algorithms.
- Potential for dynamic response

In progress: Use of ML for IO characterization, profile building, and discovery of association with system conditions

Open question: monitoring and analysis of “burst-buffer”/data-staging storage technologies