The Deployment of LDMS on Cray Systems

Deborah Rezanka – Los Alamos National Laboratory

Mentors: Mike Mason, Daniel Illescas and Howard Pritchard (HPC-ENV)

Monitoring HPC Clusters

A High Performance Computing (HPC) Cluster is a collection of computers, connected by a high speed network, working together to run a set of jobs. As HPC progresses towards the goal of Exascale computing these Clusters have grown so large in scale that Administrators must rely on monitoring tools to keep track of the health and efficiency of the systems. When monitoring HPC systems it is crucial that the monitoring process does not interfere with system performance. What is monitored depends on what questions need to be answered. LDMS monitoring gives access to node level data. This can be used to profile individual computer applications, understand how cluster resources are being utilized, and help drive smarter business decisions based on data.

What is LDMS?

The Lightweight Distributed Metrics Service (LDMS) is a low-overhead, low-latency framework for collecting metric data on HPC Clusters. Each node of a cluster runs its own LDMS daemon to collect the desired metrics. Metric information can be transported using Remote Direct Memory Access (RDMA). RDMA supports zero-copy networking by enabling the network adapter to transfer data directly to application memory, eliminating the need to copy data between application memory and data buffers in the operating system.

LDMS has many plugins which allow for customization of data collection and storage. Built in sampler plugins include meminfo, lustre_sampler, lnet_stats, procstat, vmstat and cray_aries_r_sampler. New plugins can be developed to address site specific use cases.

Examples of some questions that can be answered using LDMS data are:

- What percentage of the memory is being used by an application.
- Percentage of cpu used for running non-kernel code.
- How many packets were dropped on a node during a run.

Deployment of LDMS on Cray Systems

Until now monitoring HPC clusters at Los Alamos National Laboratory (LANL) with LDMS version 3 has only been deployed on Commodity Technology Systems (CTS-1) running the Tri-Lab Operating System Stack (TOSS). LDMS version 4 is now available and is being deployed on LANL Advanced Technology Systems (ATS-1) using the Cray Linux Environment (CLE). One challenge deploying LDMS on LANL Cray is that LDMS does not support more than one munge domain (shared security realm). The work-around was to have the connection between the level 1 aggregator nodes and the level 2 aggregator communicate over the external network using a different form of authentication. The main challenge with LDMS monitoring on Cray systems is the volume of data being collected. Trinity can produce up to four terabytes (TB) of LDMS data per day.

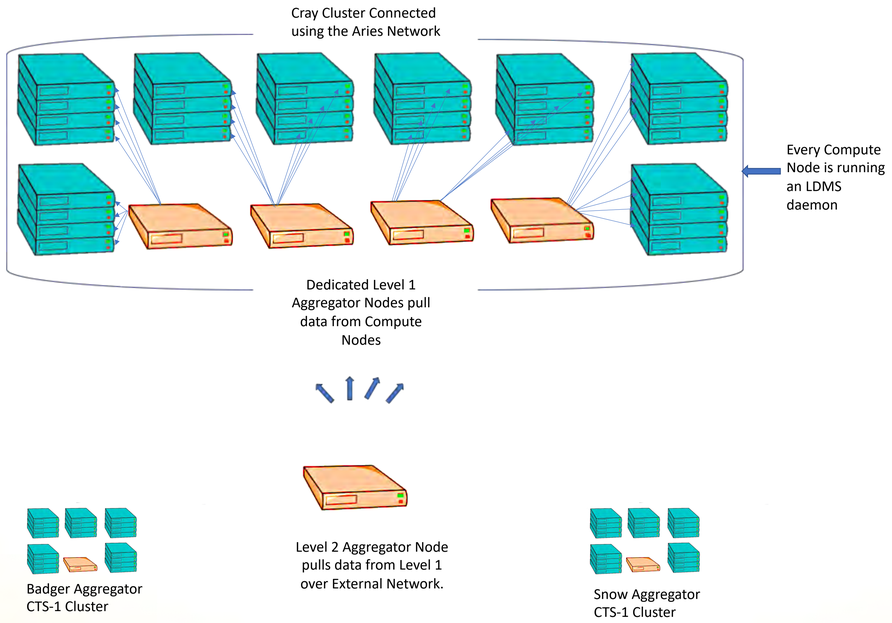

Grafana view of LDMS data using the OpenTSDB data source

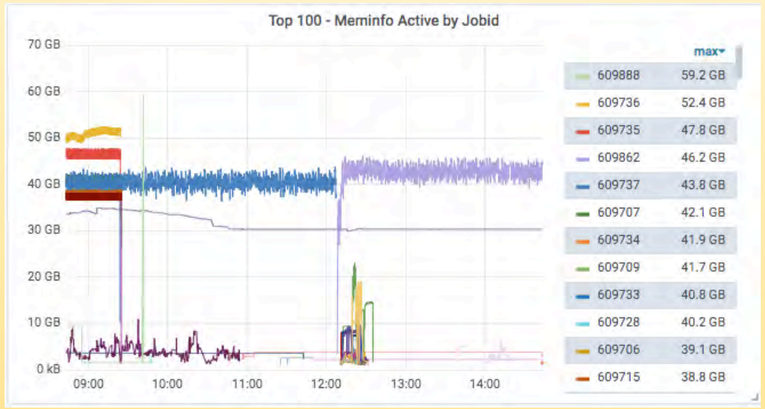

Data Flow

Each node on the Cray Cluster runs its own LDMS daemon. Only the most recent data is held on the node in memory. Data is pulled by RDMA, over the Cray Aries network, by the level 1 aggregator nodes within the cluster. A Level 2 aggregator node, outside the cluster, pulls the data from Level 1 over an external network. LDMS has several storage aggregator plugins. Some LDMS storage formats are RabbitMQ, Scalable Object Store (SOS), comma separated value (CSV) and flatfile. RabbitMQ is a message broker that transfers the data to our Tivan data analysis system. The SOS plugin saves the data to local storage on the final aggregator node.

OpenTSDB vs SOS

Two types of databases are currently being used to store LDMS data, SOS and OpenTSDB. Research is being conducted into the “best” way to analyze the data. Grafana, open source software for time series analytics, is being used to visualize the data with both SOS and OpenTSDB

OpenTSDB is a scalable, distributed time series database that runs on Hbase. Data can be queried through command line tools, a web application programming interface (API) and viewed as a GnuPlot graph. Several open source tools, such as Grafana and Bosun, can also access OpenTSDB data. LDMS version 3 and 4 OpenTSDB data is currently being visualized using Grafana via our Tivan data analysis system. OpenTSDB allows the user to perform many transformations on the data directly through the Grafana query interface.

SOS is a high-performance, indexed, object-oriented, relational database. SOS was designed for LDMS data. The query interface is similar to SQL which is a standard query language for relational databases. Sosdb-grafana is a custom plugin for Grafana. So far, attempts to visualize SOS data on Grafana have not been very successful. Data transformation needs to be performed outside of the SOS database and then added back to it in order to be displayed.

Conclusions and Future Work

More work is needed on the SOS plugin for Grafana to get it production ready and able to do more flexible queries. Jupyter notebook is one option to pursue for performing data analytics on SOS databases.

Work is moving forward between Los Alamos, Livermore and Sandia national laboratories to develop common tools for the data analysis of HPC monitoring data through the Common Computing Environment initiative (CCE). One goal for the CCE is the development of a suite of monitoring tools for HPC.

To accommodate the huge amount of monitoring data that exascale computing will generate, infrastructure for monitoring will need to be an integral part of the system, not an afterthought. Having aggregator nodes that can communicate with compute nodes over the high speed network using RDMA is key to LDMS having a minimal impact on HPC performance. The next generation HPC Cluster should include dedicated monitoring nodes for this reason.

Here is the query used in Grafana to produce the graph to the left. The gexp Type refers to OpenTSDB’s implementation of Graphite expressions.

Los Alamos NATIONAL LABORATORY EST. 1943

LA-UR-19-27104

National Nuclear Security Administration