

Tools to improve your efficiency


Customers are often facing lower than expected efficiency for their application workloads. Tools to investigate runtime behavior of applications are precious to understand where the bottlenecks are

As I/O is one of the least known part of job executions, ATOS has developed a comprehensive set of tools to understand and improve efficiency of running applications in complex workloads.

This collection of software and hardware tools already available are :

- Bull IO Instrumentation as an application I/O profiler
- Bull IO Pattern Analyzer as a multiWjobs comparison tool and accelerators usage automation
- Bull Fast IO Libraries as software accelerator modules
- Bull Smart Burst Buffer as a versatile Flash storage accelerator

A smart and focused use of these tools will provide valuable information on running applications and facilitate investigations on unoptimized jobs among complex workloads.

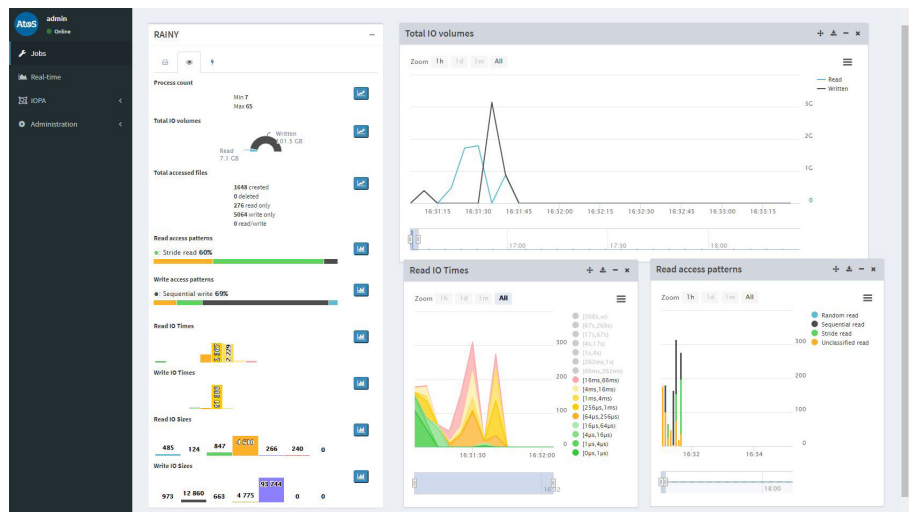


Bull IO Instrumentation

Bull IO Instrumentation is an intuitive profiling tool that helps the system administrators, developers and product support experts analyzing the IO activity generated by HPC jobs.

A role based access control is included to provide different level of information to various predefined group of users: system administrators, operators or regular users.

This tool collects, stores and displays a whole set of IO related metric through a very ergonomic graphic interface.




Bull IO Instrumentation has been developed as a highly scalable application to minimize the impact of collecting samples of its metrics. It is implementing a hierarchical 2 levels MongoDB database.

I/O interception is transparently done on compute nodes with I/O activity and statistics are sent to IO gateways on a few seconds period. Then a master node is providing consolidated metrics through a standard Web interface. These functions can be implemented on service nodes which are very common on large scale compute clusters.

Instrumented jobs will provide information as:

- Volumes read and written




- I/O patterns (sequential, random ...)
- IO durations and sizes
- ...


New metrics helping to identify applications with heavy I/O traffic -like a focus on parallel MPI jobs or Lustre filesystem activities are planned.

Bull IO Pattern Analyzer

Bull IO Pattern Analyzer is a plugin of Bull IO Instrumentation to compare the behavior of a group of jobs. It is mostly targeted to analyze various runs with different data sets of a chosen application to identify I/O acceleration opportunities. Ultimately, it will automatize the use of IO accelerators.



Based on such histograms related to various runs of a specific application, graphs –so called “Bubble” graphs can help to identify singularities in a very efficient way.



Ultimately, this knowledge can be taken in account through Machine Learning algorithms to calculate potential acceleration and automatically select the appropriate accelerator and parameters for further executions through the SLURM orchestrator.

/home_rfs/2015_10/BU...	2017-11-21 13:43:21 UTC	12 minutes 4 seconds	0	🔍
/bin/bash	2017-11-21 10:22:27 UTC	2 hours 39 minutes 9 seconds	Calculate	🔍
ExtrapolData	2017-11-15 10:04:51 UTC	8 minutes 20 seconds	100%	🔍

Bull Fast IO Libraries

Bull Fast IO Libraries is a set of software modules that addresses HPC malformed IOs with various transparent system optimizations.


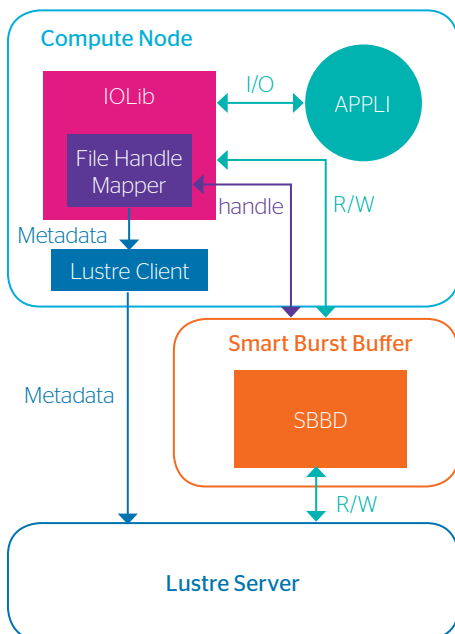
The available version is focused on prefetching and small Read IO aggregation. Future developments will handle large read patterns and also optimize write sequences.

Bull Smart Burst Buffer

Bull Smart Burst Buffer is a Flash hardware accelerator that can be included in the I/O data path and accelerate I/O on specific files for targeted applications.

ATOS will offer the Flash accelerator which addresses the most various use cases as possible. So the software stack embedded in this Flash Appliance will support 2 modes of storage allocation through the High Speed interconnect :

- Transparent caching on a file by file basis for selected applications. Data is written asynchronously to the shared filesystem
- Static OnDemand allocation of fast persistent storage per applications' request. Data is private to each node and stays on NVMe Flash storage



For the 1st mode, transparent caching will rely on I/O calls interception through the scheme implemented in the Bull IO Instrumentation library.

For the 2nd mode, applications will be able to explicitly request –through SLURM static allocation per node of NVMe storage through the High Speed Interconnect fabri.

Bull Data Management tools will improve your efficiency

For more information or to organize a test please contact: hpc@atos.net