



Contribution ID: 40

Type: not specified

Advanced Monitoring Techniques for the Atlas TDAQ Data Network

Friday, 9 May 2008 10:15 (30 minutes)

We describe the methods used to monitor and measure the performance of the Atlas TDAQ data network. The network consists of four distinct Ethernet networks interconnecting over 4000 ports using up to 200 edge switches and five multi-blade chassis. The edge networks run at 1Gbps and 10Gb/s are used for the detectors raw data flow as well as at the cores of the data flow networks. The networks feed event data to farms of up to 3000 processors. Trigger applications running on these processors examine each event for acceptability and assemble the accepted events ready for storage and further processing on Grid linked data centers. We report in detail on the monitoring and measurement techniques deployed and developed.

Summary

Monitoring and measurement must go beyond merely assuring that the installed system network is functional and that its performance meets the design specifications. The system will regularly be operated above its design point as applications are refined to optimally consume all available computing and network resources. Even before that stage users will typically reduce the filtering criteria to accept an increasing number of events until the point that something fails. This could be for example saturation due to insufficient computing power or network overload and subsequent backpressure. Thus for such a network it is the norm is to be overloaded, links running continuously at 100% are to be expected and packet discards in the thousands per second will happen. During all this, the network monitoring system must still be able to distinguish between real and self inflicted problems as well as identify traffic anomalies that may be preventing the system from going even faster. While redundancy is built in there will still be performance loss caused by equipment failure and identification of the appropriate repairs must be made rapidly and accurately. Experience has shown that analysis of individual, or aggregate traffic flows is extremely useful for system diagnosis and performance measurement but the sheer size of the system imposes issues of scaling for any of the traditional measurement techniques. We describe the approach taken to meet these challenges.

Primary authors: AL-SHABIBI, Ali (CERN, University of Heidelberg); MARTIN, Brian (CERN); LEAHU, Lucian (CERN, Politehnica Bucharest); CIOBOTARU, Matei (CERN, UC Irvine, Politehnica Bucharest); BATRANEANU, Silvia (CERN, Politehnica Bucharest); STANCU, Stefan (CERN, UC Irvine, Politehnica Bucharest)

Co-authors: DARLEA, Georgiana (CERN, Polytech' Savoie, Politehnica Bucharest); IVANOVICI, Mihail (University of Brasov, Romania)

Presenter: CIOBOTARU, Matei (CERN, UC Irvine, Politehnica Bucharest)

Session Classification: Networking infrastructure and computer security

Track Classification: Networking infrastructure and computer security