

End-to-End Network/Application Performance Troubleshooting Methodology

Authors:

Wenji Wu (Fermilab), Andrey Bobyshev (Fermilab), Mark Bowden (Fermilab), Matt Crawford (Fermilab), Phil DeMar (Fermilab), Vyto Grigaliunas (Fermilab), Maxim Grigoriev (Fermilab), Don Petravick (Fermilab)

Abstract:

The computing model being developed for the LHC is globally distributed and grid-based. In such a computing model, the LHC physics experiments data must be reliably and efficiently transferred from CERN to Tier-1 regional centers, processed, and distributed to Tier-2 regional centers around the world. Network performance issues can be a major impediment to the success of this complex, multi-tiered data grid. Factors that affect overall network/application performance exist on the network end systems themselves (application software, operating system, hardware), in the local area networks that support the end systems, and within the wide area networks utilized. Since the computer and network systems are globally distributed, it can be very difficult to locate and identify the factors that are negatively impacting application performance. In this paper, we present and discuss an end-to-end network/application performance troubleshooting methodology developed and in use at Fermilab. The core of our methodology is to narrow down the problem scope. We apply a divide and conquer strategy. The overall complex problem is split into two distinct sub-problems: network end system diagnosis and tuning, and network path analysis. After satisfactorily evaluating, and if necessary resolving, each sub-problem, we conduct end-to-end performance analysis and diagnosis. The paper will discuss tools utilized as part of the methodology. The long term objective of the effort is to enable end users to conduct much of the troubleshooting methodology themselves, without having to call upon network and end system “wizards”, which are always in short supply.