

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 models for LHC experiments are globally distributed and grid-based. In such a computing model, the experiments' data must be reliably and efficiently transferred from CERN to Tier-1 regional centers, processed, and distributed to other centers around the world. Obstacles to good network performance arise from many causes and 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. Since the computer and network systems are globally distributed, it can be very difficult to locate and identify the factors that are hurting application performance. In this paper, we present an end-to-end network/application performance troubleshooting methodology developed and in use at Fermilab. The core of our approach is to narrow down the problem scope with 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 we use as part of the methodology. The long term objective of the effort is to enable end users to conduct much of the troubleshooting themselves, before (or instead of) calling upon network and end system "wizards," who are always in short supply.