

Performance Analysis of Linux Networking –Packet Receiving

Tuesday, 14 February 2006 14:20 (20 minutes)

The computing models for HEP experiments are becoming ever more globally distributed and grid-based, both for technical reasons (e.g., to place computational and data resources near each other and the demand) and for strategic reasons (e.g., to leverage technology investments). To support such computing models, the network and end systems (computing and storage) face unprecedented challenges. One of the biggest challenges is to transfer physics data sets –now in the multi-petabyte range and expected to grow to exabytes within a decade –reliably and efficiently among facilities and computation centers scattered around the world. Both the network and end systems should be able to provide the capabilities to support high bandwidth, sustained, end-to-end data transmission. Recent trends in technology are showing that although the raw transmission speeds used in networks are increasing rapidly, the rate of advancement of microprocessor technology has slowed down over the last couple of years. Therefore, network protocol-processing overheads have risen sharply in comparison with the time spent in packet transmission, resulting in the degraded throughput for networked applications. More and more, it is the network end system, instead of the network, that is responsible for degraded performance of network applications. In this paper, the Linux system's packet receive process is studied from NIC to application. We develop a mathematical model to characterize the Linux packet receive process. Key factors that affect Linux systems' network performance are analyzed.

Primary author: Dr WU, Wenji (Fermi National Accelerator Laboratory)

Co-author: Dr CRAWFORD, Matt (Fermi National Accelerator Laboratory)

Presenter: Dr WU, Wenji (Fermi National Accelerator Laboratory)

Session Classification: Computing Facilities and Networking

Track Classification: Computing Facilities and Networking