Advanced Networking Infrastructure Projects for Computing and Collaboration

June 14, 2007

PI: Gurcharan S. Khanna, Rochester Institute of Technology
http://www.rit.edu/~gskpop
gurcharan.khanna@rit.edu
585-475-7504

Sponsor: CISCO

Executive Summary

The Research Computing group at RIT is building a research computing infrastructure to support the growing needs of researchers across all disciplines. High performance networks are a crucial part of the overall plan that is targeted to address two specific needs for research use:

- High-speed/low-latency interconnects for high performance message passing and data transfers within clusters and across clusters in grids
- High-speed/low-latency multipoint connections for live sustained uncompressed high definition multicast video streams

Current solutions for these needs include expensive, proprietary systems that are very specialized for their targeted needs. We propose to evaluate commodity, standardized solutions that are less expensive, more universal, and easier to support, such as 10 Gigabit Ethernet. We feel that the performance advancement of such generic networking solutions make them an attractive alternative to traditional dedicated solutions.

We propose to install, benchmark, and evaluate these in terms of a cost/benefit analysis, especially by being able to implement a single solution for these multiple needs. While such standard solutions as 10 Gigabit Ethernet are virtually commodity products now, they are not universally being adopted as readily as one might expect. Our intent is to demonstrate the viability of integrating easily available high performance network components into an advanced research computing infrastructure.

RIT currently does research on advanced collaboration environments (ICE Lab, http://www.rit.edu/~rc/ice.html), grid development (NYSGrid, http://www.nysgrid.org), and networking protocols (NSSA, http://www.nssa.rit.edu). Personnel experienced in these research areas will

support and guide the deployment and testing of the proposed advanced networking capabilities. This project is planned to take six months from initial design to first testing of results, with another six months for a complete evaluation.

Partners

RIT currently deploys a campus-wide distribution of CISCO networking products, including its 1 Gigabit Ethernet backbone, 1 Gigabit building routers, and 1 Gigabit to the desktop/lab in research areas. RIT and the University of Rochester share a 10 Gigabit Ethernet ring around the city of Rochester. By deploying 10 Gigabit networks for collaboration and high performance computing, the research component of the RIT community will enjoy more productive research accomplishments that will be partly due to CISCO advanced infrastructure support.

Proposed Projects

1) Compute cluster backplane switching fabric: we are currently in the process of disassembling and moving a 52 node IBM linux cluster to a new location. The current Black Diamond Extreme Blade switches are large, consume lots of power, and run at 1 Gbps. We would like to take this opportunity to replace this switch with a current generation 10 GigE switch and evaluate the performance and behavior under real research computing conditions.

We have every expectation that the improvement over the current technology employed will be significant, but may fall short of other more specialized protocols such as InfiniBand and MyriNet technologies. However, the price/point for a standardized 10 GigE system may prove to be extremely favorable.

2) The Research Data Design Team will be looking at a comprehensive strategy for dealing with research data. This includes data throughput, storage, data management, security, and web access. For the specific area of data throughput, the research computing staff are looking at the performance gain that can be achieved with upgrading to 10 GigE switches and determining what tuning needs to be done to make it optimal. Our goal is put all research systems on 10 GigE networks for transport of very large research data files, both on campus and off-campus via the NYSGrid.

3) MPI (message passing interface) between two different systems across subnets is generally avoided due to high relative latency. This project

would evaluate the size of jobs that could be run effectively across subnets by combining two or more clusters in different subnets with 10 GigE switches and routers. The network performance requirements are similar to 2. Successful results would allow much larger jobs to be run by spanning loosely coupled clusters.

4) Real-time uncompressed High Definition video for interactive collaboration. The ICE Lab in CASCI is about to start on Phase II of an advanced collaboration technologies project that is focused on real-time interactive collaboration tools with a special emphasis on high quality video.

We are working with partners at the University of Washington and the University of Southern California to implement and enhance their systems to transport high definition video over IP for real time conferencing. These systems require 1.5 Gbps each way for each participant in the conference for full 1080x1920 video. We will be contributing to the effort to enhance the software and hardware to be more efficient in latency performance as well as user interface design.

This conferencing system will be used for on-campus video connectivity between related research groups as well as for off-campus distance learning and other collaboration projects already in place.

5) Professor Xiaojun Cao in the NSSA Department is currently working on creating improved transport mechanisms beyond TCP and UDP to take advantage of high speed networks over ethernet and IP. His experimental and simulated work will provide actual confirmation and testing of his theories and models on current 10 GigE systems.

Expected Results

We anticipate that a successful evaluation of 10 Gigabit Ethernet implementations in the above projects will result in a less costly, easier to maintain, and more greatly extensible network fabric for computing and collaboration by using advanced CISCO products in place of more highly specialized solutions currently being deployed.

Personnel

Gurcharan S. Khanna, Director of Research Computing, Director of the Interactive Collaboration Environments Laboratory. Responsible for overall direction and management of the project.

Andrew Elble, Senior Network Engineer, ITS. Responsible for design, configuration, and installation of the new networks.

Professor Xiaojun Cao, Professor, Networking, Security, Systems Administration Department. Consulting for networking protocol development and optimization.

Ed Brannin, Manager, ICE Lab. Responsible for testing and evaluating devices and applications in the new network infrastructure environment. Responsible for managing student workers on the project.

Student employees of the ICE Lab. Part-time and full-time (Co-op).

Funding

Staffing

- ICE Lab Manager, 50% FTE, six months, $15,000
- Student worker, full-time Co-op, six months, $12,000

Equipment (exact models TBD)

- 48 port network switch for IBM cluster backplane interconnect
- 10 Gigabit ethernet PCI-Express network cards (5)
- 10 Gigabit ethernet switches (5)
- 10 Gigabit ethernet router cards (5)
- 10 Gigabit Ethernet border router card (1)

References

http://researchchannel.org/tech/ihdtv.asp

http://ultragrid.east.isi.edu/

http://www.gloriad-kr.org/hdtv

https://sitola.fi.muni.cz/igrid