DOE/MICS/ECPI Project Summary

Project Title: Network QoS for Experimental Magnetic Fusion Energy Research
PI: David P. Schissel and Reza Shakoori
PI Institution: The DIII-D National Fusion Facility at General Atomics

DOE Collaborator: Chin Guok, ESnet
Undergraduate Student: Gracie Elqura, UCSD
Project Website: http://www.fusiongrid.org/network/

Abstract:

The ultimate goal of this project is the deployment to an operating magnetic fusion experiment (DIII-D) of a robust and reliable QoS capability for networks that is routinely used during experimental operation. Such a Network QoS capability will represent a dramatic shift for the fusion community. Presently, computational resources for near-real-time data analysis in support of experimental operation are limited to local dedicated systems. Such a network QoS capability could lead to, for example, the usage of a massive supercomputer such as at Seaborg at NERSC or the National Leadership Computing Facility at ORNL to support near-real-time analysis of fusion experimental data; something that today is not even considered within the experimental fusion community.

Project Description:

In today's magnetic fusion experimental sessions, hardware/software plasma control adjustments are made in near-real-time as required by the scientific program. These adjustments are debated and discussed among the experimental team with decisions for changes to the next pulse being informed by data analysis conducted within the roughly 30-minute between-pulse interval. This mode of operation places a large premium on rapid data analysis to dramatically improve the efficiency of experimental sciences. Experience on today's experiments should aid in the design of systems for ITER, the next major step to be located in France in an international program aimed at proving the scientific viability of controlled fusion as an energy source. ITER is a burning plasma magnetic confinement experiment, is the Office of Science's number one facility priority, and will be the largest and most expensive scientific instrument ever built for fusion research. Like the current generation of experiments, ITER will put a premium on real-time interactions with data and among members of the geographically distributed research team.

This project has a close working relationship with the On-Demand Secure Circuits and Advance Reservation System (OSCARS) Project and the Network Project at NERSC where the computer science research for on-demand provisioning of guaranteed bandwidth over ESnet is being performed. The project's first goal is to demonstrate in a test OSCAR's circuit over ESnet the transmission of MDSplus stored fusion data from NERSC to General Atomics under QoS conditions. MDSplus is by far the most widely used data system in the international fusion program. Based on a client/server model, MDSplus provides a hierarchical, self-descriptive structure for simple and complex data types, and is currently installed and used in a variety of ways by about 30 experiments, spread over 4 continents. MDSplus is also used to securely transfer data on FusionGrid. Thus, the demonstration of QoS via MDSplus will indicate that any fusion data transmitted via MDSplus should be able to be given priority. With the ubiquity of MDSplus in the fusion community, this becomes a very powerful capability.

For this project, a dedicated computer has been installed on a private VLAN directly attached to the DIII-D core Cisco 6500 that is directly connected to the Juniper M7i ESnet border router. A test Label Switched Path (LSP) was then established between this computer (GA) and NERSC for initial QoS testing. Only traffic from the specific test computers (source/destination IP addresses) was injected into the LSP. Since the LSP was using the production DIII-D network, tests were of short duration and performed where practical later at night. A DIII-D C-based utility for measuring MDSplus data throughput was used for the LSP testing. The client code makes multiple serial requests for a 10 MB size of piece data from the MDSplus server and the average time is used to calculate network throughput. The DIII-D test network machine was configured as an MDSplus server and the NERSC network test machine was configured as an MDSplus client. The test was run where the MDSplus testing utility was run on the NERSC machine to fetch data from the DIII-D MDSplus server. Using the production DIII-D network, an average throughput of 23 Mb/s was observed. Utilizing a 130 Mb/s UDP blast from DIII-D to NERSC resulted in the MDSplus throughput dropping by an order of magnitude (3.2 Mb/s). When this same MDSplus traffic was given priority (Expedited Forwarding), the network throughput climbed back up to the maximum value of 23 Mb/s. Thus MDSplus data transfer with network QoS has been demonstrated.

Although successful, this initial test is not sufficient for production deployment. Due to the large number of off-site collaborators at DIII-D, the quantity of MDSplus data flowing out of the DIII-D network is large. Yet, only certain packets will be deemed worthy of QoS, those attached to certain analysis codes deemed critical by the experimental team. Thus, a method to mark certain MDSplus packets is necessary and the project is testing the usage of DSCP bits to satisfy this demand. The project's next steps are to install the equilibrium fitting (EFIT) code that calculates plasma shape at NERSC and to test the ability to transmit these packets with QoS by setting the DSCP bit using the OSCARS' reservation capability. Using a mock-up of the DIII-D pulse cycle, EFIT calculation will be automatically triggered for between-pulse EFIT calculations using QoS. When DIII-D begins operations again in April 2006, this capability will be tested to support actual experimental operations.

In parallel with the network testing described above has been the infrastructure preparation to run the fusion code TRANSP between-pulses. TRANSP is used for time-dependent analysis and simulation of fusion plasmas. Between-pulse TRANSP was run for the first time taking on average 7 minutes by tailoring algorithms and input data calculations. Utilizing the secure FusionGrid TRANSP service, this time included sending the input and output MDSplus data between DIII-D and PPPL where TRANSP was run as well as the code execution time. More than half of the total time was spent in network transmission so a QoS network service has the potential to greatly reduce the total time. Since our solution with EFIT described above is very general, it focuses on the data transmission mechanism (MDSplus) and not the code, the demonstration with TRANSP should be relatively straightforward.

With the successful demonstration of EFIT and TRANSP as QoS data analysis services for experimental fusion research, work will shift to implementing this into the standard fabric of daily life at DIII-D including analysis reservation interfaces, documentation, and training of the scientific staff. Additionally, the software developed in this project at DIII-D will be transferred and made functional at the two other large operating fusion machines in the US, Alcator C-Mod and NSTX. Our progress will also be reported to the ITER Organization.