

Science DMZ Implemented at CU Boulder

The University of Colorado, Boulder campus was an early adopter of Science DMZ technologies. Their core network features an immediate split into a protected campus infrastructure (beyond a firewall), as well as a research network (RCNet) that delivers unprotected functionality directly to campus consumers. Figure 1 shows the basic breakdown of this network, along with the placement of measurement tools provided by perfSONAR.

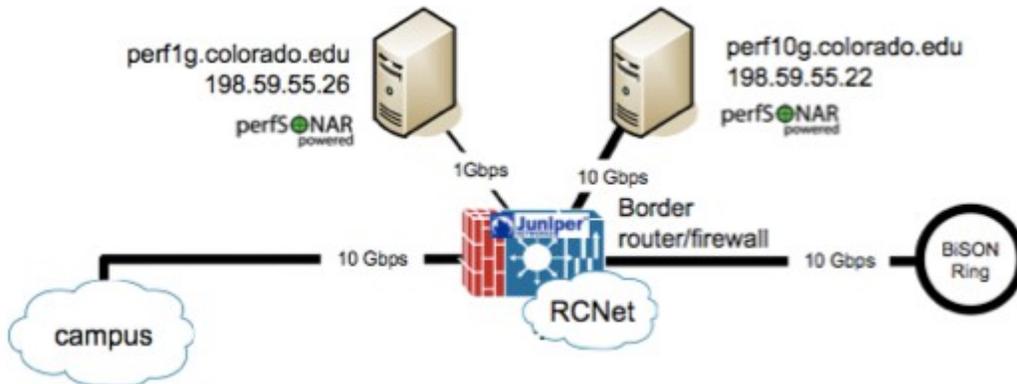


Figure 1: University of Colorado campus core network, showing the core split of functionality to support a Science DMZ.

The physics department, a participant in the Compact Muon Solenoid (CMS) experiment affiliated with the LHC project, is a heavy user of campus network resources. It is common to have multiple streams of traffic approaching an aggregate of 5 Gbps affiliated with this research group. As demand for resources increased, the physics group connected additional computation and storage to their local network. Figure 2 shows these additional 1 Gbps connections as they entered into the portion of the RCNet on campus.

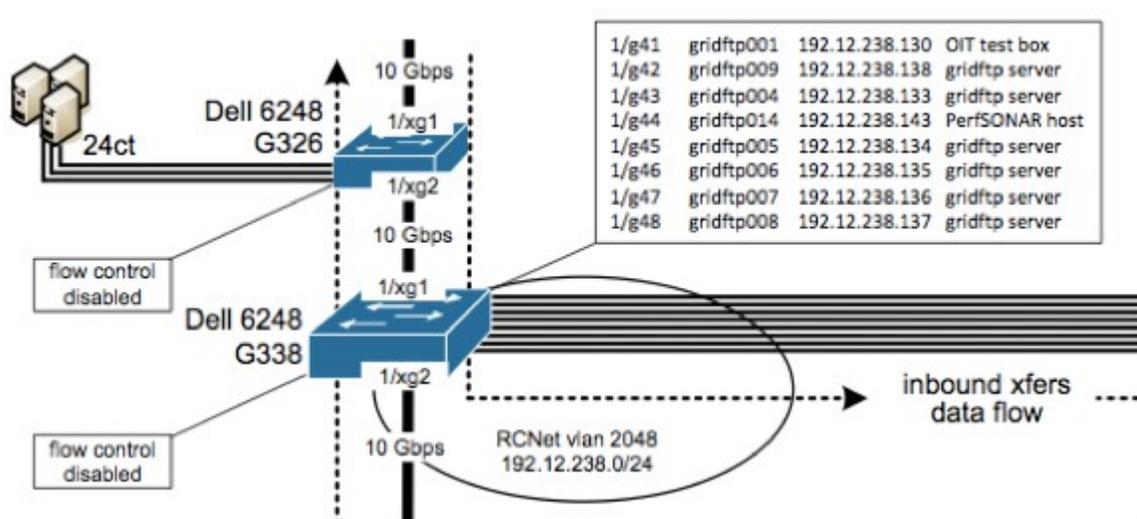


Figure 2: University of Colorado Network showing physics group connectivity.

Despite the initial care in design of the network, overall performance began to suffer during heavy use times on the campus. Passive and active perfSONAR monitoring alerted to low throughput to downstream facilities, as well as the presence of dropped packets on several network devices. Further investigation was able to correlate the dropped packets to three main factors:

1. Increased number of connected hosts,
2. Increased network demand per host,
3. Lack of tunable memory on certain network devices in the path.

Replacement hardware was installed to alleviate this bottleneck in the network, but the problem remained upon initial observation. After additional investigation by the vendor and performance engineers, it was revealed that the unique operating environment (e.g., high "fan-out" that featured multiple 1 Gbps connections feeding a single 10 Gbps connection) was contributing to the problem. Instead of operating in the default "cut-through" mode, the switch was forced to drop down into a "store-and-forward" mode where buffers limit capacity; this behavior reduced performance by delaying or dropping excess network traffic.

After a fix was implemented by the vendor and additional changes to the architecture were implemented, performance returned to near line rate for each member of the physics computation cluster.

Originally this page was a link to the *fasterdata* web page. When ESnet reorganized, the link was broken. I am reproducing the content rather than merely fixing the link. The new link as of June 2015 is:

<https://www.es.net/science-engagement/case-studies/science-dmz-case-studies/science-dmz-implemented-at-cu-boulder/>