Why 100Gbps Matters to the Internet2 Community

While many of the core IP technologies deployed on the Internet2 Network are in use on the commercial Internet today, the needs and capabilities of the user communities are different. Much of the end-to-end Internet2 infrastructure (national, regional, campus) provides >1Gbps flow capability to Internet2 users. As large-scale research projects come online, and as properly tuned 10G-connected PCs become the norm, the Internet2 infrastructure will continue to require greater capacities in the wide area.

A Network With Great Potential – It's time

Edge capacity between the Internet2 backbone and the Internet2 Connectors have largely kept pace with the fastest commercially-available interface speeds. A sampling of five Internet2 core nodes in mid 2001 showed a potential of **14Gbps** worth of edge-capacity into the 2.5Gbps Internet2 backbone. In mid 2009, that potential had risen to **248Gbps**- a seventeen-fold increase.[1] Yet, in many places, backbone capacity is constrained to 10 and 20Gbps between Internet2 core nodes. The current oversubscription rate (approx 25:1) is at the same level it was in 2001 when Internet2 began the migration to the 10Gbps network backbone.

An Explosion in Traffic

In many respects, this potential bandwidth capacity is being realized. Internet2 has measured a steady increase in total amount of traffic into the edges of its backbone network. In 1999, Internet2 typically transported **2 Terabytes** worth of data between its network edges. In 2009, a typical month sees **6.7 Petabytes** worth of data transferred[2]- a 3,300% increase.

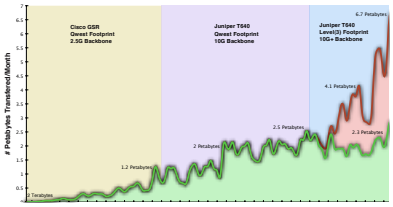

It's All About the Large Flows

In the last several years, Internet2 has grown its network capacity through the addition of parallel 10G backbone links between core nodes. This provides 20Gbps of overall capacity, though large flows risk being sub-optimally distributed amongst the parallel links. While true that >10G flows aren't possible in a parallel 10G solution, *the more pressing concern is a large number of >1G<10G flows.* As the individual flows are hashed across the parallel links, there is a risk that additional large flows end up being distributed to the same backbone link- causing one flow to receive less bandwidth before the aggregate 20G potential is used. Internet2 supports several applications that have individual flow potential in the 5-6Gbps range. With LHC-related applications regularly pushing 4.8Gbps[3], the risk of sub-optimal hashing of the individual flows is high in a parallel-link architecture.

A Cost Efficient Delivery Platform

As the Internet2 service portfolio diversifies into both IP and circuit-delivery services, *growing separate parallel 10G links in a separate circuit and IP network fails to deliver capabilities required for cutting edge applications and is less efficient financially than moving to a larger single 100G network*. A combined IP and Circuit delivery infrastructure, operating at 100G, clearly provides an easier and more cost-efficient growth curve for both services without the complexity of apportioning resources toward one service over the other.

---

[1] See http://www.internet2.edu/presentations/spring09/20090428-network-robb.pdf, slide 12 for an *average* of edge interface speeds

[2] See http://www.internet2.edu/presentations/spring09/20090428-network-robb.pdf, slide 10

[3] See http://www.usatlas.bnl.gov/dq2/throughput