OptiWay Hybrid Routing

Problem Statements

- The cost of LAN infrastructure is determined by the peak bandwidth the network must support
- LAN cost goes up non-linearly with bandwidth
- A few power users require the majority of the bandwidth in a network
- Financial constraints force a sub-optimal compromise
  - *For most users you pay for more bandwidth than you need*
  - *For a few users you provide nowhere near their needed bandwidth*

Typical LAN environments consist of hundreds of computers and servers networked together using IP over Ethernet. Copper twisted pair along with 100BT Ethernet switches form the physical network. Most users on the network demand little; however, typically a few dozen power users require tremendous network bandwidth compared to other users. These power users may require a lot of bandwidth because they are doing intensive computation, interacting with high bandwidth video, or are accessing a lot of data from storage/servers. Power users do not have to be "people", they can be storage devices, servers or systems, such as radar or sonar systems. The standard 100BT physical network is insufficient for these users. In addition, the normal users are adversely impacted by the network needs of the power users, because network throughput goes down dramatically when the power users demand network resources.

The Glimmerglass Hybrid Solution

What is necessary is a way of overlaying express paths through the network for handling high bandwidth transactions. A switched optical overlay network, shown in Figure 1, solves this problem in a cost effective manner. A Glimmerglass System 300 is at the center of the overlay network. It interconnects optical Ethernet NICs installed into each power user's computer and also into appropriate servers. These optical NICs are in addition to the existing copper 100BaseT NICs. When a power user needs a high bandwidth connection (e.g., to download a large file from a server), the System 300 forms a transparent connection between the computer and the appropriate server. The Ethernet and IP layers automatically make use of the additional bandwidth. Network traffic is thus diverted from the 100BT network, allowing the normal users to continue work without a reduction in network throughput.

Advantages

- Incremental. Only power users need to be upgraded as they are identified. All existing infrastructure (100BT copper plant and switches) are preserved.
- Low-cost. Upgrading entire network to fiber is very expensive and unnecessary. In particular, the high cost of optical L2 Ethernet switches is avoided.
- Transparent. System 300 provides transparent connections. Switch doesn't add any latency or drop packets, unlike L2 switches.

- Future-proof. System 300 ports can equally carry 40Gb as 10Gb connections. Power user PCs can be upgraded to 40Gb without any additional investment in the switch.

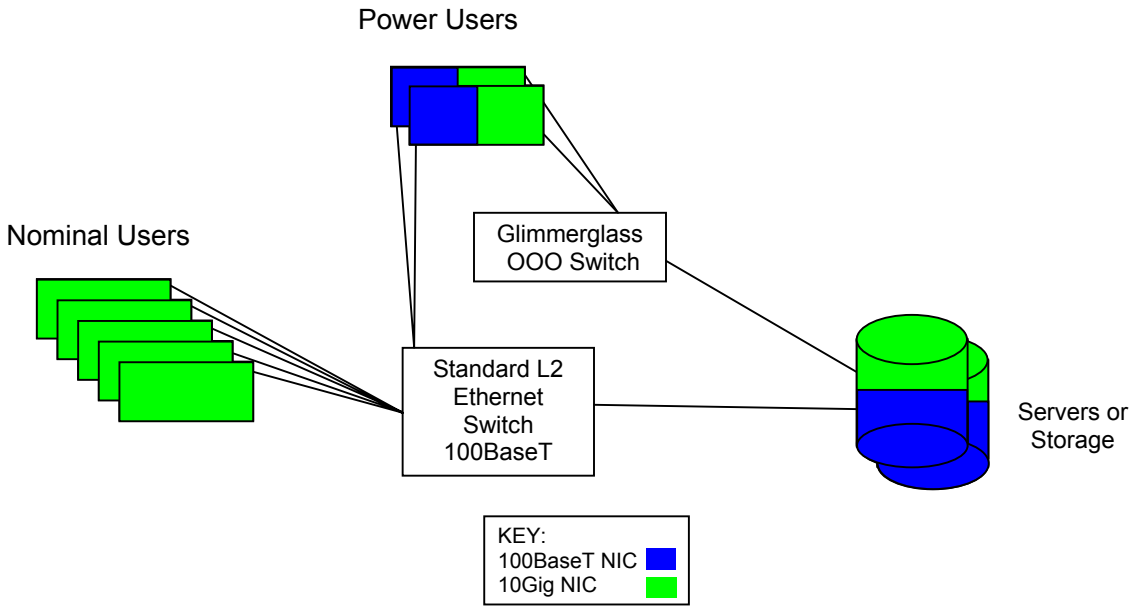

Figure 1: Glimmerglass Hybrid Network Configuration

Details

The optical network is established as a separate IP subnet (e.g., different network address) from the existing copper network. The routing metrics are configured to favor the optical network when up. All common commercial OSs support such configuration "out of the box".

The System 300 is instructed to establish connections either manually (using a GUI) or automatically via Grid-enabled applications (e.g., GridFTP).

Alternatives

1) Upgrade entire network to optical. This is very expensive, consumes a lot of weight, space and power, and most users will see no benefits. A comparison of a Glimmerglass Hybrid optical network and a 10 Gig electrical network is given in Figure 2.
2) Interconnect power users's optical connections with L2 switch.
   a) Very expensive compared to L1 switch. Paying for packet-by-packet non-blocking connectivity when not needed.
   b) Not transparent: packet loss, delay in L2 switch degrades performance.

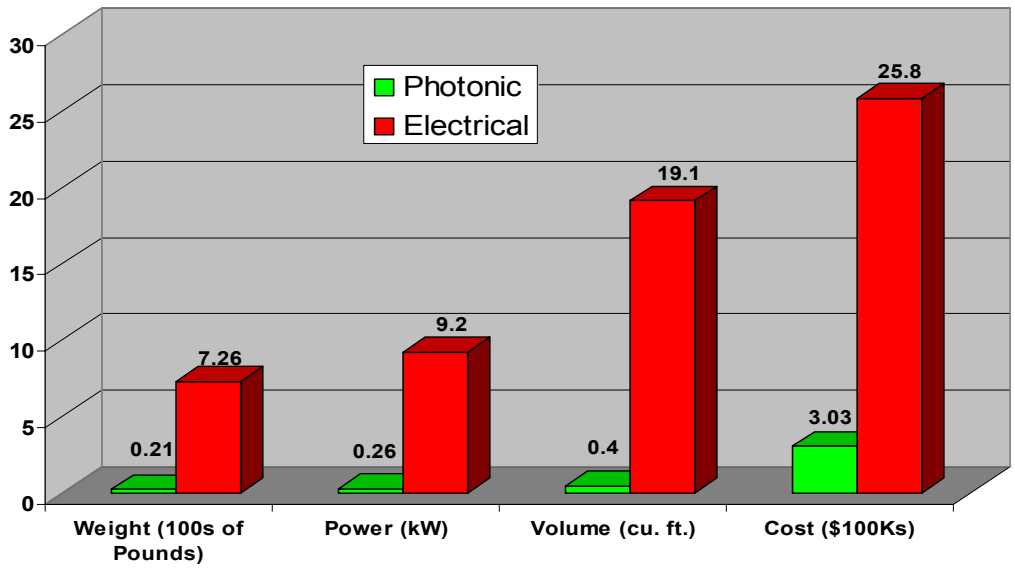

Figure 2: Comparison of Glimmerglass Hybrid Photonic network with 10 Gigabit electrical network. Total network size = 96 nodes. Number of power users/servers = 32 nodes.

Conclusions

A photonic hybrid network provides increased throughput and scalability both in number of power user ports and in bandwidth. In addition, the weight, power, volume and cost are dramatically reduced compared to an electrical network.