

Monitoring and Accelerating GridFTP

Martin Swany, Indiana U.

Dan Gunter, LBL

Jason Zurawski, Internet2

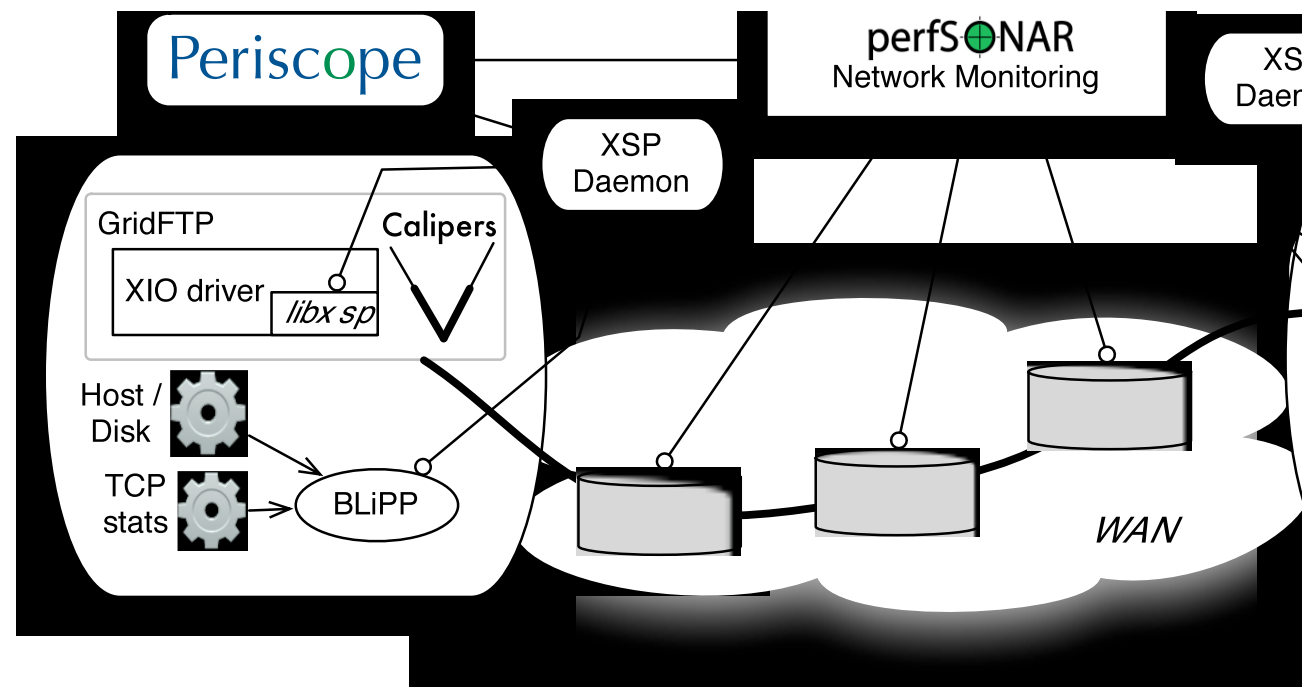
XSP - eXtensible Session Protocol

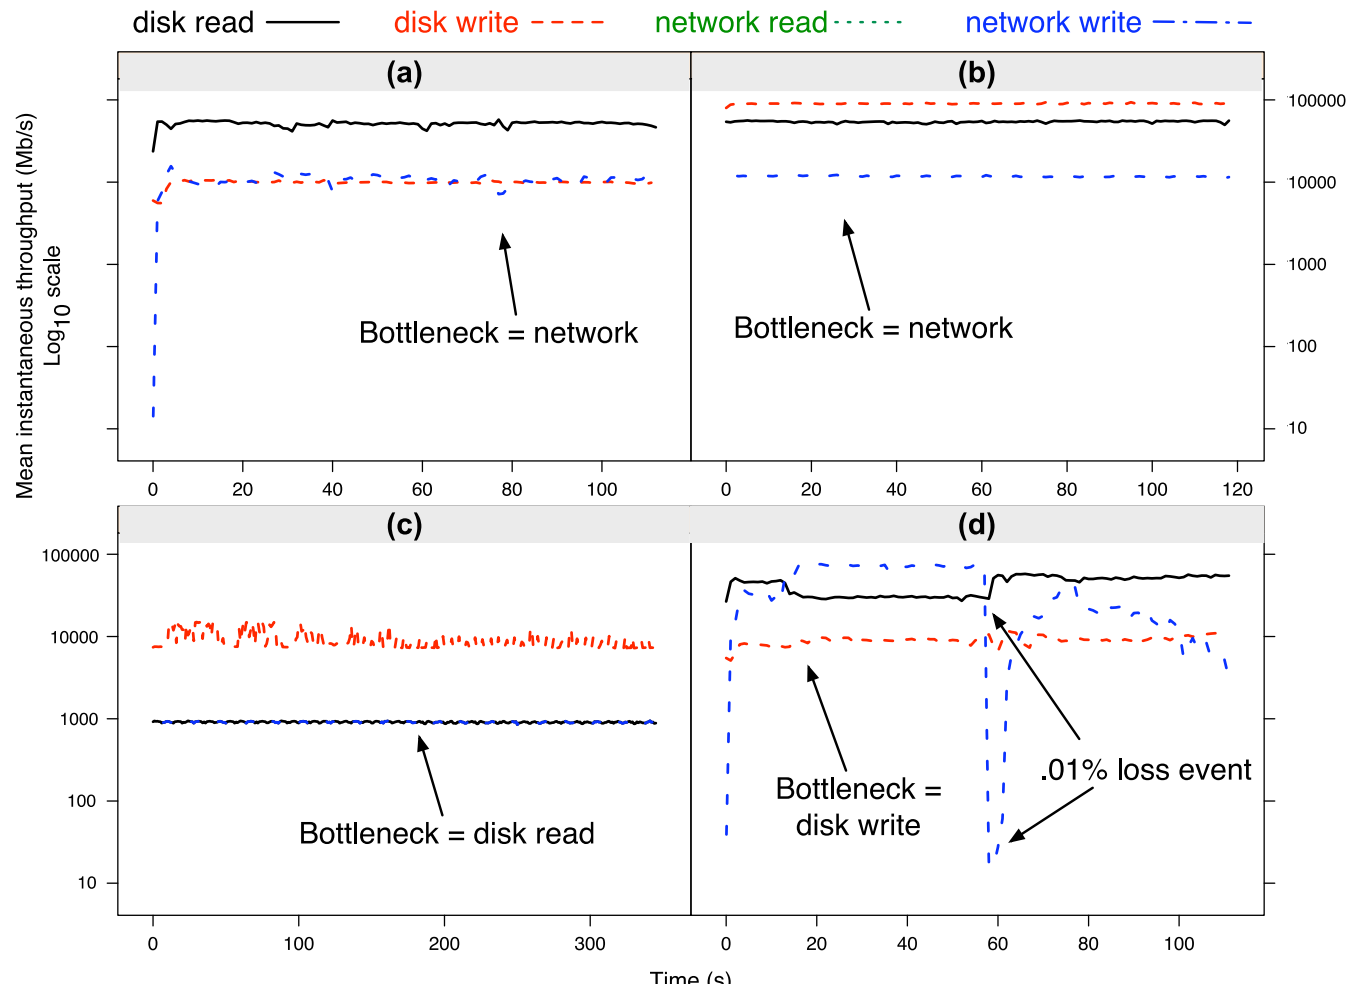
- Session layer (Layer 5 in the OSI model)
- Implemented as a Globus XIO driver
 - Loadable with `-dcstack` and `-fsstack`
- Use cases:
 - Monitoring with NetLogger, Calipers and Periscope
 - Dynamic network provisioning
 - WAN Acceleration with Phoebus

Monitoring

- NetLogger monitors read and write system calls, and Calipers summarizes these in situ
- The XSP collector daemon collates and forwards to Periscope
 - Periscope is a perfSONAR front end and cache

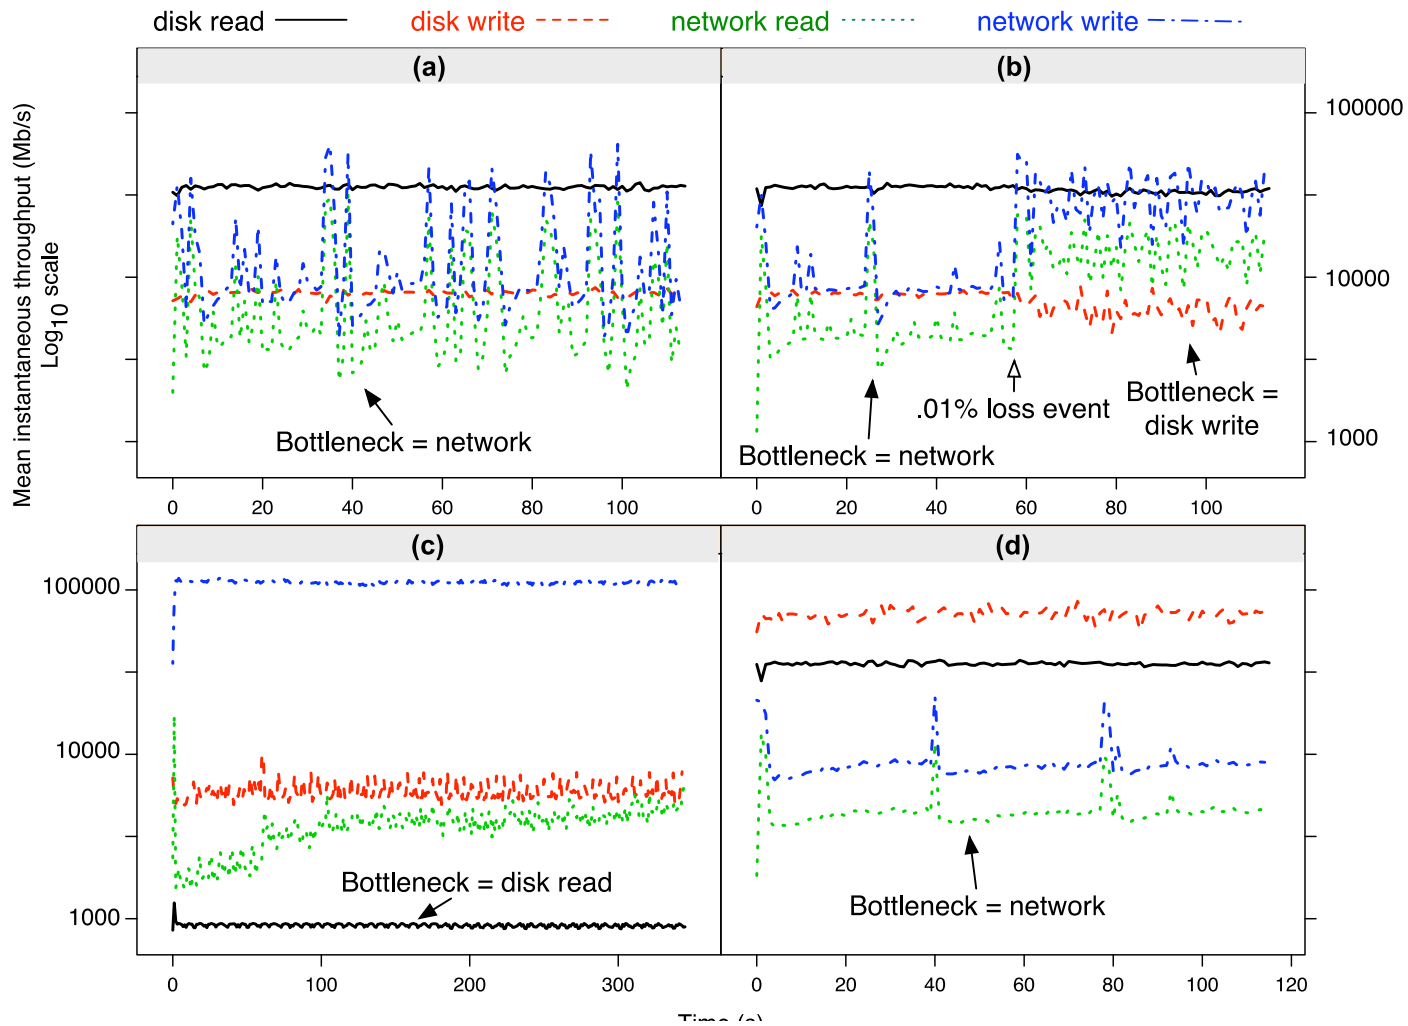
- BLiPP:
 - Basic
 - Lightweight
 - Periscope
 - Probes





TCP throughput

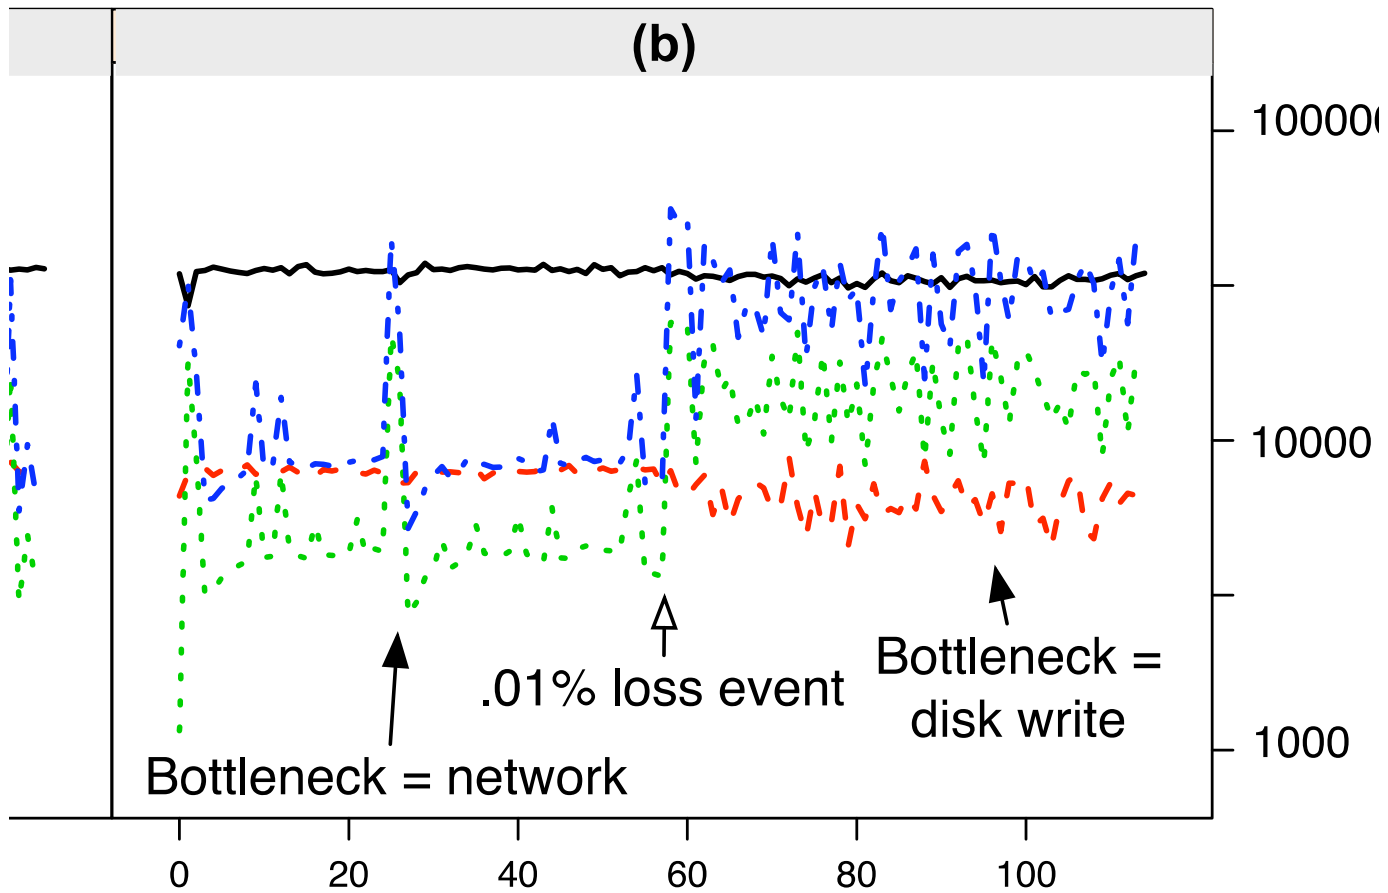
Time series of throughput for representative TCP experiments: (a) 1 stream memory-to-disk with 100ms latency, (b) 1 stream memory-to-memory with no latency, (c) 1 stream disk-to-disk with no latency, (d) 4 streams memory-to-disk with 100ms latency and 1% loss added at 60 seconds.





UDT throughput

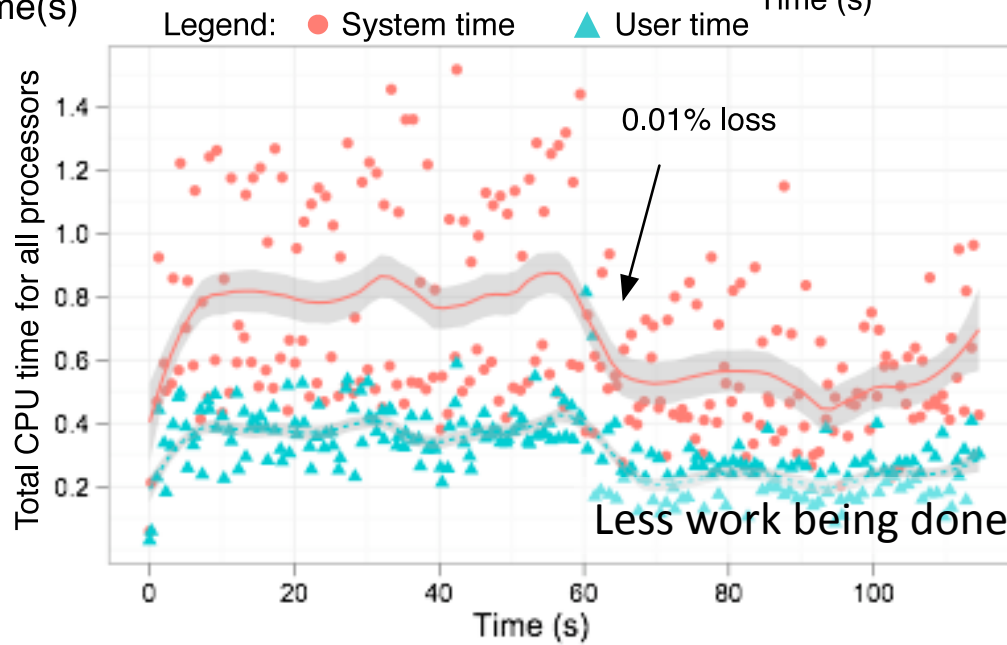
Time series of throughput for representative UDT experiments: (a) 4 streams memory-to-disk with 100ms latency, (b) 4 streams memory-to-disk with 100ms latency and 1% loss added at 60 seconds, (c) 4 streams disk-to-disk with 100ms latency, (d) 4 streams memory-to-memory with 100ms latency.

Wait, what?

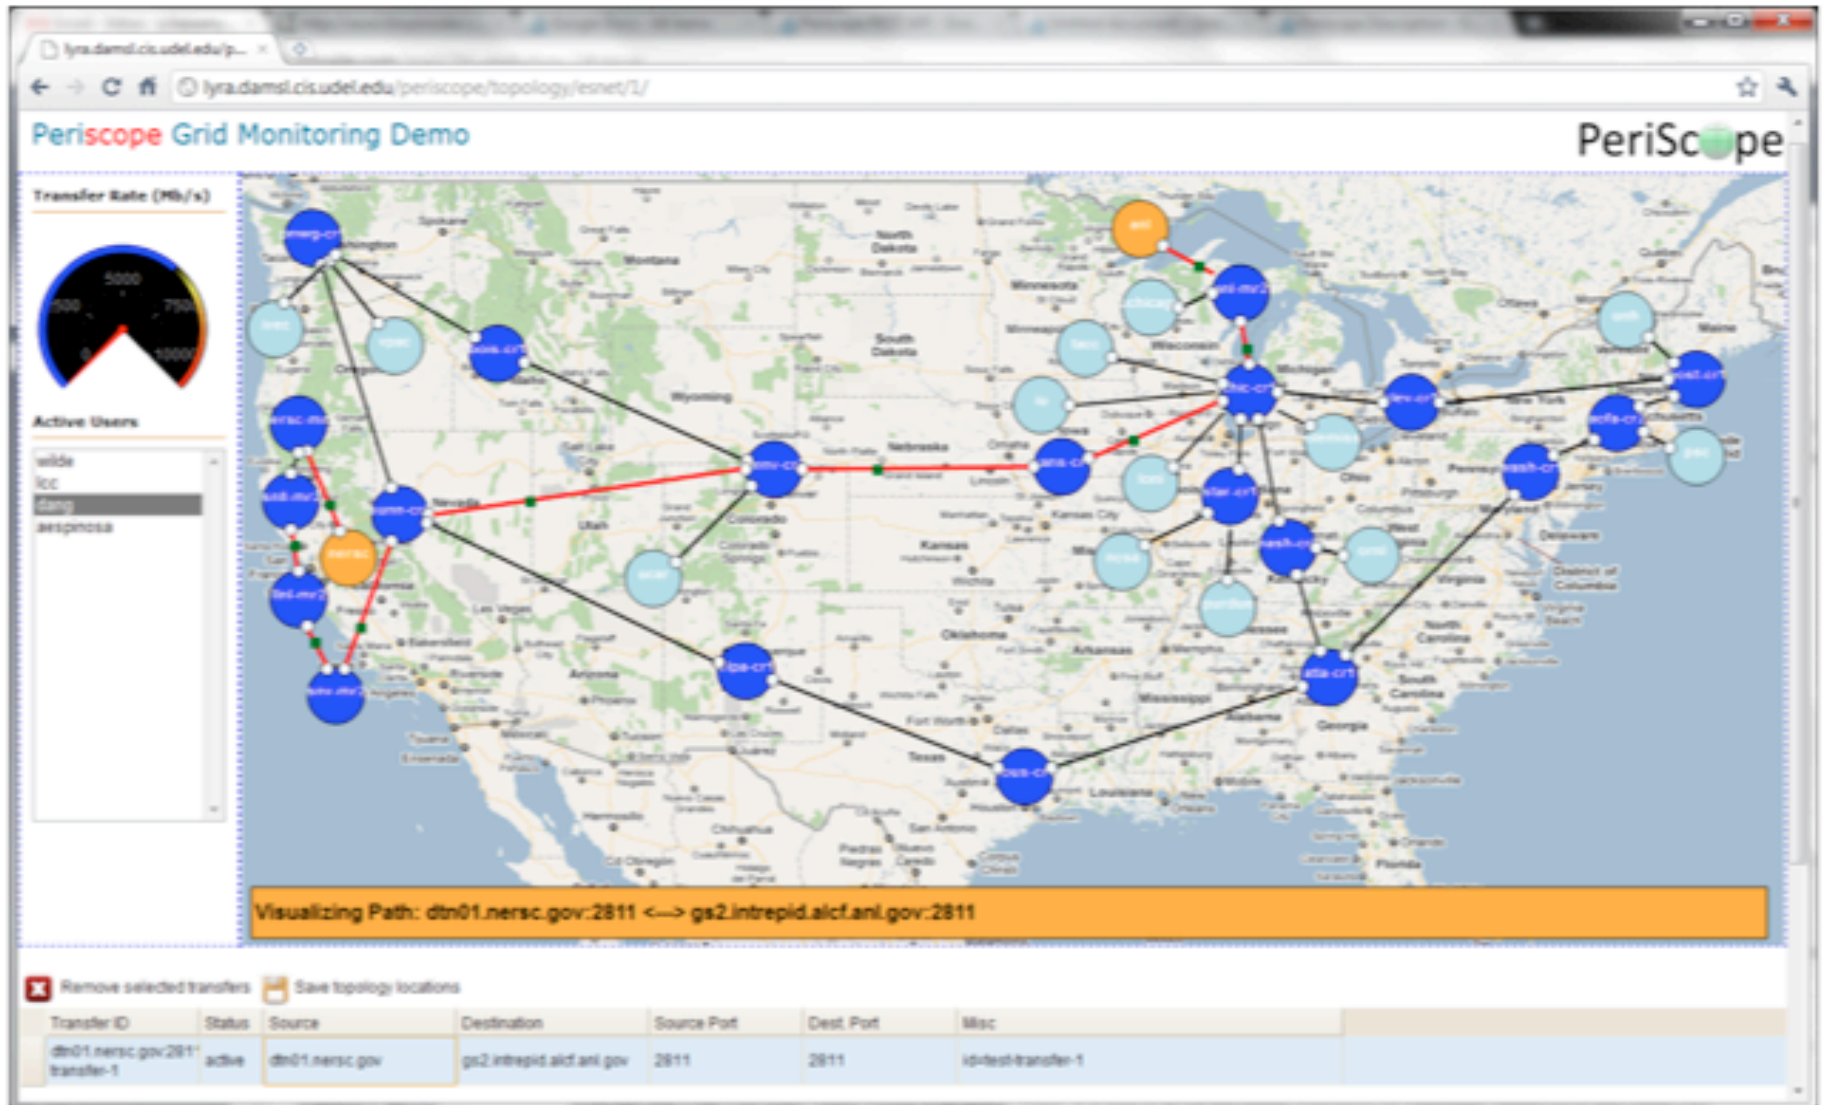




Half as many read(s).
Others return zero,
not counted




Topology-aware Monitoring



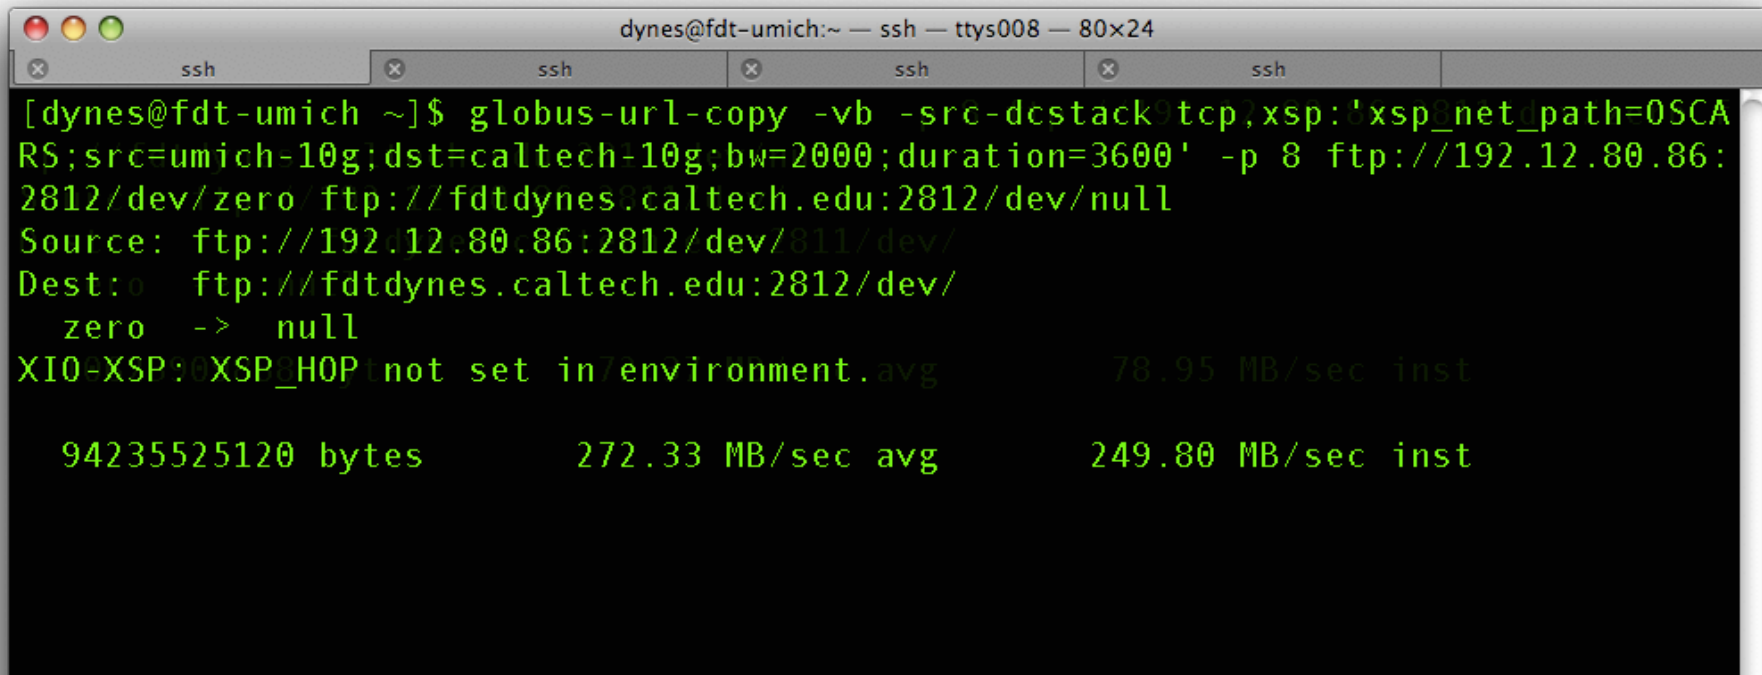
XSP and Dynamic Circuits

- All things old are new again – circuits
 - Still packets over traffic engineered paths
- ESnet OSCARS, Internet2 ION and OS³E/NDDI
 - NDDI is based on OpenFlow, an emerging protocol to allow fine-grained control of network devices
- SC11 Demo: monitor and react



GridFTP+XSP using the Phoebus WAN Accelerator System

- Phoebus is an open source WAN accelerator funded by the DOE and now the NSF
- Phoebus uses XSP to communicate via gateways that can tune, adapt and translate protocols
 - TCP tuning, UDT, RDMA over the WAN



```
dynes@fdt-umich:~ — ssh — ttys008 — 80x24
ssh ssh ssh ssh
[dynes@fdt-umich ~]$ globus-url-copy -vb -src-dstack@tcp,xsp:'xsp_net_path=OSCAR;src=umich-10g;dst=caltech-10g;bw=2000;duration=3600' -p 8 ftp://192.12.80.86:2812/dev/zero ftp://fdtdynes.caltech.edu:2812/dev/null
Source: ftp://192.12.80.86:2812/dev/811/dev/
Dest:  ftp://fdtdynes.caltech.edu:2812/dev/
      zero -> null
XIO-XSP:@XSP_HOP|not set in/environment.avg          78.95 MB/sec inst
94235525120 bytes          272.33 MB/sec avg          249.80 MB/sec inst
```

Conclusion

- Quite a few topics focused on the performance of GridFTP
- Flexible and scalable monitoring for troubleshooting
- Adapt performance using emerging network technologies and protocols
- Questions?