Towards Bridging the Gap between Peak and Average Loads on Science Networks

Raj Kettimuthu
Argonne National Laboratory & University of Chicago

Sam Nickolay, University of Chicago
Eun-Sung Jung, Hongik University
Ian Foster, Argonne and University of Chicago

Data Deluge

Cosmology

Light Source Facilities

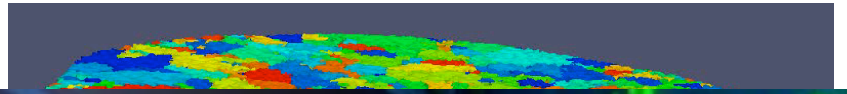

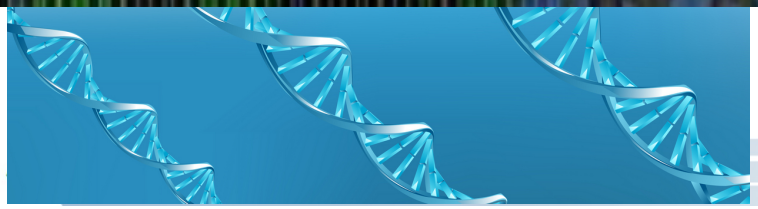

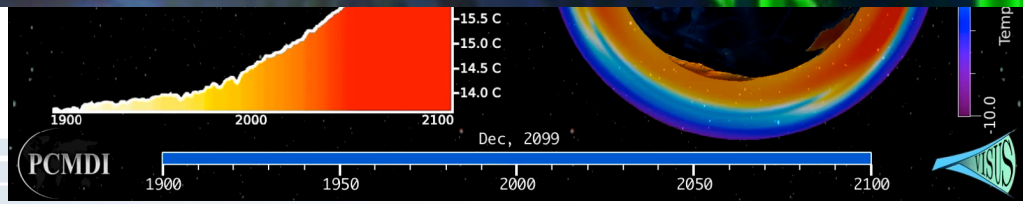

Science Workflows

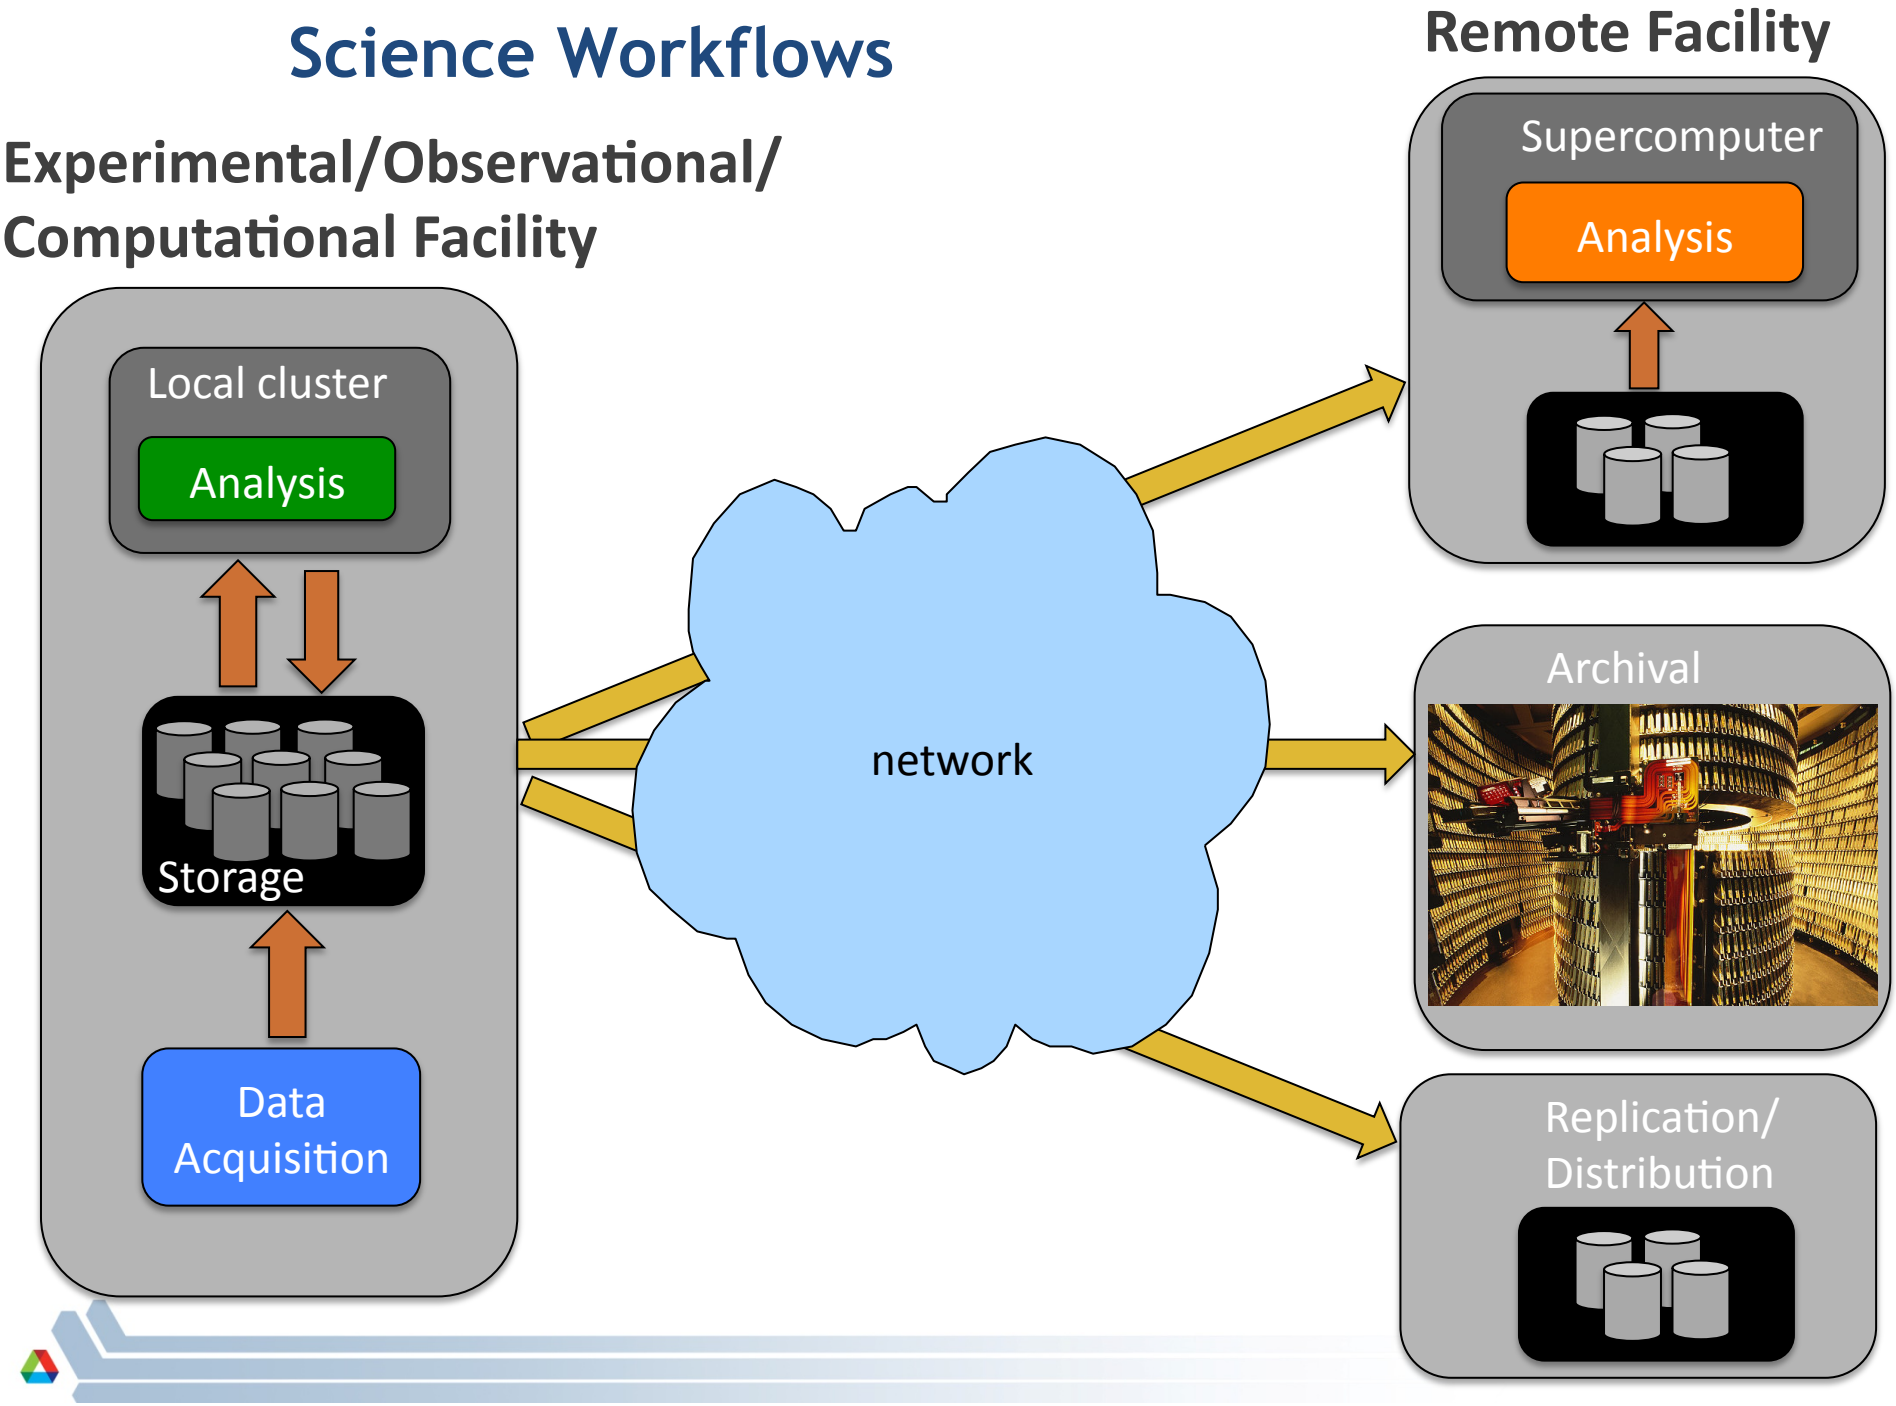

On-demand Transfers

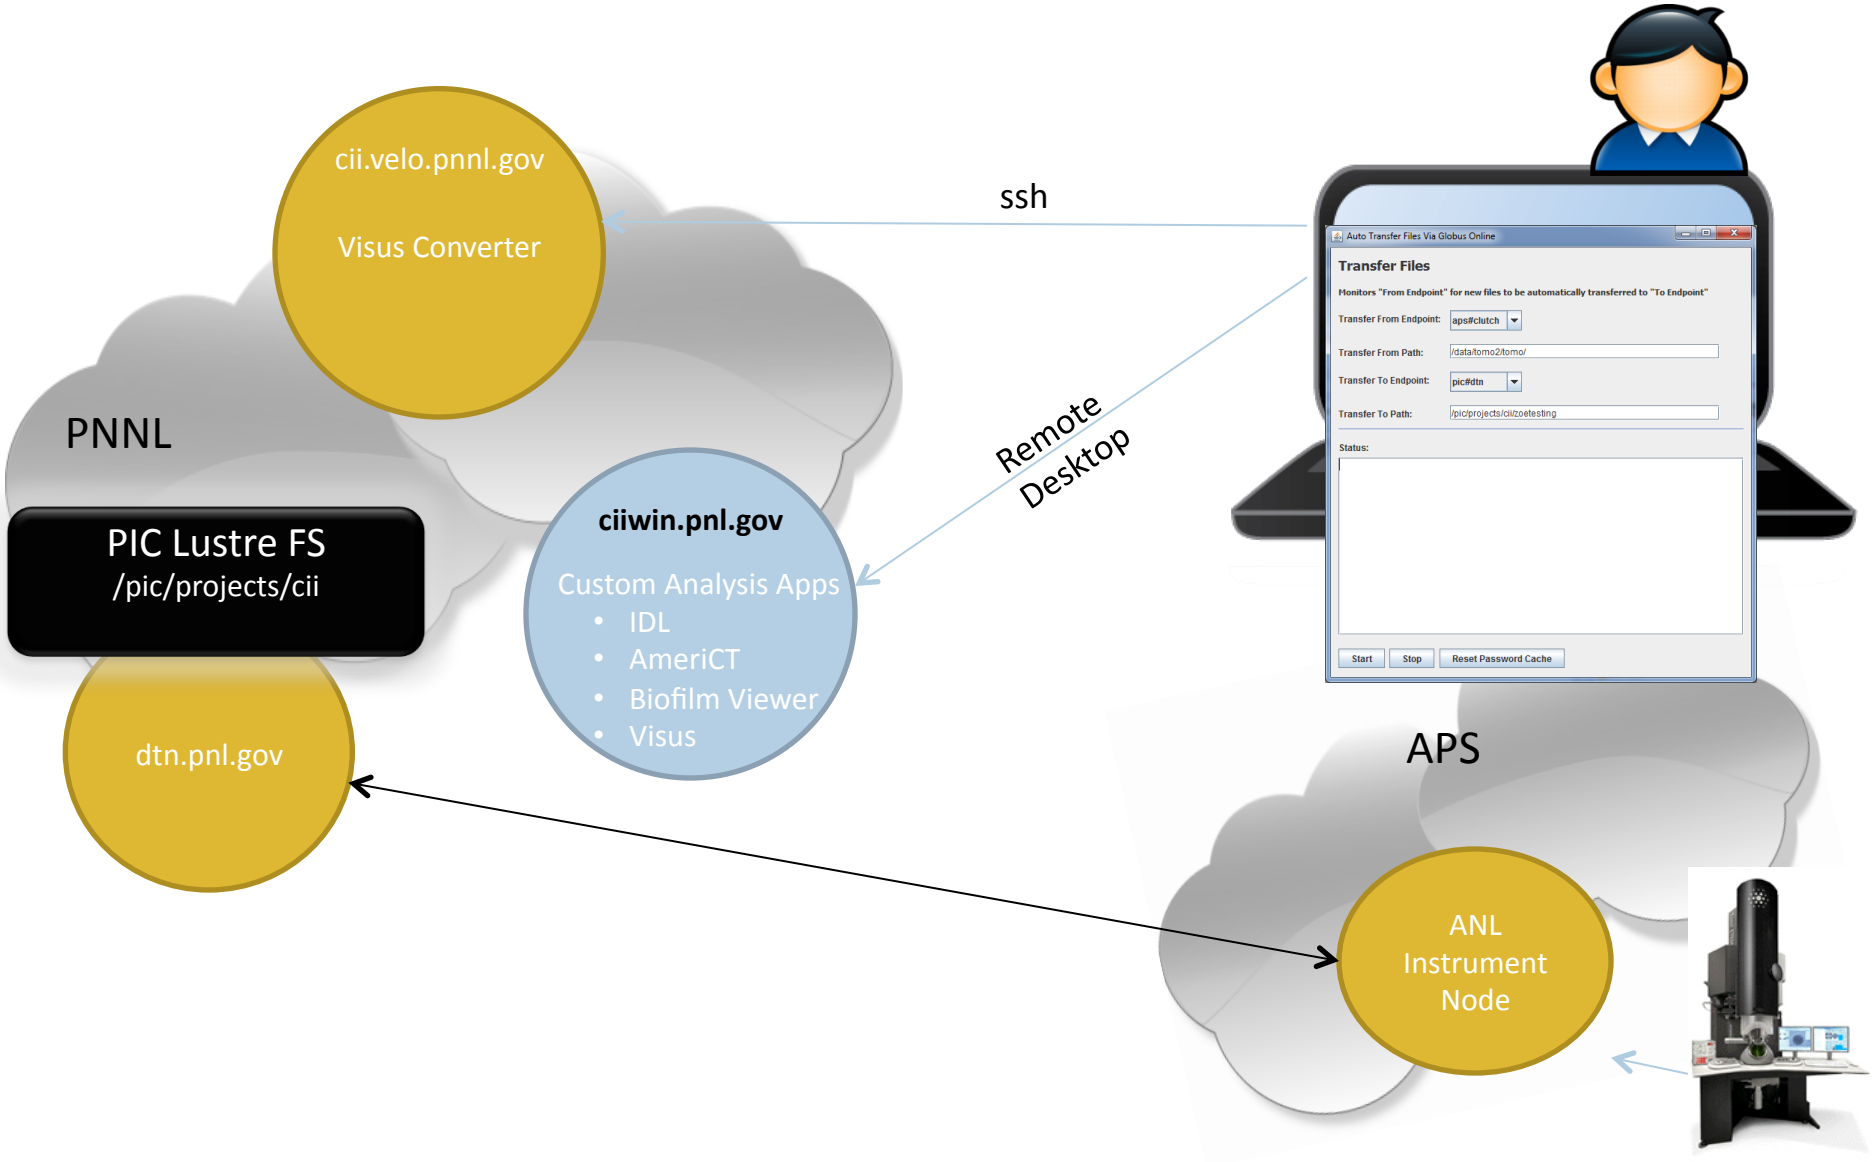

Source: my.es.net

Sat 20 Aug 2016 - Mon 19 Sep 2016

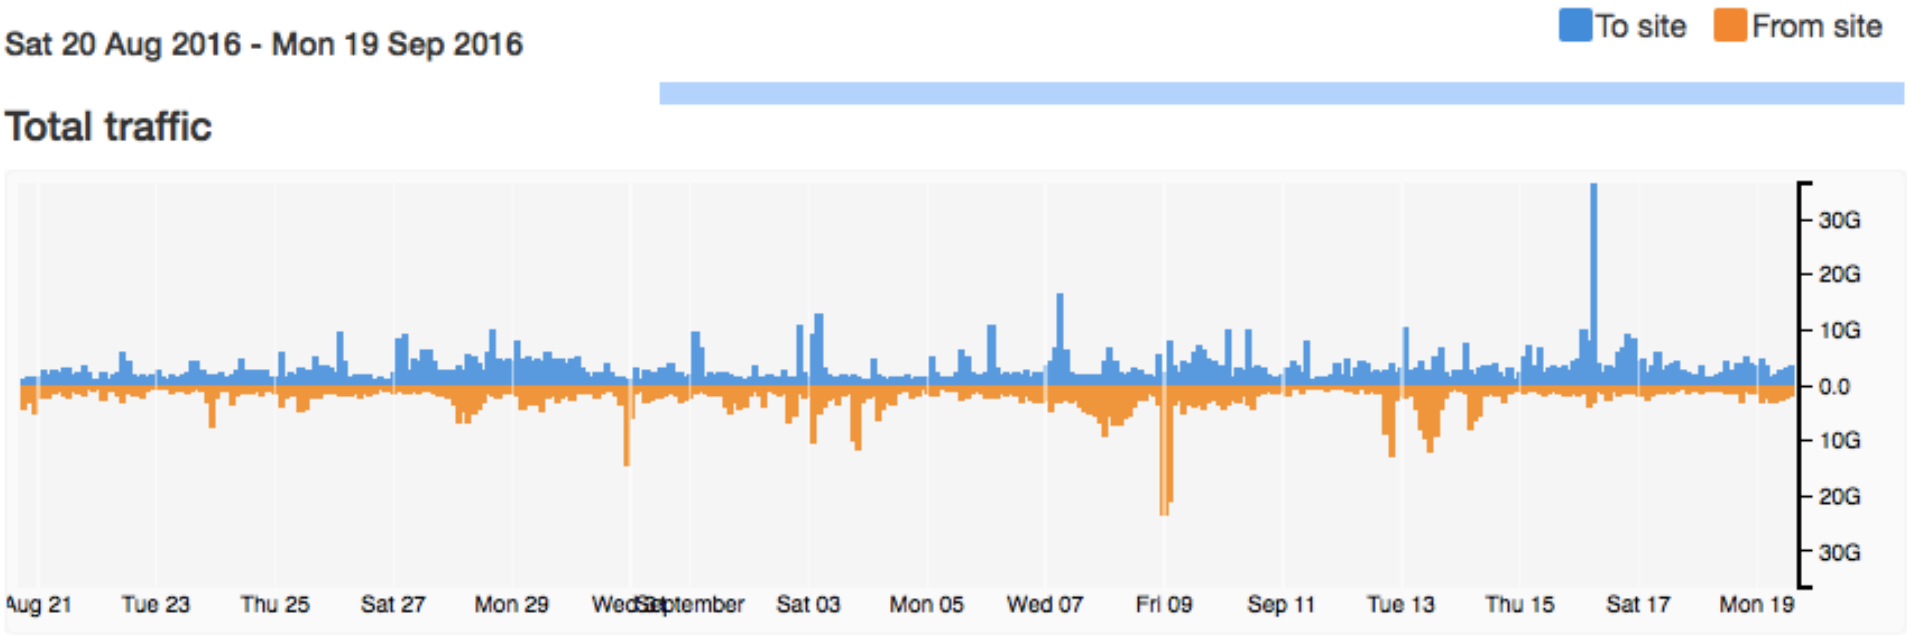

Sat 20 Aug 2016 - Mon 19 Sep 2016

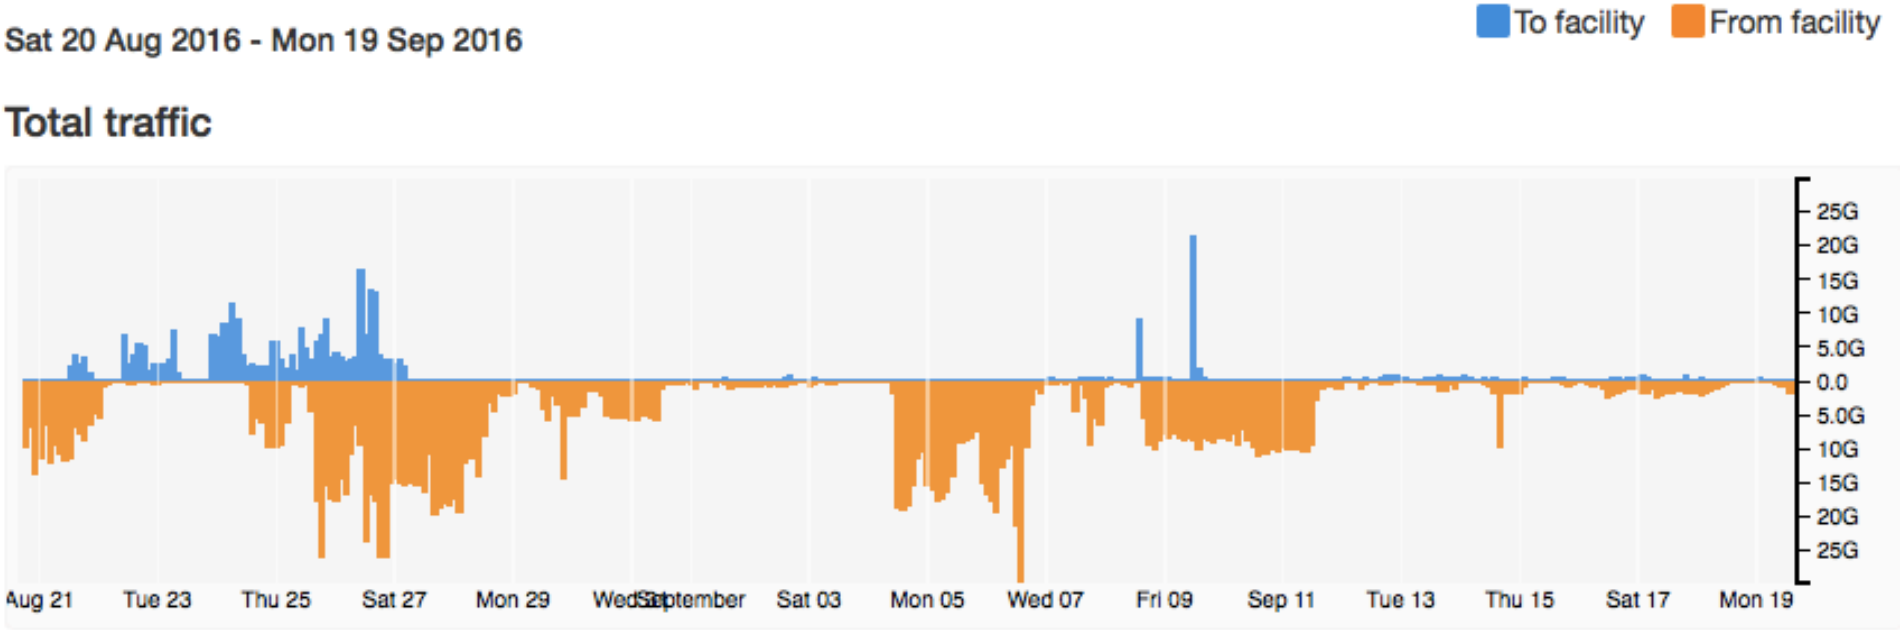

GridFTP Usage Data for Top Servers

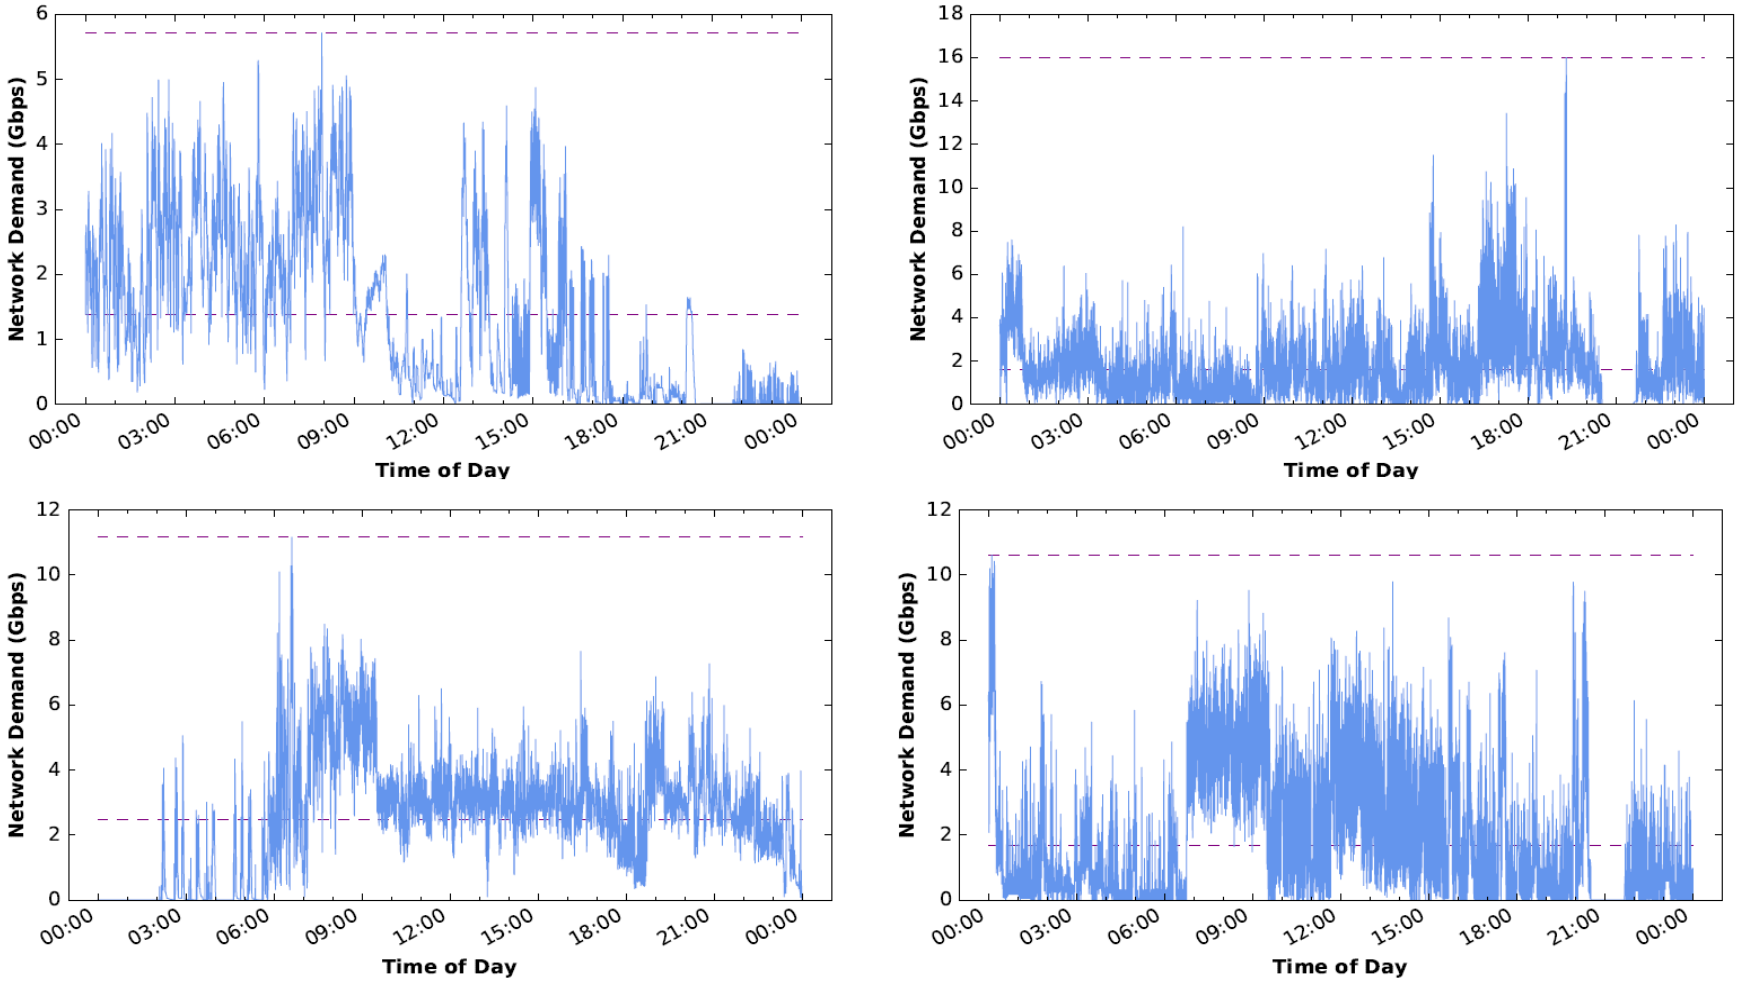

Figure 2: Network Demand for $trace_1$, $trace_2$, $trace_3$, $trace_4$.

Simulation Study

- Keeping the capacity of (overprovisioned) network constant, study the impact of increasing load on transfers
- Used 4 trace logs from Globus GridFTP servers
  - Varying peak (5.7Gbps - 16.0Gbps) and mean throughput (1.4Gbps - 2.5Gbps)
  - All have a mean throughput between 10% and 25% of the peak
- Simulate transfers in the logs in online fashion
  - 24-hour logs and scheduling interval in simulation is 1 second
- Group transfers into on-demand and best-effort
- Use slowdown metric

$$BS = \frac{\text{Waittime} + max(\text{Simulation Runtime}, \text{Bound})}{max(\text{Log Runtime}, \text{Bound})}$$

Simulation Variables

- % OD transfers – $\{10\%, 30\%, 50\%, 70\%\}$
- % OD bandwidth – absolute values depended on % OD transfers
  - $\{\%\ OD\ tranfers,\ \%\ OD\ transfers+10\%,\ \%\ OD\ transfers +20\%,\ 100\%\}$
- Transfer Load Ratio – $\{1.0x, 1.5x, 2.0x, 2.5x\}$
  - Although ratios are the same for all trace logs, the resulting loads are different
- 64 different configurations for each trace log
- Baseline control Experiments – 100% BE transfers, 0% OD transfers
  - Used to compare the relative performance of our scheduling algorithm

Average Slowdown for Trace 1

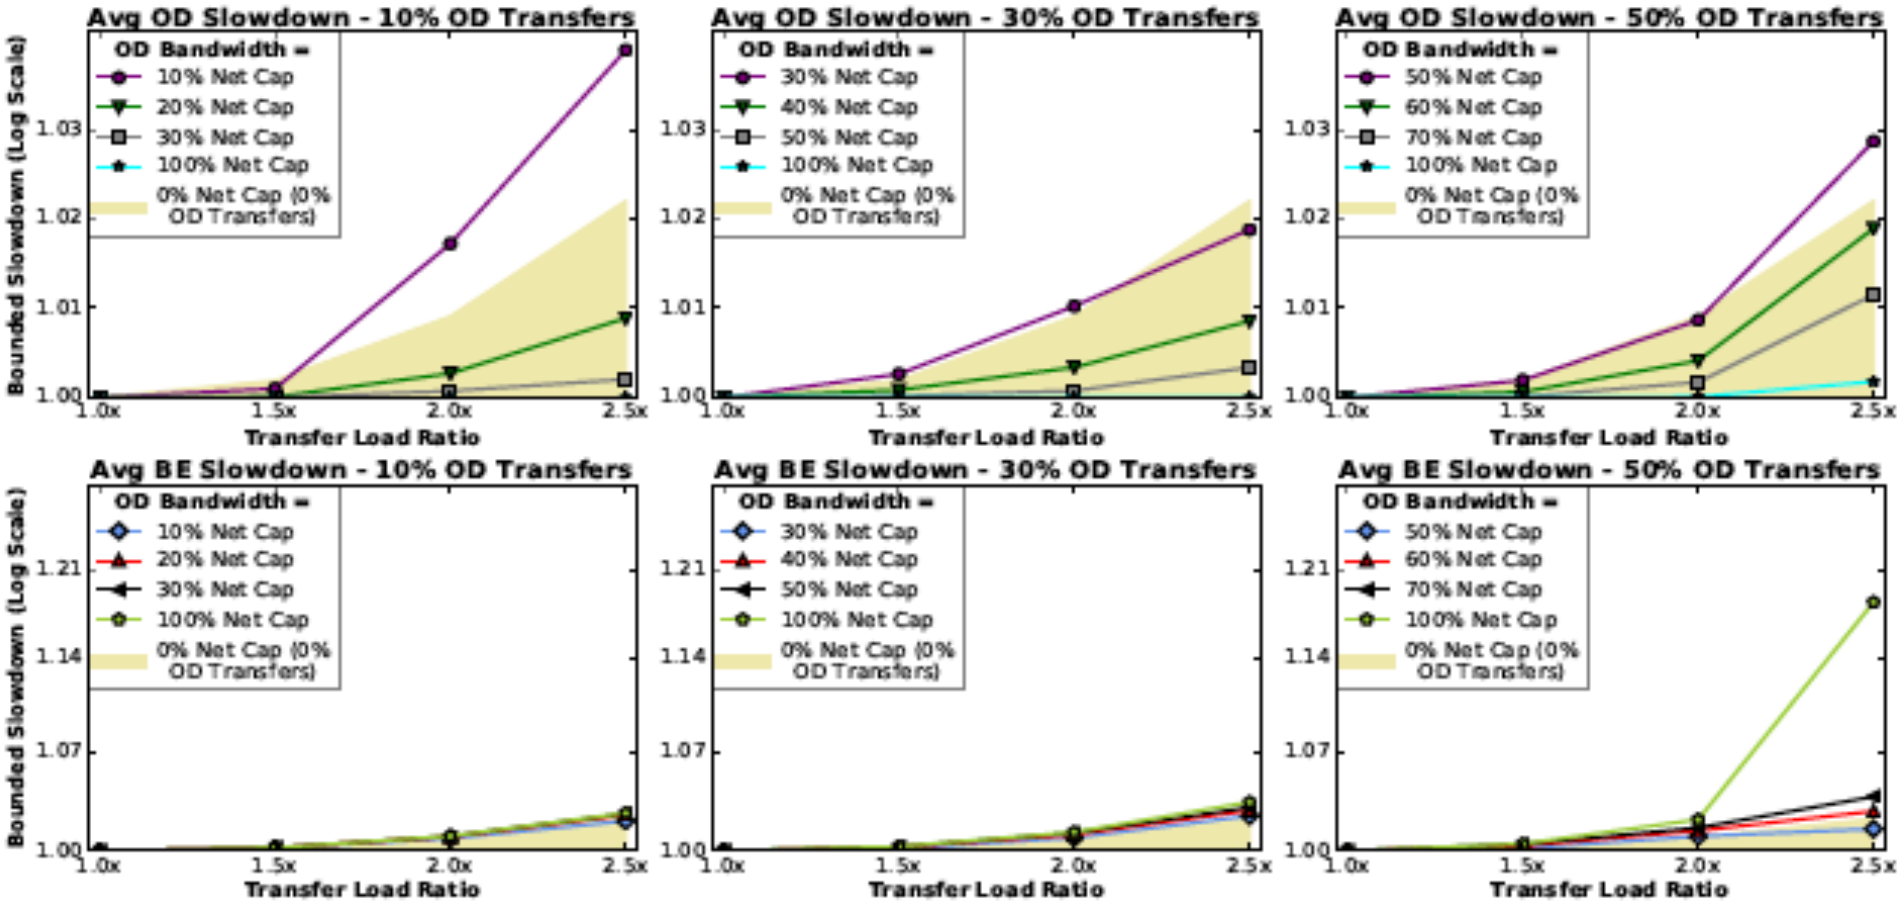

Max Slowdown for Trace 1

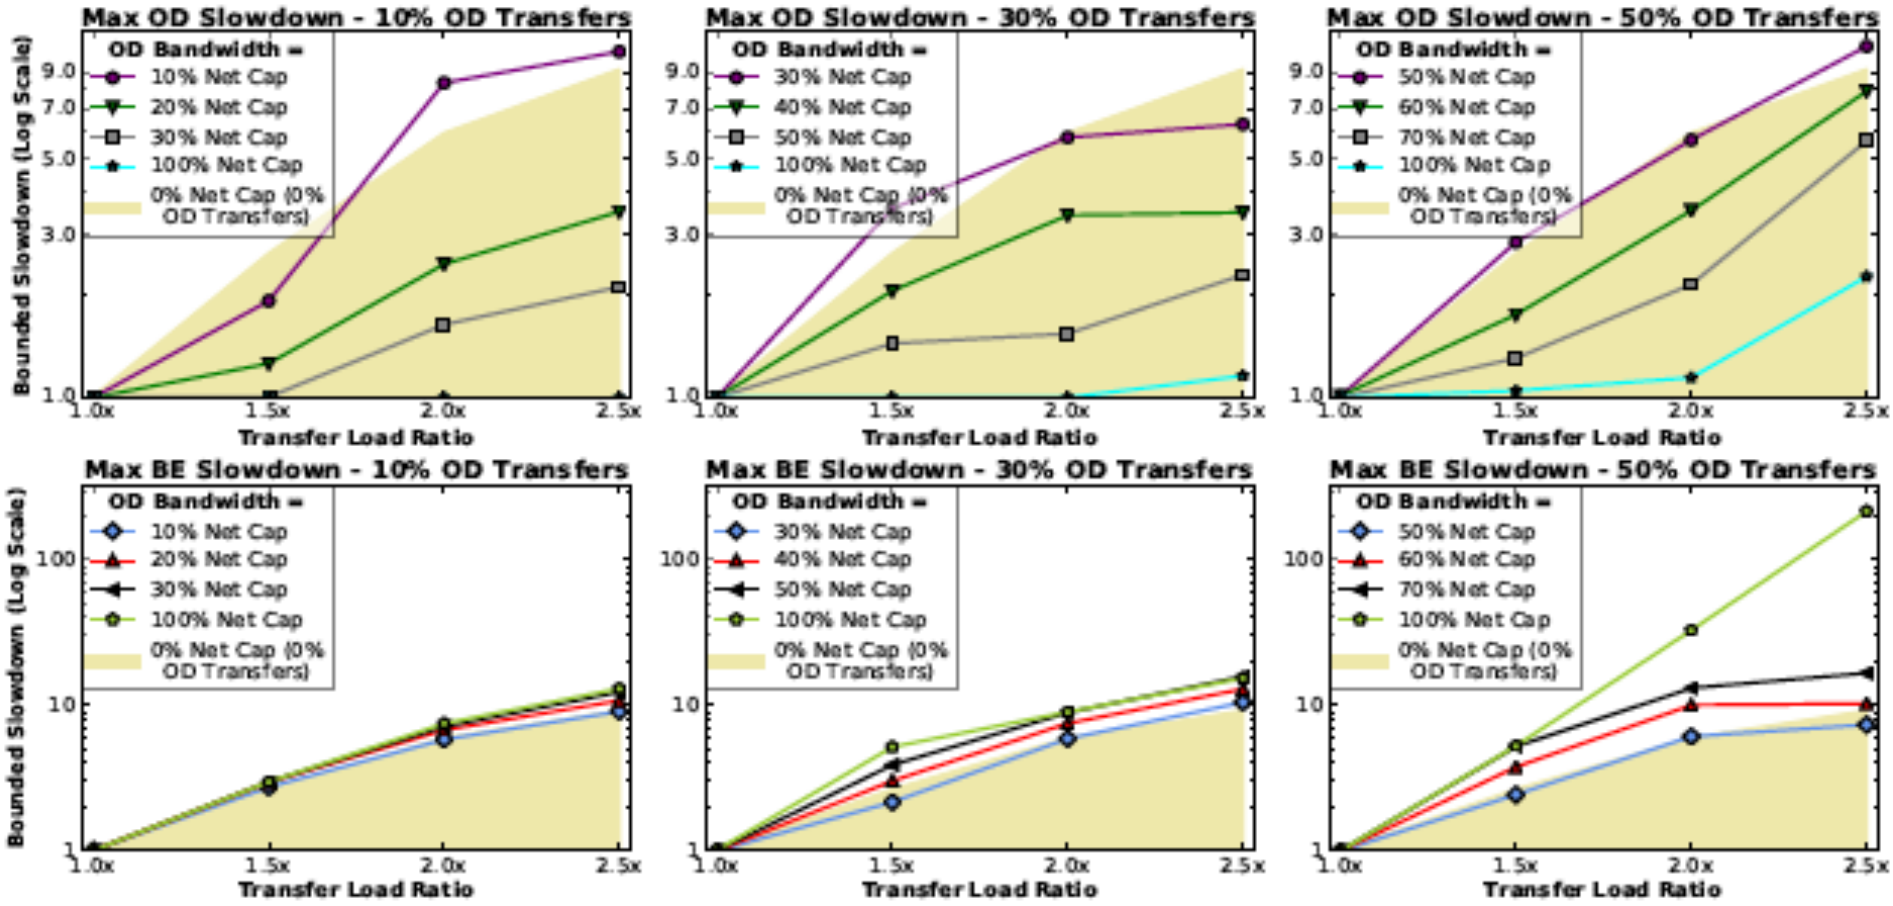

Average Slowdown for Trace 2

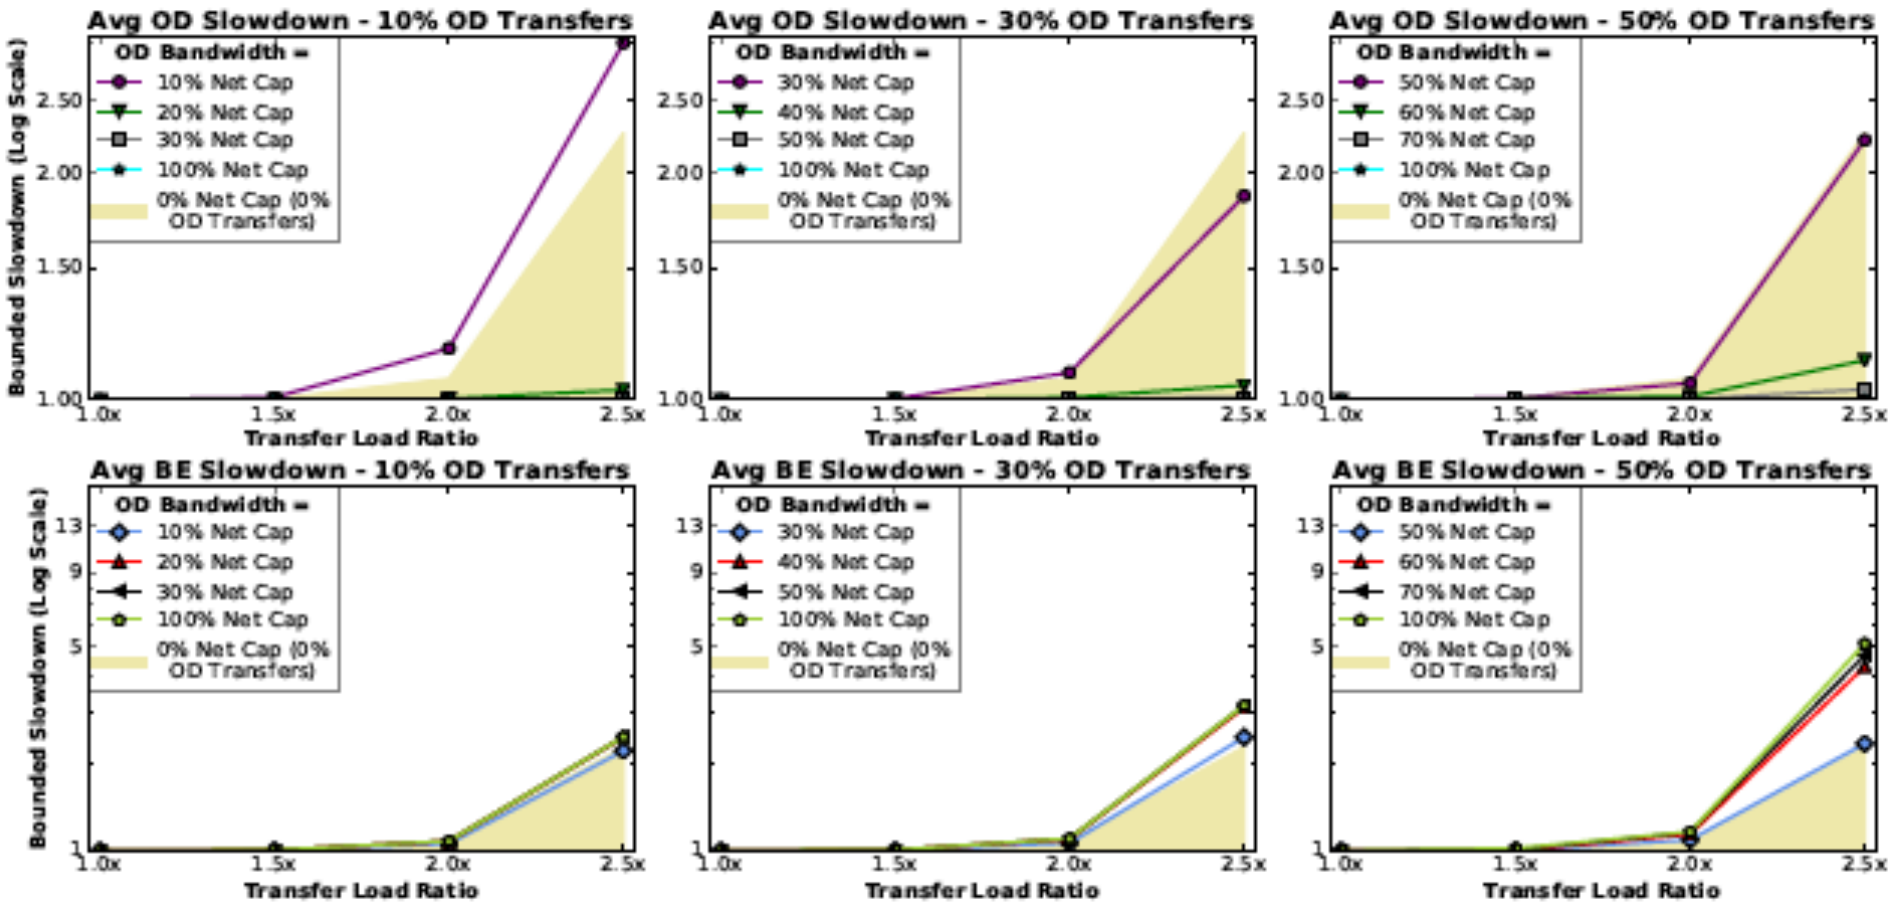

Average Slowdown for Trace 3

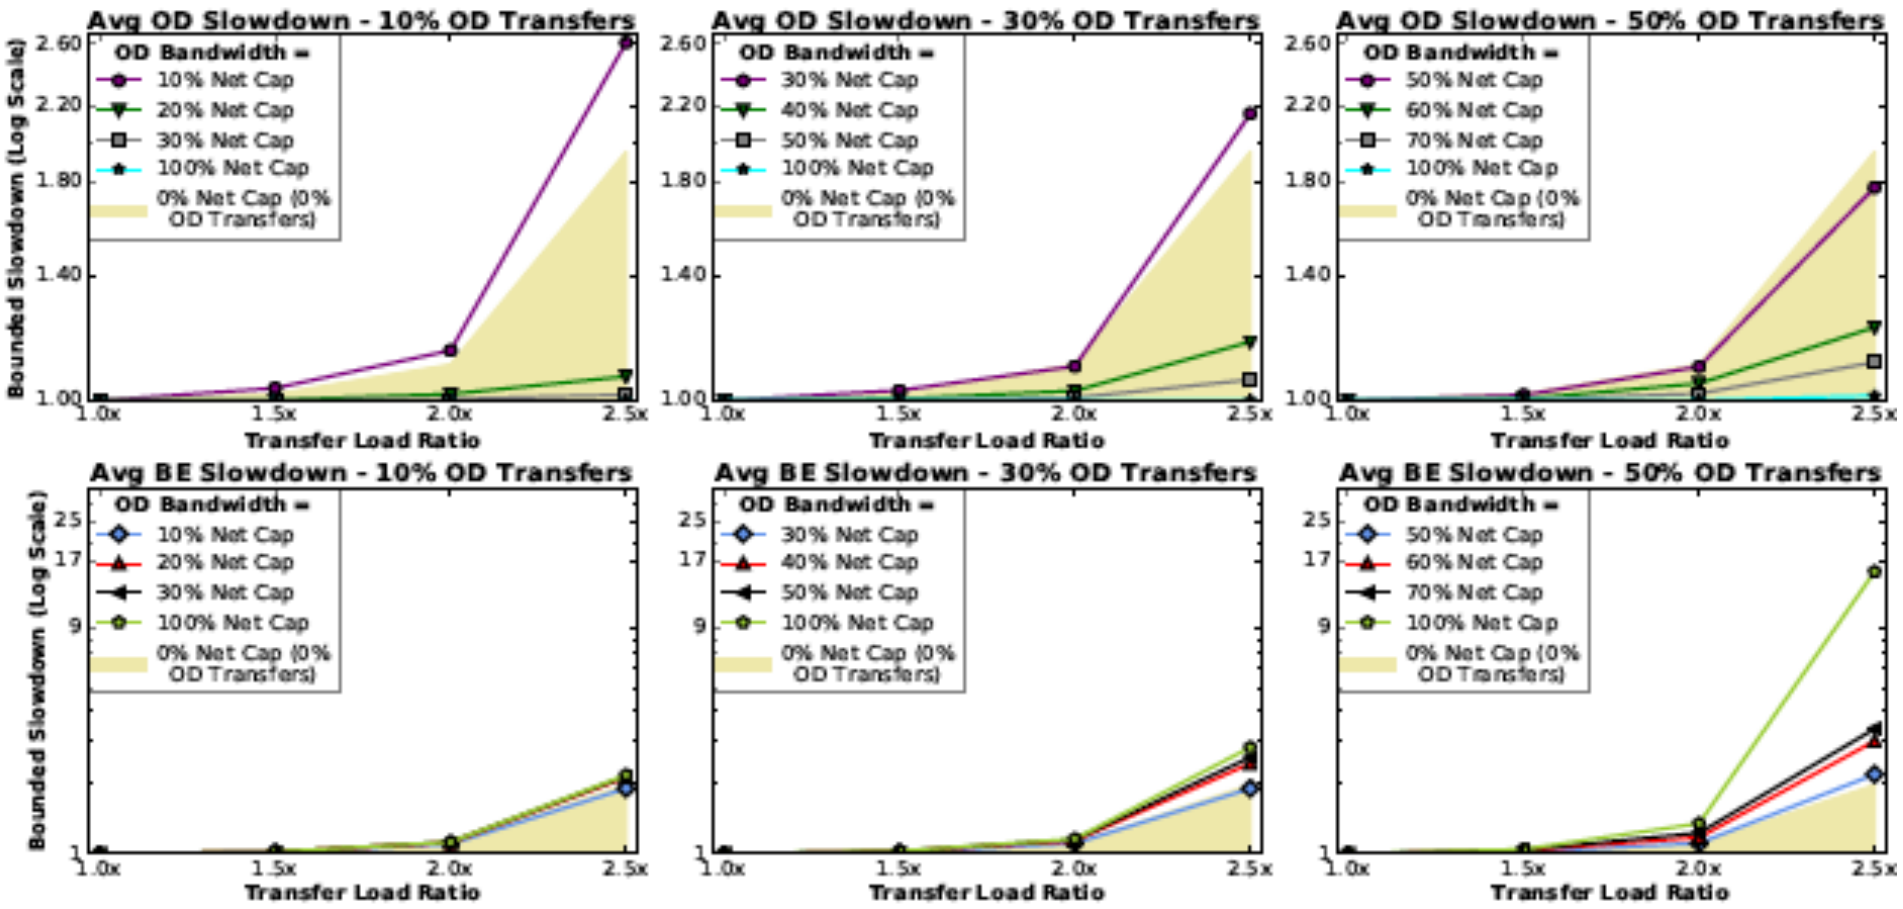

Load Variation

- Trace 2 has a higher mean demand, peak demand, and mean to peak ratio than Trace 3
- Trace 2 - Mean: 2.5 Gbps, Peak: 11.2 Gbps, Mean to peak ratio: 0.22
- Trace 3 - Mean: 1.7 Gbps. Peak: 10.6 Gbps, Mean to peak ratio: 0.16)
- Trace 2 has lower OD and BE slowdown values
- Trace 3 is bursty
  - Concurrency and throughput coefficient of variations for Trace 2 were 0.67 and 0.69
  - For Trace 3, they were 0.87 and 1.07 respectively.

Summary

- Study to motivate measures to reduce the huge gap between peak and average loads in research and education networks.
- Using real world logs, simulated high transfer loads by keeping capacity at current levels and studied the impact.
- Showed current network capacity can handle up to 2x the current load with minimal impact to the data transfers
  - When the peak load is 5x or more than the average load
- When transfers are categorized into on-demand and best-effort, impact on on-demand transfers can be made negligible
  - Keeping the impact on best-effort transfers minimal

Questions