How can peak load in DOE's science network be contained?

NAME: Rajkumar Kettimuthu, Eun-Sung Jung, Sven Leyffer, Ian Foster
AFFILIATION: Argonne National Laboratory
EMAIL: kettimut@anl.gov
EXPERTISE: Transport & application protocols, Analytical & data-driven models, Game theory

Backbone networks are typically overprovisioned in order to sustain peak loads so that all science users continue to have adequate network resources for their science workflows. With the exploding data volumes in almost all science domains and the increasing reach of easy-to-use data movement tools to users, can we sustain the same overprovisioning strategy in the future? We argue instead that researchers should investigate measures to spread the load and keep the peak under control.

MOTIVATION

Under the current mode of free access to the network, the traffic at peak load may range from flows that need to be transferred in near-real time, for example, for computation and instrument monitoring and steering, to flows that are less time-critical, for example, archival and storage replication operations. Thus, peak load does not necessarily indicate the capacity that is absolutely required. Introducing a charge-model for network usage with dynamic pricing that based on demand will create incentives for less time-critical flows to move to periods of low demand. Some electric utility providers and cloud compute resource providers adopt a similar approach. ComEd has a pricing program that lets users pay a rate based on hourly market prices (fixed based on demand) for electricity. Amazon uses spot instances to maximize the utilization of its available resources; the price for spot instances fluctuates based on supply and demand for instances. Access to every expensive facility (including supercomputers and science instruments such as light source) in the DOE complex is regulated through an allocation process. Since a high-speed wide-area network is an expensive resource, too, why don't we regulate access to it through allocation? Arguably, a critical consideration is that these resources are scheduled in advance whereas the dataset sizes that need to be transported on demand might vary significantly.

POTENTIAL BENEFITS

1) Reduced cost by building networks to sustain only the necessary peak load.
2) Improved overall utilization of the network through incentives for using the network in off-peak periods.

KEY DOE NETWORK/TRANSPORT CHALLENGES

1. Credit and token system: Should we allocate credits to users that can be used to request the needed network resources?
2. Track usage on per user basis: Does tracking require changes to the network and/or transport protocols? How should retransmissions be handled?
3. Routing: Should the network layer route packets in the least expensive path by default? How does it impact performance for latency-sensitive applications?
4. Dynamic pricing and transport: How frequently does the burn rate (price) change? How does dynamic pricing affect long-running transfers? How should one handle transfers whose size is unknown until the transfer is completed? How can users control the charge for these transfers, and how does such control affect the transport protocol?
5. Quality of service: How do we incorporate QoS in this charge model and what are the implications for the network and transport layers?

DOE RESEARCH CHALLENGES

Such a system requires understanding the interactions of markets and networks and preventing individuals from exploiting the system. Hence, we need to assess different pricing mechanisms by building models, for example, using game theory, which in turn raises further challenges: (1) formulating suitable leader-follower games to design pricing mechanisms that minimize congestion or maximizing network throughput or quality of service under different usage scenarios; and (2) solving large-scale leader-follower games with realistic network models and demands.