Integrating Storage and Data Transport for Furthering Science

Satyajayant Misra (*New Mexico State University*), Beichuan Zhang (*The University of Arizona*)

With the increase in scientific data collection volumes and the measurement precision, the magnitude of the diverse scientific data generated and transferred by networks is increasing rapidly. Looking at anticipated data volumes at three levels: new dataset size, LAN transfer requirements, and WAN transfer requirements, projections show that within two years these values will be 6 TB/day, 5 TB/day, and 200 TB/day respectively. At the 2-5 year time window, these values will be 20 TB/day, 40 TB/day, and 600 TB/month at 1 GB/sec to the end users (i.e., analysis locations). At the 5+ years horizon they would be 100 TB/day, 200 TB/day, and 3 PB/month at 10 GB/sec [1]. These values represent almost an order of magnitude growth rate every two years. This growth rate is much faster than the Moore's law rate for semiconductors' as well as the scaling of the network infrastructures, thus presenting a significant challenge to the science/networking communities.

Meanwhile storage resources (e.g., memories, hard disks, solid-state drives, etc.) are becoming significantly more available in terms of capacity and price to the point where there are proposals for routers to cache/store significant amounts of data that they forward. Integration of the abundance of storage resources both at the end sites and in the network can significantly help networks meet the ever increasing demand for data transport. An example of the outcome of such symbiosis is: if researchers at the Fermi Lab download a large experimental dataset from CERN over the transatlantic link, then the network-storage can work together to route other researchers in the US interested in the dataset to the Fermi Lab and the routers in the core, rather than requiring them to use the high latency transatlantic link again. Caching and collaborative data sharing will reduce overall network load, spread the load to different regions, and improve user-perceived delay and throughput. Science networks have to be designed to go beyond current core TCP principles, namely end host-to-end host connection and congestion control, so that they can interact with multiple caches or data stores.

While caching has been explored in the past such as in web caching and CDN, the past focus was more on reducing server load and reducing end-to-end delay. In the context of DoE networks, the technical challenges are: (1) to maximize the throughput of data transport, and (2) to efficiently utilize network resources under the unprecedented volume of data transport. This calls for a data transport model that supports ***native multi-source and multi-path content retrieval***. In the ideal scenario, a user will want to download a large dataset, which may be stored or cached, either in its entirety or subsets, at multiple locations, the network is able to retrieve the dataset from a nearby location or multiple locations, using multiple paths if needed, avoid congestion, and switch data sources for better throughput or less congestion when necessary. This model is a dramatic change from the current system, which was designed and optimized for only single-path, single-source file download.

Efficient utilization of the combined network and storage resources can be done by considering three facets: *cache placement, content routing, and congestion control*. These facets are inter-related—results from solving one impact the other two. For example, cache placement affects the routing decision as we would like to retrieve data from one or more nearby sources; routing decision in turn impacts link utilization, i.e., congestion levels; and to alleviate congestion, better cache placement and/or use of alternative routing paths is desirable. In existing research, these three problems are mostly considered in isolation, and often under different contexts. To address DoE network's data transport challenges, we need to develop integrative approaches and solutions to address these problems within one framework. The emerging Information-Centric Networking (ICN) paradigm, Named Data Networking (NDN) in particular, may provide such an architectural framework. In ICN/NDN, as long as the user satisfies its requests, it doesn't care whether the content is retrieved from disks, via network links, from different paths, or different sources. In this sense, it's a promising architecture that unifies the bandwidth resource and storage resource to address the data transport challenges.

References

[1] J. Zurawski. Bridging the technical gap: Science engagement at ESnet. In *Great Plains Network Annual Meeting*, 2015.