TITLE: Challenges to Support Real-time Applications: a science version of Internet of Things

NAME: Alex Sim and John Wu

AFFILIATION: Lawrence Berkeley National Laboratory

EMAIL: ASim@lbl.gov

EXPERTISE: Data Management, Network Performance modeling, Storage Resource Management

KEY DOE NETWORK/TRANSPORT CHALLENGE

DOE is the leading sponsor of research activities in the physical sciences. Its research portfolio includes large experimental facilities, computing facilities, and a high-speed network serving as the glue that binds all others into a grand tool of discovery. As these facilities evolve in the next decade, we foresee a number of disruptive trends that would challenge the existing networking technology. Next, we focus on the challenge and real-time applications and the associated growth of metadata.

In the past two decades, the volume of data has grown exponentially and this growth in data volume was the primary driver in the network usage. In the next decade, we anticipate a dramatic rise of real-time applications. In commercial application, this is appearing as the Internet of Things. In DOE science community, these real-time use cases will be initially dominated by distributed collaboration on large-scale experiments, such as ITER and LSST, where many participants may contribute to the analysis of an ongoing experiment and provide feedback to the control of the next stage of the experiments. As experiment and simulation become more complex, a significant amount of computing resources might be needed to determine whether or how to terminate the experiment or steer the simulation. Such computing resources might not be available right next to the on-going experiment or simulation. Furthermore the data necessary for the analysis might be located far away. In all these cases, there is a need to transfer a potentially large amount of data in a limited time window. As the work on Internet of Things proceed, many of the large science experience will adopt the automated measurement devices that can make its own decisions in real-time and therefore increasing the need for supporting real-time applications.

Related to the real-time application is the radical growth of metadata. Such metadata include information to locate the necessary data to make real-time decisions and results of the real-time analysis. So far, metadata has been relatively small compared to raw data. However, the volume of metadata will increase because more analysis results and linkages among the existing data collections are being created by large analysis. As more automated data collection processes are developed, more systems are performing analysis in situ and keep only the features and other metadata about the raw data records. A number of visionaries have realized the usefulness of metadata services and are starting to produce commercial tools for metadata collection and management. All together, we anticipate a dramatic increase of metadata in the near future.

DOE RESEARCH CHALLENGE DIRECTION

The increasing metadata and real-time applications will drive the network technology in a very different direction than it is going right now. For example, the metadata records are typically relatively small in size, and are often replicated in different forms in a variety of locations. Requests of metadata are usually in the form of queries, which lead to small random accesses to underlying storage systems and small messages passed around the network. The performance of these small messages is latency bound. Reduce and control such latencies is essential to enable real-time applications and associated metadata accesses.

Emerging data storage systems such as burst buffers promise to reduce the data access latency, research is needed for the networking software to take advantage of this new feature in storage and memory systems.