CRAY STORAGE IN EARTH SCIENCES

Michael Feldman

White paper

January 2016

MARKET DYNAMICS

The computational domain of earth science presents a number of challenges for researchers. Demand for higher-resolution and higher-complexity simulations of climate, weather and ocean systems creates a constant demand for greater numerical performance. But, because computation advances faster than storage performance, many of the more critical challenges lie in storage and data management. Exacerbating the storage challenge is the explosion in remote-sensing activity, which produces voluminous amounts of real-time observational data. As a result, designing and building storage systems for these applications tends to be a complex undertaking.

Earth science applications can be divided into three broad categories: operational numerical weather prediction, earth system research and remote sensing. Although there is overlap between these sub-segments with regard to storage and data management demands, they each have distinguishing features that require different types of solutions.

For example, operational numerical weather prediction is focused on performance and consistency since these applications must deliver forecasts in constrained timeframes. (It would be useless, for example, to produce a daily weather forecast that took 48 hours to model.) Also, because people rely on these forecasts to make critical decisions about their activities or business processes, these systems must maintain very high availability, usually greater than 99 percent. In order to guarantee on-time delivery of forecasts many weather centers configure fully redundant systems. As a result, the underlying storage infrastructure must support fail-over capabilities.

Due to the time-sensitive nature of weather prediction and the fact that data is often assimilated in four dimensions to build the forecast models, the simulations are very data intensive. Therefore, I/O must be highly parallelized and scalable to meet the time-constrained deadlines.

Earth system research, on the other hand, involves simulations of longer-term duration involving climate, weather patterns and ocean behavior. This includes observations accumulated over long periods of time, which requires the application to ingest and process vast amounts of data. Typically the data must be reduced during processing in order to fit within the I/O bandwidth and storage capacity constraints of the available hardware.

Remote sensing produces a somewhat different set of storage requirements since it is dominated by streamed data that must be stored and processed in real time. In many cases, remote sensing data undergoes a certain amount of preprocessing, which then feeds into weather and climate modeling systems. Sources are devices such as meteorological satellites, radar and other types of in-situ observational instruments.

Because of the very large amounts of data required to drive earth sciences applications, storage is usually the most complex component of the computer system. I/O bandwidth and capacity requirements tend to be extreme; in the case of remote-sensing applications, I/O latency can be a limiting factor as well. Parallel file I/O

is a feature across the application set, and in the case of weather prediction, accessing data across multiple file systems is often required.

OPPORTUNITY FOR CRAY

Cray supercomputers (XC™ and CS™ series) are well established in the earth sciences market, especially for the larger weather and climate centers that require extreme levels of computing. The company's recent expansion of its storage solutions — in particular, the Sonexion® line — gives it a more complete solution for this market. Besides Sonexion, which represents Cray's high-density, integrated storage solution, Cray also offers Lustre®-based solutions for DataDirect Networks (DDN) and NetApp (E-Series), all of which are integrated with the Lustre File System by Cray. These systems are described in greater detail in a companion paper from Intersect360 Research describing Cray storage solutions for the energy sector.[1]

All of Cray's storage solutions can also be used in conjunction with third-party HPC clusters and supercomputers. This is useful for customers who purchase compute and storage separately, which is often the case for remote sensing applications and occasionally for weather forecasting and earth science research systems as well.

The company also provides extensive expertise for Lustre, and includes integrated solutions of the file system for its storage solutions as part of its data management platform. Through an offering called Cray Cluster Connect, Cray offers end-to-end Lustre support, from the storage hardware to the client, and can provide performance optimizations based on the storage/compute configuration and the nature of the application.

In addition to Lustre, Cray's data management platform includes a set of data mover services, which offers connectivity to third-party storage systems and an interface to hierarchical storage management. The idea is to provide a seamless method for data movement between Lustre, NAS and non-Lustre file systems.

As a result, Cray storage is well-suited to a range of situations in the various sub-segments of the earth science domain, and Cray engineers can provide guidance and support for those users less familiar with Lustre and the complexities associated with scalable storage.

The company's ability to provide advanced data management solutions and the necessary support to integrate with existing storage, file systems and compute infrastructure has resulted in significant wins in this sector: one at Germany's National Meteorological Service, the Deutscher Wetterdienst (DWD), another at the Korea Meteorological Administration (KMA), and another at the European Centre for Medium-Range Weather Forecasts (ECMWF).

KMA employs a multi-tier storage system for its weather forecasting and climate research missions. It consists of 879 TB of DDN S2A 9900 volumes (Tier 0), 1.6 PB DDN S2A 9900 volumes (Tier 1), 507 TB of a DDN S2A virtual tape library, and 4 PB of a Spectra Logic T950e tape library. The Tier 0 and Tier 1 storage is set up as a globally shared Lustre file system. KMA subsequently expanded the Lustre storage with a 1.5 PB Cray Sonexion 1600 system that adds more than 24 GB/s of additional bandwidth. The storage is connected to dual Cray® XE6™ supercomputers. System availability is reported to be at least 99.5 percent.

In the case of DWD, the agency purchased two Cray Sonexion 1600 Lustre storage systems for its numerical weather prediction center. The storage systems were installed in conjunction with a pair of Cray® XC30™

[1] Intersect360 Research, "Cray Offers Scalable High Performance Storage for Energy Sector," January 2016.

Intersect360 Research
P.O. Box 60296 | Sunnyvale, CA 94088 | Tel. [888] 256-0124
www.Intersect360.com | info@Intersect360.com

supercomputers (one is used for production, the other for research and backup). The Sonexion systems at DWD have a combined capacity of more than 3 PB and an aggregate bandwidth of 72 GB/s.

INTERSECT360 RESEARCH ANALYSIS

Earth science simulations are among the most demanding of "big data" applications. By its very nature, the large amount of data involved necessitates storage and data management schemes that can reach extreme scale in both capacity and I/O bandwidth. And since meteorologists, climate researchers, oceanographers and other earth scientists demand ever-more-accurate models of these physical processes, there is constant pressure to collect and process even greater amounts of observational data and produce higher-fidelity simulations.

Such an environment is hostile to commodity storage, instead demanding the more advanced solutions on the market in terms of performance, capacity and interoperability. The complexity of meeting these demands means there are a limited number of storage solutions that are appropriate in these situations and just a handful of vendors with the expertise to build and support such systems.

Cray, with its long history in extreme-scale systems, is one such vendor. The company is well established in providing supercomputers to weather, climate, and oceanographic centers and now offers a complementary set of offerings for the storage side. Its use of high-end storage systems (Sonexion, DDN, NetApp) integrated with its Lustre and data management software stack provides a complete, integrated solution for earth science workloads. In addition, Cray's expertise in delivering such systems successfully over many years makes it a trusted choice for this application domain.

Intersect360 Research
P.O. Box 60296 | Sunnyvale, CA 94088 | Tel. [888] 256-0124
www.Intersect360.com | info@Intersect360.com