Improving Performance for Data Access Web Services

Summary of a panel discussion held 6-January at the 2015 ESIP Winter Meeting

The stage was set for the subject panel by the appended abstract, distributed beforehand as part of the ESIP conference materials. Introduced by Dave Fulker, the session context was set by Jeff Walter (NASA/ESDIS). Jeff spoke about the U.S. Big Earth Data Initiative (BEDI) and about NASA's goals both to further exploit and to enhance the performance of web-services for data access, enabling seamless, responsive workflows for a large, diverse community of users. Then Dan Pilone, Jason Werpy, Jim Frew and James Gallagher gave 15-minute presentations on problems and/or solutions related to the performance of data-access services.

The talks were otherwise quite varied, but common to all four was consideration of how source data are stored at the point of web-service provision, including cases where client queries require information from very large numbers of source files (often called granules). Such cases are of particular note because they are common and embody performance impediments that are significantly independent of the web-service protocol, per se. More specifically:

A. Earth-science data archives increasingly comprise granules that number in the tens or even hundreds of thousands.
B. The granularity of archives often is determined by patterns of collection, so patterns of retrieval can literally be orthogonal to the storage structure. Therefore, absent clever implementation, the latency of end-user retrievals can exceed the sum of thousands of file-access latencies.
C. A web-service protocol can hide this granularity by, for example, offering clients a (virtual) aggregation of the underlying source files. Though logical and convenient, such hiding can exacerbate the performance issue because users typically expect behaviors consistent with file-based data access.

Here are a few highlights from the panelists' talks.

- **Dan Pilone (Element84) -** Intending to provoke discussion, Dan argued that better utility and performance can be gained by departing from the current pattern of storing Earth-system data only in a single location and, hence, a single form (which differs little from that of 9-track tapes used three decades ago!). Pointing out that NASA holdings are at least an order of magnitude smaller than those at Facebook, Dan suggested that sub-second response times are entirely practical for all forms of scientific data if we commit to migrating them into cloud contexts such as Amazon Web Services, and then exploiting the potential of parallel computing and multiple representations. He also touched on performance gains already being demonstrated at NASA in the Global Image Browsing Services (GIBS) and the Common Metadata Repository (CMR).
  Audience feedback led to a suggestion for measuring web-service value with a metric like "data returned / data used to calculate the response." This ties service-effectiveness to reducing the volume of data needed to answer a particular query.
- **Jason Werpy (NASA/LP-DAAC) -** Noting that providers often measure delivered data *volumes*, Jason suggested that a better metric would be data *value* delivered, where this entails pushing "algorithms to archive" (such as server-side averaging, e.g.). He described the Data Access Protocol (DAP) as an important step toward "web-addressable data," but need for parallel computation arises with large numbers of granules (as above). Jason described experimenting with improved performance via DAP implementations on clusters, where multiple source granules can be accessed in parallel. The price for these gains is the need to transfer source data into storage systems directly accessible by the cluster or cloud. (Note the consistency with Pilone's comment about copying source data.)

- **Jim Frew (UC-Santa Barbara) -** Departing from Werpy's DAP emphasis, Jim described performance and other benefits gained from using an array database (SciDB) whose query mechanism includes an algebra of array operations. Such operations permit users to request computations as well as to select data, and SciDB is designed (from the ground up) to exploit parallel-execution environments, such as clouds. A key aspect of Jim's approach is assigning every source-image pixel to a single element in a huge, very high-resolution (~10 meter) latitude-longitude array. Such assignment capitalizes on SciDB's sparse-array optimization, and it efficiently hides granularity as well as variations among observing instruments. This simplifies aggregation and other applications of array operations because all spatial arrays have identical shapes. The price for this simplicity, flexibility and performance is a need to transfer source data into SciDB's internal representation. (Note here too the consistency with Pilone's comment about copying source data.)
- **James Gallagher (OpenDAP) -** Returning the discussion to DAP (invented by OpenDAP two decades ago), James reinforced several points made by earlier speakers but also talked about the extent to which performance problems can result from the actions of clients and not simply from server behaviors. In particular, he highlighted clients and client libraries designed on the assumption that data retrievals are from local files. Because file-access times are much less than typical Internet latencies, web-service performance is severely suboptimal in clients that, for example, were originally designed to make large numbers of (small) file retrievals. Fixing all clients is far beyond what OpenDAP can do, but James observed that clients will evolve toward best practices if server interfaces and configurations are advanced to offer improved performance.

After the talks, Fulker invited comments and questions, projecting the slide below to help keep the discussion focused. A vigorous, 20-minute question-and-answer period ensued, engaging nearly a half-dozen members in the audience (of 50 or so people, filling the room beyond its seating capacity). Some comments/questions were about budgetary constraints that limit consideration of data migration, and some were about the technical details of concepts put forth by the speakers.

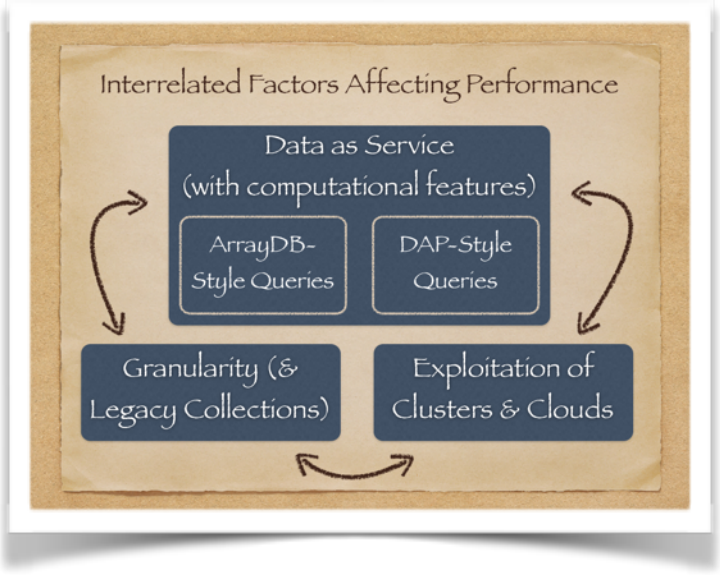

Published Abstract: *This panel discussion will feature experts on data retrieval via Web services, emphasizing technologies and methodologies that will help OPeNDAP and OGC web service users (e.g., NASA/ESDIS) achieve sub-second response times for important aspects of its data discovery and retrieval services. The panel will present material, answer questions from the audience and note suggestions from the audience. The format of the discussion will be informal and invite participation from all attendees. Likely technologies to be discussed are data organization techniques, data replication, data storage technologies that are not based on file systems and user interfaces for on-the-fly data aggregation.*