3 | DATA MANAGEMENT SYSTEM

HISTORICAL CONTEXT OF CURRENT OPERATIONS

IRIS operates one of the most actively used scientific data centers in the world. The Data Management System (DMS) ingests an exponentially increasing volume of observational time series data every year, from an increasing number of seismic networks and stations. Currently, more than 20 terabytes (TB) of new data are being added annually to the holdings. DMS continues to deliver an increasing amount of data to the research community through batch requests to the archive, streaming data in near-real time, and using advanced application program interfaces (APIs). APIs allow remote clients to access metadata directly in an Oracle Database Management System as well as access the time-series data in mass storage systems. Current projections indicate that the DMS will deliver 80 TB of data to the research community this year. This output-to-input ratio of four attests to the importance of the DMS to the community it serves.

The DMS consists of the Data Management Center based at the University of Washington in Seattle, Washington, and two primary Data Collection Centers (DCCs): the U.S. Geological Survey DCC operated by the Albuquerque Seismological Laboratory (ASL), and the University of California, San Diego, DCC operated by the International Deployment of Accelerometers (IDA) project. The University of Washington receives DMS support for host activities that include data quality assurance and the development of specific applications and algorithms for use at the DMC. IRIS supports additional centers in Central Asia, enabling access to high-quality data sources there. This organization provides a stable pipeline for the flow of data from a variety of sources, though a consistent quality-control process, and into the data archive.

Although it remains challenging to operate the infrastructure and systems that provide access to the thousands of seismic stations, data are now received, archived with backups both on and off site, and distributed via automated mechanisms that work seven days a week, 24 hours a day.

Primary storage for most of the time-series data is a large, disk-based RAID system, but some voluminous, infrequently accessed datasets are still stored on two tape-based robotic systems. The DMC continually improves the software systems that lie at its heart, distributes and supports key software applications used by the global seismological community, and supports and distributes applications that are used by global data centers.

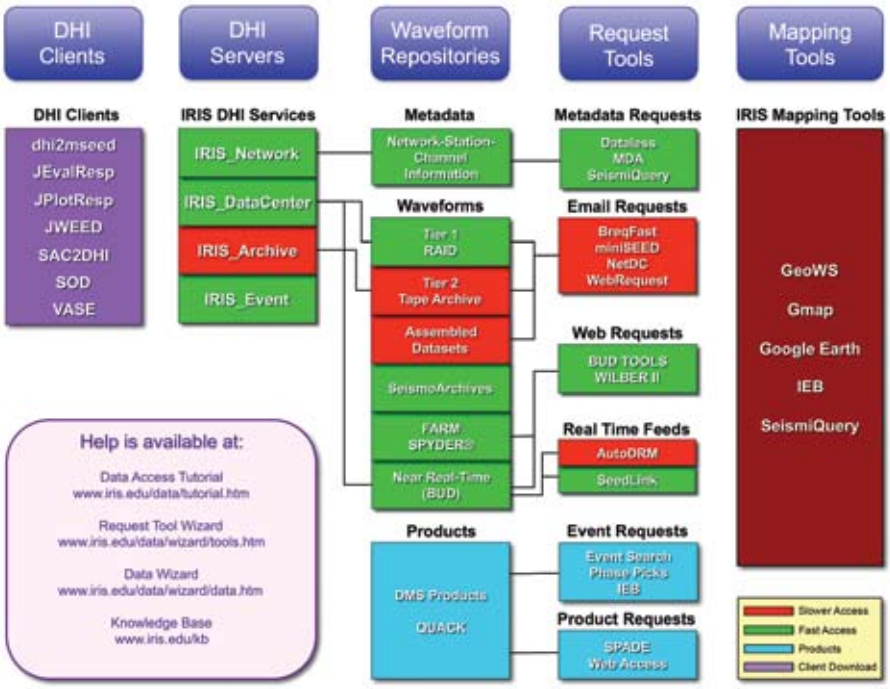

Figure A3.1. Organization of user services and access methods at the DMC.

DATA COLLECTION

Data are received at the DMC through a variety of different paths and protocols depending on the data source.

- *GSN:* ASL and IDA produce data from core GSN stations in SEED format, maintain the metadata describing the seismological observatories, and perform quality control of the waveforms and metadata. DCCs forward the data in real time as well as after completing quality review, usually within a few days.
- *PASSCAL:* PASSCAL experiment data primarily come through the PASSCAL Instrument Center (PIC). Experiment data are sometimes received in near-real time but normally are received with delays of months to years after an investigator releases the data and the PIC completes quality review.

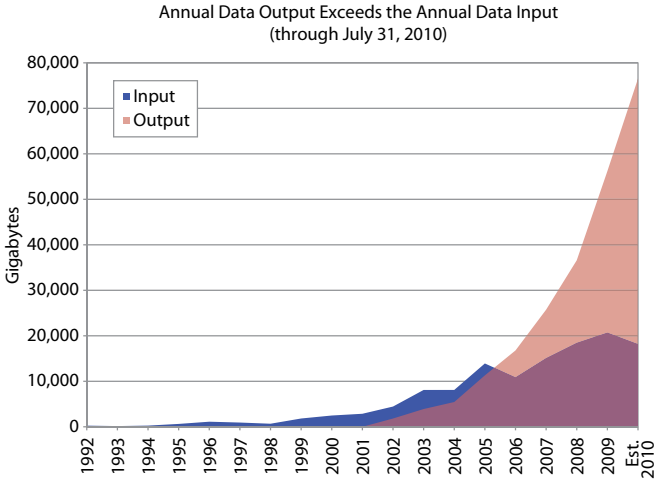

Figure A3.2. The IRIS DMC is an extremely active data center and has become a central component enabling seismologists to do their research. Beginning in 2006, the amount of data flowing out of the DMC exceeded the amount of new data arriving at the DMC, with an estimated 80 terabytes of data being shipped from the DMC in 2010 versus 20 terabytes of new data flowing into the DMC. A ratio of 4:1 is extremely high for a scientific data center and attests to the active use of the IRIS DMC.

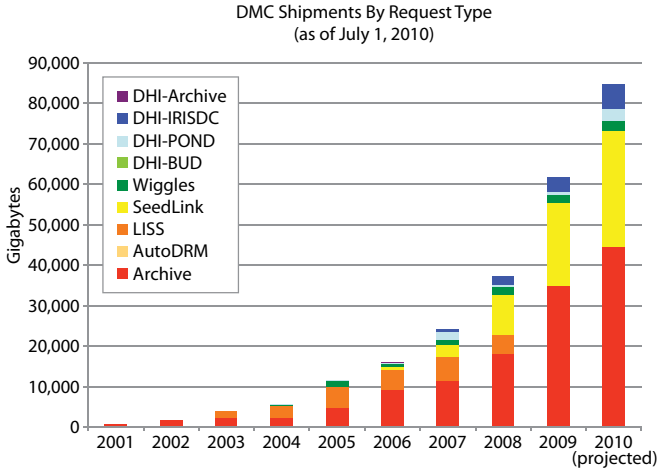

Figure A3.3. This figure shows the amount of data, in gigabytes, flowing out of the DMC by year and by various request mechanisms. We are currently projecting 80 terabytes will be shipped in 2010. The figure shows data shipments from the archive by traditional mechanisms (labeled Archive), real-time feeds (labeled LISS and SeedLink), and by the Data Handling Interface, (labeled as one of four DHI mechanisms). Users of the IRIS DMC use all three primary request mechanisms heavily.

- *USArray:* Transportable Array data are received from the Array Network Facility (ANF) in San Diego. Quality-controlled data from ANF, consisting primarily of data with some gaps to be filled, are received within a few days of real time. Flexible Array experiment data are received from the Array Operations Facility (AOF) in Socorro, New Mexico. Data from the USArray Reference Network are transmitted to the DMC from the USGS National Earthquake Information Center (NEIC) in Golden, Colorado. Three analysts at the DMC review USArray data for quality.
- *U.S. regional networks:* USGS makes data available from all regional networks it supports in the United States, including the U.S. National Seismic Network as well the Advanced National Seismic System. NEIC makes some data available in real time and forwards a quality-controlled version several days later. Most regional networks provide data directly to IRIS and do not make a quality-controlled version available.
- *FDSN:* The International Federation of Digital Seismograph Networks (FDSN) is comprised of approximately 65 different organizations in more than 50 countries. IRIS receives data from the majority of FDSN networks in real time through a variety of protocols. Quality control for FDSN data at the DMC is limited to automated quality assurance processes that are applied to data as they are received in real time.
- *Other networks:* IRIS has bilateral arrangements to collect, manage, and freely redistribute data from many other non-FDSN networks around the globe. Quality control for these data is also limited to automated processes applied at the DMC as they are received.

One of the most significant changes in data reception is that more than 98% of data from **permanent** networks is now received in near-real time. A decade ago, most data were received through non-real-time methods. DMC operates an automated quality-control system that continuously calculates metrics related to data quality for all data received in real time (see http://www.iris.edu/servlet/quackquery). This information is available to any researcher through a variety of web interfaces.

DATA DISTRIBUTION

Originally, IRIS's goal was to service user requests for data given specific lists of seismic stations, recording channels, and time windows. It was expected that a request for data for one earthquake from all GSN stations—about 27 MB of data—could be met within 24 hours, and that there would only be a few such requests daily. Actual services have surpassed this initially daunting requirement by orders of magnitude. In 2009, end users received roughly 60 TB of data, more than two million times more data than for the single benchmark data shipment envisioned in 1987. The number of customized requests in 2009 was more than 600,000—more than 2,000 requests per day.

Data are served through three fundamentally different techniques: responses to formatted email requests such as Breq_fast and NetDC (66%), real-time data feeds (21%), and well-defined interfaces implemented on servers that interact with client applications on individuals' computers (13%). The set of request mechanisms available to researchers is rich and powerful, and targets different types of user requests effectively.

As the number of seismic networks and stations continue to grow, data services continue to be adapted to the needs of users. Gone are the days when users "knew" the names of the stations whose data they sought. Users today are more interested in defining broad regions from which they wish to extract data with characteristics that match their research needs. Most of the significant improvements to data services in the future will be driven by users' needs and employ newly developed tools that let the users get the data they need, sometimes with preprocessing applied. To provide these services, powerful web services are being developed through which users can access data very simply without writing their own applications.

DEVELOPMENTS UNDER THE CURRENT COOPERATIVE AGREEMENT

DMS has always been a leader in the seismological community in developing new tools through which information can be accessed. Additionally, DMS has pioneered methods by which a distributed data center concept can be implemented. An email-based system called NetDC and an Internet-based system called the Data Handling Interface are capable of accessing data from distributed centers seamlessly. Because both of these techniques have limitations, the DMC has been developing modern web services techniques through which a remote client can access time series and the metadata describing them. During the past five years, the DMC has continued developing the tools and systems the scientific community needs for their research, as described below.

DATA COORDINATION

Global

As the FDSN archive for continuous data, the DMC collects data from the FDSN Backbone Network of 200 stations and aims to collect data from one or more stations in every network operating with an FDSN code. The archive currently includes data from 124 of these permanent seismic networks, including 92 networks from which data are received in real time.

Regional

Data sharing is important for seismogenic zones that cross national boundaries, and IRIS focuses resources to improve data exchange in a few critical regions, such as a Central Asia. Representatives from the DMS and the German Research Center for Geosciences (GFZ) met last year with seismologists from Kazakhstan, Kyrgyzstan, Uzbekistan, and Tajikistan, and formed the Central Asia Data Exchange (CADE) group, partly based on GFZ's Central Asia Real-Time Earthquake Monitoring Network (CAREMON) project. The CADE group agreed to readily share data among the Central Asian nations and the outside world via IRIS and GFZ and to work together to seek funding for ongoing network operation within their countries.

DATA HANDLING

Offsite Active Backup

NSF's 2008 management review of IRIS stated that seismological research now relies so heavily upon data services that research would be impaired by a prolonged DMC outage. Prior to this review, all time series, software, and information in the Oracle database were being replicated at an active backup location. Based on the review directive, IRIS began developing the capacity to fully service requests from the remotely located active backup. Although it is now possible to service some user requests from the active backup, development continues. Currently, all routine request processing is still done at the primary DMC in Seattle. In the event of catastrophic failure, however, access to data can be provided from the active backup location.

Enhanced Support for Real-Time Data

One of the primary efficiencies that made it possible to manage exponentially increasing amounts of data and serve an ever-increasing amount of data to the community has been the development of automated systems for real-time data ingestion. Recently, the DMC also developed the capability to distribute data in near-real time using the SeedLink protocol developed by GFZ. Extant SeedLink server systems were incapable of handling the volume of real-time data that users require, so the DMC developed a new system that follows the SeedLink protocol. DMC also developed a "turnkey"

SeedLink server that can be given freely to any network, removing a technical obstacle for networks that wish to share their data in real time.

Production of Merged Data

Often, near-real-time data ("R" data) are received out of order and in duplicate. Additionally, some networks make available a quality-controlled version ("Q" data) of their data within a few weeks after data recording. DMC implemented methods to merge R data with Q data, including complicated steps to remove gaps and overlaps, in response to requests from users.

Seismic Analyst Review

Because of the increased need to manage and assess the quality of USArray data, three PhD-level seismologists now provide expert review of the seismic data from a research perspective. This type of staff support is one of the key reasons why the data quality from the USArray components is so high. With available funding, the capacity to generate automated quality metrics for seismic data can be extended while retaining staff to review the waveforms for quality.

DEVELOPMENTS

Products

A workshop in 2004 (http://www.usarray.org/files/docs/pubs/USArrayProducts.pdf) identified the need to develop a variety of value-added products using data from USArray sources. This concept has been extended beyond EarthScope data to data from many other sources. A vibrant product development activity is now in place (see http://www.iris.edu/dms/products/ for current product availability), and it is clear that this is a priority area for users. Representational State Transfer (REST) web services are the basis for a product management system that allows discovery and delivery of all products meeting specific product type, geographic area, or time constraints. To support this increased effort in products, two product specialists were added to the DMC staff in the past year using EarthScope project funds. These PhD-level staff provide a seismological perspective that is extremely useful in the development of new products as well as an assessment of product generation and quality.

Web Services

DMS has been a leader in the development of new approaches to information dissemination. During the past two years, significant progress has been made in developing a series of REST-style web services that provide access to time-series data, event catalogs, and metadata that describe the events and seismic observatories. The newly developed web services will permit researchers to receive time-series data in a variety of formats, with user-selected processing already applied to the data. A wide variety of processing services are being developed, such as mean removal and down-sampling, filtering, gain correction and unit conversion, instrument response deconvolution, and time-series integration and differentiation. Such services can simplify access to information and encourage use of IRIS data by other communities.

NEW OPPORTUNITIES AND DIRECTIONS

The new tasks the DMS proposes in this 27-month proposal focus on helping the research community address the science identified in the *Seismological Grand Challenges* report. Activities are anticipated in three primary areas:

1. *Continue to ease access to even more seismological data* by providing data curation facilities for networks without existing archival infrastructure, and by acting as a data broker to other seismological data centers with holdings that complement those of the DMC.
2. *Broaden web service development activities to integrate data with groups outside of IRIS.* More complex scientific problems will require simple access to more diverse datasets. While some of these datasets can fit within the IRIS data model very easily, datasets from other domains are best accessed through well-coordinated web services. Web services will allow IRIS data to be more accessible to other communities for their specific uses. Additionally, the IRIS community will be able to more easily access data from other disciplinary data centers that support web services. DMS will develop systems that enable better integrated data access by our internal community as well as those scientists in external fields.
3. *Produce additional products that will aid researchers in pursuit of better understanding of Earth systems.* As the complexity of scientifically interesting research increases, it is clear that the production of higher-level products from which research may begin will assist researchers in studying more complex problems without always having to do routine and mundane data processing.

EASE ACCESS TO MORE SEISMOLOGICAL DATA

Data Brokering Services

In the past, DMS has supported the installation and support for distributed data centers by installing data-access technologies at the specific data centers. While this approach works for some data centers, it does not work in general. For this reason, we propose to develop an FDSN-sanctioned data brokering service. Instead of installing a new data-access technique at a specific center, we will instead provide a service that works with data-request mechanisms already in place and supported by a specific data center. A request will be received by the brokering service, it will be translated into the request method supported at each of the relevant data centers, data will be assembled by the brokering service, and the resulting data volumes meeting the user's specific request will be returned to the requestor. We realize that the full capability cannot be provided at all data centers, but in terms of receiving data in SEED format, the brokering service will very likely provide significant new data to meet a researcher's request.

Value-Added Services for Network Operators

Over the more than two decades of its operation, the DMC has been very successful at opening up data from networks all over the world. To foster even more widespread data sharing, networks with inadequate funding, but a commitment to open data sharing, will be offered a series of value-added services available from the DMC. Of particular interest are the various metrics that DMC computes in its quality-assurance system for all real-time data flowing into the DMC. As an example of their use, managers of the Australian National Network operated by Geosciences Australia have stated that the metrics measured by the DMC QA system are invaluable to them for the operation of their network.

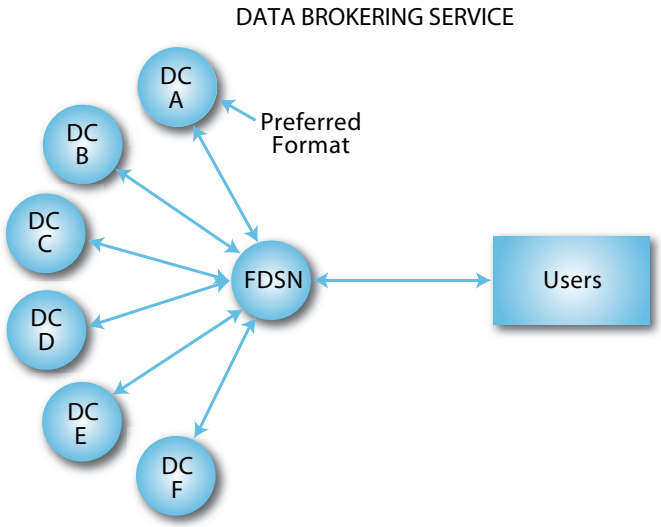

Figure A3.4. This figure depicts how a user of the Data Brokering Service would send a single request to the FDSN brokering service node. The broker would determine which participating data center manages the data being requested. The brokering service would then use the data retrieval mechanism preferred by the specific data center to retrieve the data. This process would be repeated for all data centers holding data requested by the user. The data would then be assembled at the FDSN node into a single product that would be returned to the user, satisfying their request.

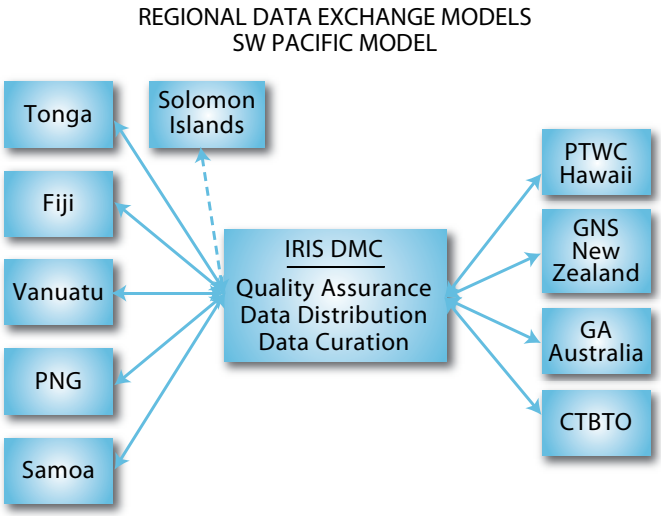

Figure A3.5. The IRIS DMC offers many value added services for network operators contributing data to the IRIS DMC. For instance, as real time data are received, a variety of algorithms are applied to the data that serve as indicators of data quality. A network operator can reference this database of quality estimates as part of monitoring their network performance. Network operators such as members of the SW Pacific Tsunami Task Group, depicted in this figure, are considering sending their data to the IRIS DMC for quality assurance, data archiving, data distribution to other Task Group members, distribution to the international community, and for long term data management. Other international efforts in Central Asia as well as in Africa and the Middle East are also considering this model of data exchange.

Recently a group of countries in the Southwest Pacific met to discuss data sharing in their region. Their decision was to work within the model (Figure A3.5). Data from each of the networks (Samoa, PNG, Vanuatu, Fiji, Tonga, and the Solomon Islands) will send their data to the DMC for quality assurance, data curation, and data distribution, eliminating the need for each country to develop their own systems supporting these capabilities. DMC has developed and has available a turnkey SeedLink server that supports real-time data distribution from a network. Data will be redistributed to neighboring countries through a SeedLink protocol already supported by the DMC. Much of the technical work for this capability is already complete but the promotion of this data-exchange model will require additional effort.

Workflows

DMS develops tools that bring powerful capabilities to the scientific community that are not typically available for purchase. These tools can be linked together in user-defined sequences or workflows. For instance, users will be able to request data, select data based upon waveform attributes (e.g., continuity, signal-to-noise ratio), filter the data, correct for the instrument response, and reformat the data before

having the data returned to them for further analysis. While the DMC will produce some workflows that are pre-configured for typical operations, we will also work with Microsoft Research in the use of the Trident Scientific Workflow Engine. A user will be able to interact through a Silverlight-enabled browser with the Trident Workflow Engine to create customized workflows.

WEB SERVICES DEVELOPMENT

Data Access and Processing

DMS proposes to continue development of web services that allow access to waveforms, event catalogs, metadata, and products. This type of access is the current paradigm for information distribution in a platform-independent, scalable manner. We intend to focus our efforts on REST-style web services and coordinate our efforts with our FDSN partners.

We will focus our data-access services on time series in SEED format, event catalogs, and metadata describing the time series and the events. Access to products managed at the DMC will also be enhanced specifically for products in the DMC Product Management System.

We will also develop a series of seismological and generic time series analysis services and expose them as web services. For instance, data rotation, data down sampling, and instrument correction will be offered. In the area of generic time-series services, we anticipate such things as demeaning, tapering, filtering (low pass, high pass, and band pass), correlation, differentiation, integration, convolution, and deconvolution.

We have the existing capacity necessary to maintain the data-access services on computational and storage systems acquired and operated by the DMC. We intend to deploy processing services in a more scalable environment, however, because we cannot independently maintain systems capable of meeting peak computational demands from the external community. First deployments will be on a condominium-style architecture operated at the University of Washington. As developments in cloud services continue, we will be prepared to move to a cloud environment when appropriate.

PRODUCT DEVELOPMENT AND SUPPORT

Project Management System

DMS will continue to support product development through its evolving Product Management System. Staff includes multiple product specialists who take community-vetted ideas and turn them into products, and a developer who focuses on the software infrastructure needed to convert product ideas into actual products.

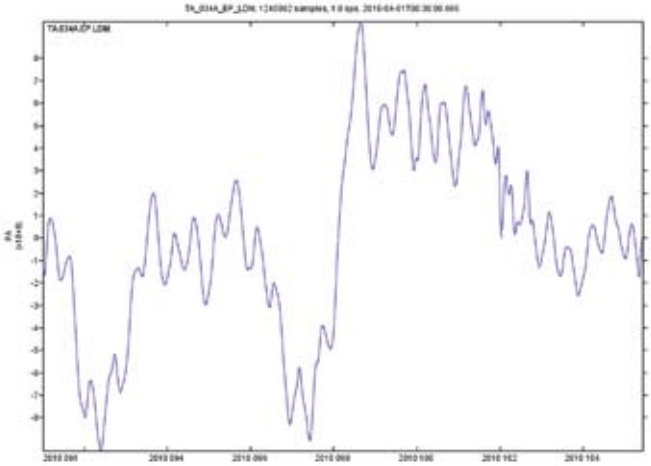

Figure A3.6. Easing Access to IRIS Data for Other Disciplines. The IRIS DMC manages data from approximately 24 different kinds of sensors in addition to seismometers. Traditionally the data have been esoteric and generally understood by seismologists only. This raised a significant impediment for use of IRIS managed data by scientists in other disciplines. The IRIS DMS has embarked on the development of web services that are capable of providing data, with a variety of user controllable preprocessing applied, and with output in very easy to understand ASCII formats. As shown in this figure barometric pressure data as recorded by the EarthScope funded Transportable Array can be provided to an atmospheric scientist in units of pascals, low pass filtered or with other signal processing algorithms applied.

Simplified Data Access for New or Infrequent Users

During the early stages of development of the DMS, the primary responsibility was to serve the needs of research seismologists with extensive experience in data processing and familiarity with the often-obscure details of data formats and instrument response. As access to data has improved and interest in seismology has expanded, services and user interfaces are being developed that respond to the needs of the non-specialist. For instance, a service that allows simple access to corrected time-series data via a URL in a web browser can be a convenient mechanism for displaying waveforms, exploring the archive, and extracting limited amounts of data. The time series displayed in the section in Figure A3.6 shows two weeks of barometric pressure variations in which the raw data have been low-pass filtered and converted into units of pressure. In addition to the screen display, the data can be exported as a series of time-value pairs for convenient input into analysis programs that can accept generic time-series data or even to a spreadsheet program such as Microsoft Excel. This approach is suited to a scientist who wants limited data access, but does not necessarily require all of the details contained in the complex SEED data format. Others may want the data returned directly as a SAC-formatted file rather than having to receive the data, run rdseed, convert it to SAC format, and then manipulate it in SAC. Still others may wish to have the data delivered to an application such as MATLAB. Simplified User Access streamlines access to data by contracting everything into a single step. It eliminates the current approach

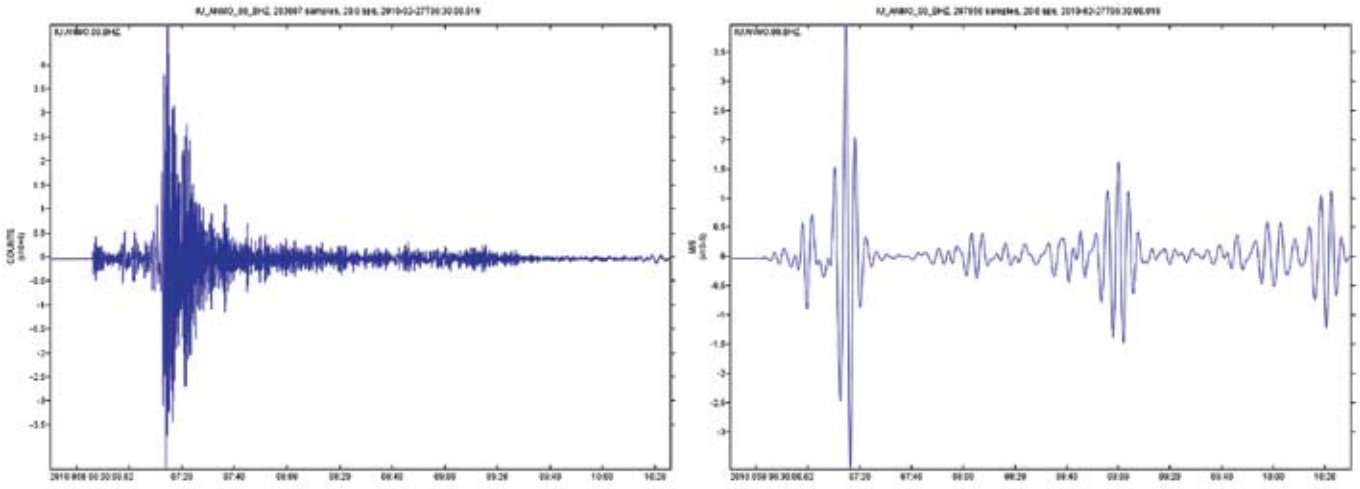

Figure A3.7. The graphic on the left shows the 2010 Chilean earthquake recorded by station ANMO in Albuquerque, NM. The vertical axis is digital counts. The right-hand graphic shows the same information after being corrected for instrument response, and low-pass filtered at 100 sec. The vertical axis is now in units of velocity (meters per second).

of request formulation, request transmission, generation of dataset, retrieval of dataset, conversion of dataset, and ingestion of dataset into the desired analysis tool. DMS will explore the needs of a variety of user communities and significantly expand these efforts to encourage broader use of the increasing variety of data stored at the DMC.

Processed Data Streams

An important attribute of the DMC archive is that primary data are stored with full fidelity, maintaining the complete resolution of the originally recorded time series. No compression or pre-processing is applied that risks degrading the original data. It has been the responsibility of the user to perform instrument corrections, unit conversion, or other processing as part of the scientific analysis. This can be inconvenient or even burdensome for users with limited software tools or processing capability. Tools are being developed to apply a standard suite of algorithms for those who would prefer to have pre-processing performed at the DMC—either for the casual user with limited local resources or the heavy user wishing to pre-process significant quantities of data. For instance, Figure A3.7 (left) shows an earthquake recorded at station ANMO. It is the raw seismogram in units of counts. A preliminary web processing service is in place at the DMC that can correct the data for the instrument gain, and demean and filter. Figure A3.7 (right) shows the same record after it has been low-pass filtered at 0.01 Hz and gain corrected. The units of this waveform are in meters per second after these corrections. A variety of other services are contemplated, including the ability to down sample data from a high sample rate to a lower sample rate. While these added services will be of great utility to a number of user groups, the DMC will always provide the ability to access the raw, unprocessed waveforms for those scientists wishing to more fully control the processing environment.

Request Filtering Using Predetermined Metrics

As the volume of data continues to increase, it becomes more important to provide automated tools and services that allow users to conveniently asses data quality and select those data most appropriate for their research needs. The real-time quality assurance system in place at the DMC calculates many metrics related to data quality as the data enter the DMC. As DMC moves to a web service infrastructure, we can begin adding Quality of Service (QoS) capabilities for data access. During the next 27 months, we plan to begin making progress in this area. We are not proposing to perform dynamic QoS services at this time, but that capability should be considered in the next five-year proposal from IRIS.

Dedicated Processing to Condominium and Cloud Service Models

DMC has historically operated all of its own storage and processing systems using IRIS systems administrators. During the next 27-month proposal, we propose to begin deploying processing web services on the University of Washington HYAK condominium-style computational resource located near the DMC in University of Washington facilities. In essence, a user of the HYAK system buys a specific number of slots providing a specific number of processors, high-speed scratch space, and slower access storage systems, including lower-performance disk and tape systems. The purchase of these slots is good for three years. At the end of that period, a user may choose to purchase an equivalent number of slots for the same price, however, the number of cores and fast

storage is likely to have more than doubled in that time and so the same funding will permit access to significantly more computational and storage capacity.

The condominium model also allows access to unused cores in the HYAK system when needed, and similarly, others may use your idle capacity. However, the owner of the slots may always preempt a non-owner when needed. This type of system will allow the DMC to meet peak demand requirements, for instance, after a large earthquake, at a much lower cost than normal.

We envision using HYAK as an initial test of cloud-like services over the period of the 27-month proposal. IRIS would be responsible for deploying services in HYAK but would be relieved of the operation of the system, thus eliminating the need to have IRIS paid systems administrators. DMS plans to provide data services on a modest, multinodal system located close to our primary storage system; however, the processing web services would be deployed in HYAK. These processing services will have a much higher requirement for shear processing power to meet the community's needs, while the data access services are similar to what DMC currently operates and for which DMC has processing capacity to meet current community needs.