

Data Management at the Texas Advanced Computing Center

June 7, 2013

Chris Jordan



THE UNIVERSITY OF TEXAS AT AUSTIN
TEXAS ADVANCED COMPUTING CENTER

Texas Advanced Computing Center

- TACC is 12 years old as of last Saturday
- Serving local, statewide, and national research communities
- Leadership in NSF supercomputing, national cyberinfrastructure (XSEDE, iPlant)
- Expanding data management and analysis offerings for last ~5 years
- Stampede - #7 on TOP500 list
- World-Class visualization facilities

Data @ TACC

- Data Intensive Computing Division
 - Data Management and Collections
 - Data Mining and Statistics
- Broadly, we provide resources, services and expertise:
 - “Corral”, 6PB disk-based storage
 - Experimental Hadoop-based platform
 - “Ranch”, 100+PB Tape Archive
 - Training, consultation, and collaboration

Data Management and Collections

- Assist researchers in managing, sharing and preserving data collections
- “Data” is anything generated in the course of research activities
- Spreadsheets, simulation output, sensor output, statistical analyses, visualization, etc
- “Management” is having a plan for this data before you start creating it

Data Management w/o Plans

- Until recently, data management planning often happened after the fact
- Common case: long-running projects with:
 - A few terabytes of data
 - Thousands to millions of files
 - Data on a few servers, external disks, laptops
 - 0 to 1 people who know what and where it is
- Adapt to new, organized workflow on the fly

Data Management Planning

- More recently, researchers want to do this ahead of time
- TACC Collaboration with UT Library and UT ITS:
 - Central resource website w/ DMP template
 - Route researchers to appropriate information and organizational unit
 - Publications handled by library
 - Large datasets handled by TACC

Data Management Training

- Provided through XSEDE and UT System
- Several courses, 1-3 hours:
 - Data Management Planning and Execution
 - Storage Systems and Networks
 - Metadata creation and management
 - Using Corral
- Will soon begin offering full-day workshops

Storage Resources

- Historically, very few dedicated “persistent online storage” systems in NSF ecosystem
- Corral first deployed 4+ years ago as 1PB resource for data collections
- Restricted to UT Austin researchers but still fully utilized
- More recent UT Data Repository expansion supports more data, wider user base

Corral Deployment Model

- Persistent system with changing collection of technology components
- Have completed one full hardware transition (1PB to 5PB) plus expansion (5PB to >6PB)
- Add specific hardware to meet varied needs (e.g. SSD arrays for large databases)
- Eventual goal to make all technology migration transparent to data users

UT Research Cyberinfrastructure

- University of Texas System assessment of needs for research cyberinfrastructure:
 - Research data storage
 - High-Performance Computing
 - Dedicated 10 gigabit networking for research
- \$23 Million total investment in 2011
 - \$15 Million for network infrastructure
 - \$4 Million for research data infrastructure

The University of Texas System

- 15 Total Institutions
 - 9 “Academic”, i.e. non-medical
 - 6 Health Institutions
- Most prominently:
 - UT Austin
 - UT Southwestern Medical School
 - MD Anderson Cancer Center
 - UT Health (Houston)

UT Data Repository Requirements

- 5 Petabyte usable capacity
- Geographically replicated
- Open to all UT System researchers
- Scalable in performance and capacity
- Support for sharing data outside UT System
- Support for data management functions

UTDR Implementation: Corral

- 2 Installations (Austin and Arlington):
 - 12 Dell Servers w/ Infiniband & 10GigE
 - 2 Data Direct Networks storage controllers
 - 5PB raw capacity 3TB SATA drives
 - 350TB raw capacity 600GB SAS drives
 - 20GB/sec measured performance
- IBM General Parallel File System
- iRODS Data management software

Corral Software Services

- Static and Dynamic web services:
 - File Hosting
 - Java/PHP/Python web applications
 - iRODS-backed web applications
- Database applications
 - Postgres and MySQL
 - MonetDB
 - Statistical analysis with SAS/R/Matlab often tied to large databases

iRODS Data Management

- Provides virtual hierarchy overlaying multiple file systems and the tape archive
- Supports collections with metadata and data-sharing needs
- UNIX Command-line interface, WebDAV, Web browser application, many APIs
- Supports high-performance data transfer (100s of MB/sec) where clients support it

Data Management Functions

- Development support for structured data, data management tools, metadata, web
- Hosting capabilities for all of the above
- Additional planning support for proposals
- Up to 5 year commitments for data storage
- Not yet truly “long-term” preservation

UTDR Allocation Model

- 5TB “free” for UT Principal Investigators
- Web-based account signup/allocation request
- Human review for allocations at present
- Plan to automate approvals using InCommon
- Larger allocations on a \$250/TB/year basis
- Allocations often coupled with HPC or analysis resource access

UTDR Sustainability Plan

- Expectation of ongoing subsidy from UT for “free” tier of storage
- Recharge tier provides for ongoing expansion/ replacement of hardware for paying users
- Subsidy not required for continued operation

Corral Adoption/Usage

- Over 100 PIs within first year of operation
- Data growth about 10% per month
- ~10% of PIs account for ~50% of usage
- Usage in almost all disciplines/departments
- Major adoption from genome researchers
 - Up to 80% of new requests in some months

Future Plans for Corral

- Implement true “long-term” preservation via Digital Preservation Network
- Expand to third site for increased replication
- Support HIPAA, other high-security data types
- Increase support of web-based data sharing

Overview of Application Space

- Many projects begin with simple need for storage of large data collections
- Storage leads to further sharing and organization
- Proximity to HPC provides for new kinds of analysis, including automated workflows
- Metadata, web development provide new sharing capabilities

Applications: Genome Sequencing

- Currently up to 80% of new data collections are DNA-related data
- Wide array of processing and analysis applications
- Multi-Terabyte datasets now common, multi-petabyte datasets not unheard of
- Significant need for data-sharing capabilities
- Comparative analyses produce major results

Applications: fMRI

- Several active collaborations in this area
- Automated ingest, processing of fMRI images
- Hosting of public data-sharing websites for fMRI data
- Collaboration on wider distribution and linkage of fMRI data across consortia
- Data collection sizes from 5-50TB

Detailed Example: ICA

- Institute for Classical Archeology
 - 2 major archeological digs with multi-decade lifespans
 - Diverse collection of geo-referenced notes, images, other digital artifacts
 - No consistent organizational scheme
 - Need to create metadata, GIS database, image archive, and linking website
 - Goal: Documentation and “processual archeology”

TACC and ICA

- Initial consultation in 2008
- 2 Initial tasks: understanding workflow and ingesting as much data as possible
- Development of standard metadata, data organization scheme, file-naming conventions
- Support for multiple websites hosting data collections for specific groups of locations
- Massive improvement in overall data management and dissemination capability

Other Major Uses

- Natural History Collections – database, web, and media sharing
- Sharing of simulation output – engineering, physics, astronomy, chemistry
- Social Science data analysis – large structured datasets, often from government sources
- “Smart Grid” applications – massive structured data inputs, analysis and simulation

Data Area Research Interests

- Benchmarking of large-scale storage systems – Object stores/HDD/SSD/etc
- Benchmarking of SQL and No-SQL solutions
- Tracking of data collections infrastructure and curation needs
- Development of automated curation workflows
- Reliable, high-performance data transfer

TACC Data futures

- 2012 O’Donnell Foundation grant
 - \$10 Million to expand data-intensive computing@TACC
 - Expect VM/service hosting cluster, parallel file system, and “data analysis” components
 - File system component to be announced in June
- 100-Gigabit network expansion
- Expanding staffing across Data division

Observations on “Big Data”

- Not all the hype is true
- But, a lot of it is...
 - Data Volume and Velocity is increasing across all disciplines, in some cases exponentially
 - New capabilities for analysis/visualization/understanding of data are being developed
 - Sharing of data across research projects enables significant new science to be done

Q & A

- Thanks to the Data Management and Collections Group
 - Maria Esteva
 - David Walling
 - Tomislav Urban
 - Siva Kulasekaran