

# Collaborative Data Intensive Science

Arun Jagatheesan

San Diego Supercomputer Center


and

[iRODS.org](http://iRODS.org) / [DiceResearch.org](http://DiceResearch.org)

# Agenda (10 min!)

- Use case: LSST
- Collaborative Data-life cycle Management
  - Scale-up and Scale-out
- Current efforts
  - DASH, iRODS
- We need more
  - Data I/O protocols with control channels
  - Storage Time Machine (if there is time for this)
- Q&A


# How many of you know what is LSST?



# LSST

- Large Synoptic Survey Telescope (LSST)
  - Survey entire sky every 3 nights
  - Dark Energy, Dark Matter, Near Earth Asteroids, ...
  - Largest digital camera in the world (3 billion pixels)
  - Images 3000 times wider than Hubble
- LSST Data Management
  - Data from Chile to US and rest of the world
  - 15 TB/night, over hundred(s) petabytes
  - Multiple data centers around the world
  - Trillions of rows database ( $\sim 15$  PB)
  - Hundreds of millions of files ( $\sim 80 \times 3 = \sim 240$  PB)

# LSST current sites



## Site Roles and their Functions

- **Base Facility**  
Real-time Processing and Alert Generation,  
Long-term storage (copy 1)
- **Archive Center**  
Nightly Reprocessing, Data Release  
Processing, Long-term Storage (copy 2)
- **Data Access Centers (DACs)**  
Data Access and User Services
- **System Operations Center (SOC)**  
System Supervisory Monitoring Control  
& End User Support/Help Desk

\* Co-located DAC: shares infrastructure with Archive Center

\*\* Co-located DAC: shares infrastructure with Base Facility

# LSST and CDLM

QuickTime™ and a  
TIFF (Uncompressed) decompressor  
are needed to see this picture.




QuickTime™ and a  
TIFF (Uncompressed) decompressor  
are needed to see this picture.



QuickTime™ and a  
TIFF (Uncompressed) decompressor  
are needed to see this picture.

QuickTime™ and a  
TIFF (Uncompressed) decompressor  
are needed to see this picture.



# LSST and CDLM

QuickTime™ and a  
TIFF (Uncompressed) decompressor  
are needed to see this picture.



QuickTime™ and a  
TIFF (Uncompressed) decompressor  
are needed to see this picture.



QuickTime™ and a  
TIFF (Uncompressed) decompressor  
are needed to see this picture.

\\i\exp\file1.fits

\\i\exp\file2.fits

/euro/exp/file2.fits

QuickTime™ and a  
TIFF (Uncompressed) decompressor  
are needed to see this picture.



/res/chile/exp/file1.fits

/exp/file1.fits

/exp/file2.fits

# Topic and current problems (related to this talk)

- Collaborative Data-lifecycle Management
  - “Data by itself is a process”
  - Data has to be social and “collaborate” with many including producer(s), consumer(s)
- Scale-out
  - Data Grid or Data Cloud or ?
  - iRODS.org
- Scale-up
  - IO latency (CPU cycle >>>> IO cycle)
  - SDSC DASH



# iRODS: Logical File System

Scale out to multiple data centers

- iRODS
  - Data Grid Management System for Digital Libraries, Persistent Archives and Data Grids
  - Open Source BSD
  - Version 2.1

# SDSC DASH

(one small step for byte,  
one giant leap for a petabyte)

- Prototype effort for data intensive computer
  - Scale-up is EXPENSIVE (supercomputer)
  - Reduce IO latency with more memory (cheap) and SSD
- vSMP node
  - Aggregate multiple nodes into a single powerful node using software : Global memory as commodity
- SSD
  - 4TB of SSD
  - 3 IO nodes

# If I had a billion bucks...

- IO latency
  - Smarter storage with CPU attached (just for storage control) and new protocols that can get control messages about h/w at a very low-level.
- Inter-processor and Inter-data center IO
  - IO for scale-up and scale-out
  - Improvements in CPU or data management software are handling the symptoms rather than the cause
- Data to Knowledge Communities
  - Data, Information, Knowledge
  - People, Communities

# Storage Time Machine

- Capacity : Infinite
- I/O latency: Almost None
- Persistence of data: 10,000 years ++;
- TCO : Almost Zero
- Scalability: Few exabytes
- Start-up time: TBA (its ok don't need to perfect)

# Agenda (10 min!)

- Use case: LSST
- Collaborative Data-life cycle Management
  - Scale-up and Scale-out
- Current efforts
  - DASH, iRODS
- We need more
  - Data I/O protocols with control channels
  - Storage Time Machine (if there is time for this)
- Q&A