Catalog Archive Server

Ani Thakar, JHU

For the CAS Team

Outline

- CAS Team
- CAS Intro
- Data Storage
- Data Loading
- Data Access
- CAS for SDSS-IV
- SciServer DIBBS project
- SDSS-IV website

CAS Team

Name	Role	~ FTE
Deoyani Nandrekar-Heinis	Software developer (ImgCutout, SkyServer, JPEG generation, VO services)	0.5
Dmitry Medvedev	Software developer (CasJobs, SkyServer)	0.5
Sue Werner	Software developer (DB tuning, partitioning)	0.25
Victor Paul	Database administrator, storage management	0.5
Jordan Raddick	Documentation lead, press officer, website lead	0.5
Bonnie Souter	Website developer (SDSS.org, SkyServer)	0.5
Rich Ercolani	IT support lead	0.25
Alex Szalay	Oversight, DIBBs (SciServer) PI	0.05
Ani Thakar	Management, data loading, logging (SDSS/VO), SciServer integration	0.25

CAS Intro

- History
  - Original CAS an OODBMS
  - Migrated to SQL Server in 2001
    - The SkyServer was born in 2001
- Design based on analysis of logs
  - 3 main types of users:
    - Lots of quick queries
    - A few "power" users
    - Visual data browsers (astro)
  - Power users slowed everyone down
    - Need to segregate query work-loads on separate servers
  - Increasing SQL usage
- Minimizing data movement
  - Bring the analysis to the data, not vice versa

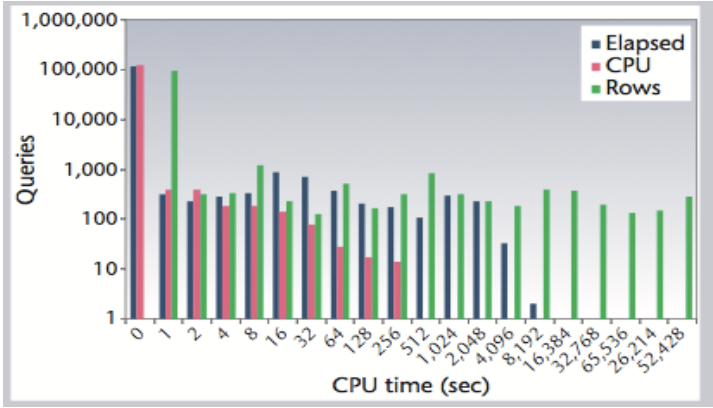

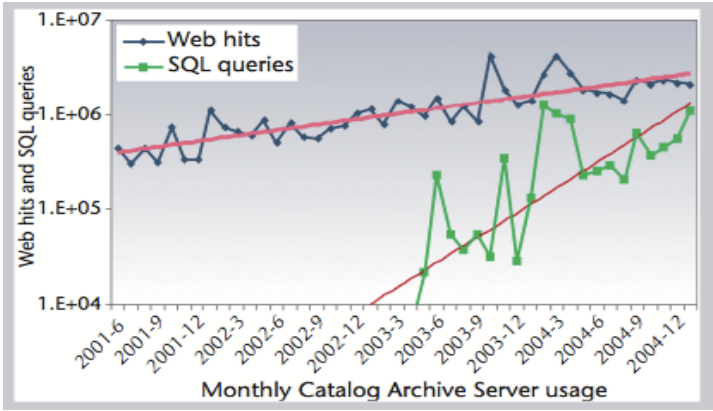

CAS Design

- Based on relational DBMS: MS SQL Server
- Layer of science schema built right into the DB
  - Extensive use of UDFs/SPs
  - HTM spatial index in C# CLR
- 3 pillars of data access
  - Synchronous: SkyServer
  - Asynchronous: CasJobs
  - Visual: ImgCutout
- sqlLoader data loading pipeline

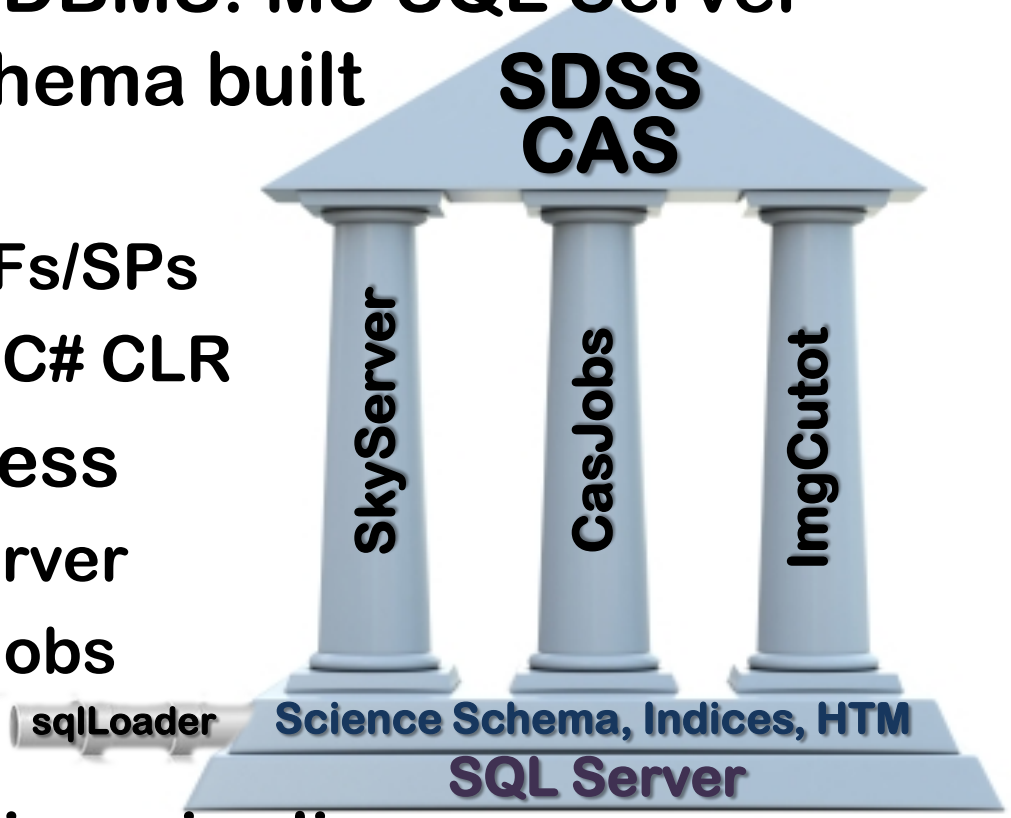

Reusable Building Blocks

- SkyServer
  - Extensive built-in science, query, metadata support
- CasJobs batch query workbench
  - Adapted and deployed in several (non) astro projects
- ImgCutout visual JPEG browsing service
  - Recently adapted to display 2MASS JPEGs
- sqlLoader data loading pipeline
- Hierarchical Triangular Mesh spatial index
  - CLR library written in C#, ported to Java, C++
- All downloadable from skyserver.org
- All MS SQL Server based, at present

The SDSS-CAS Genealogy

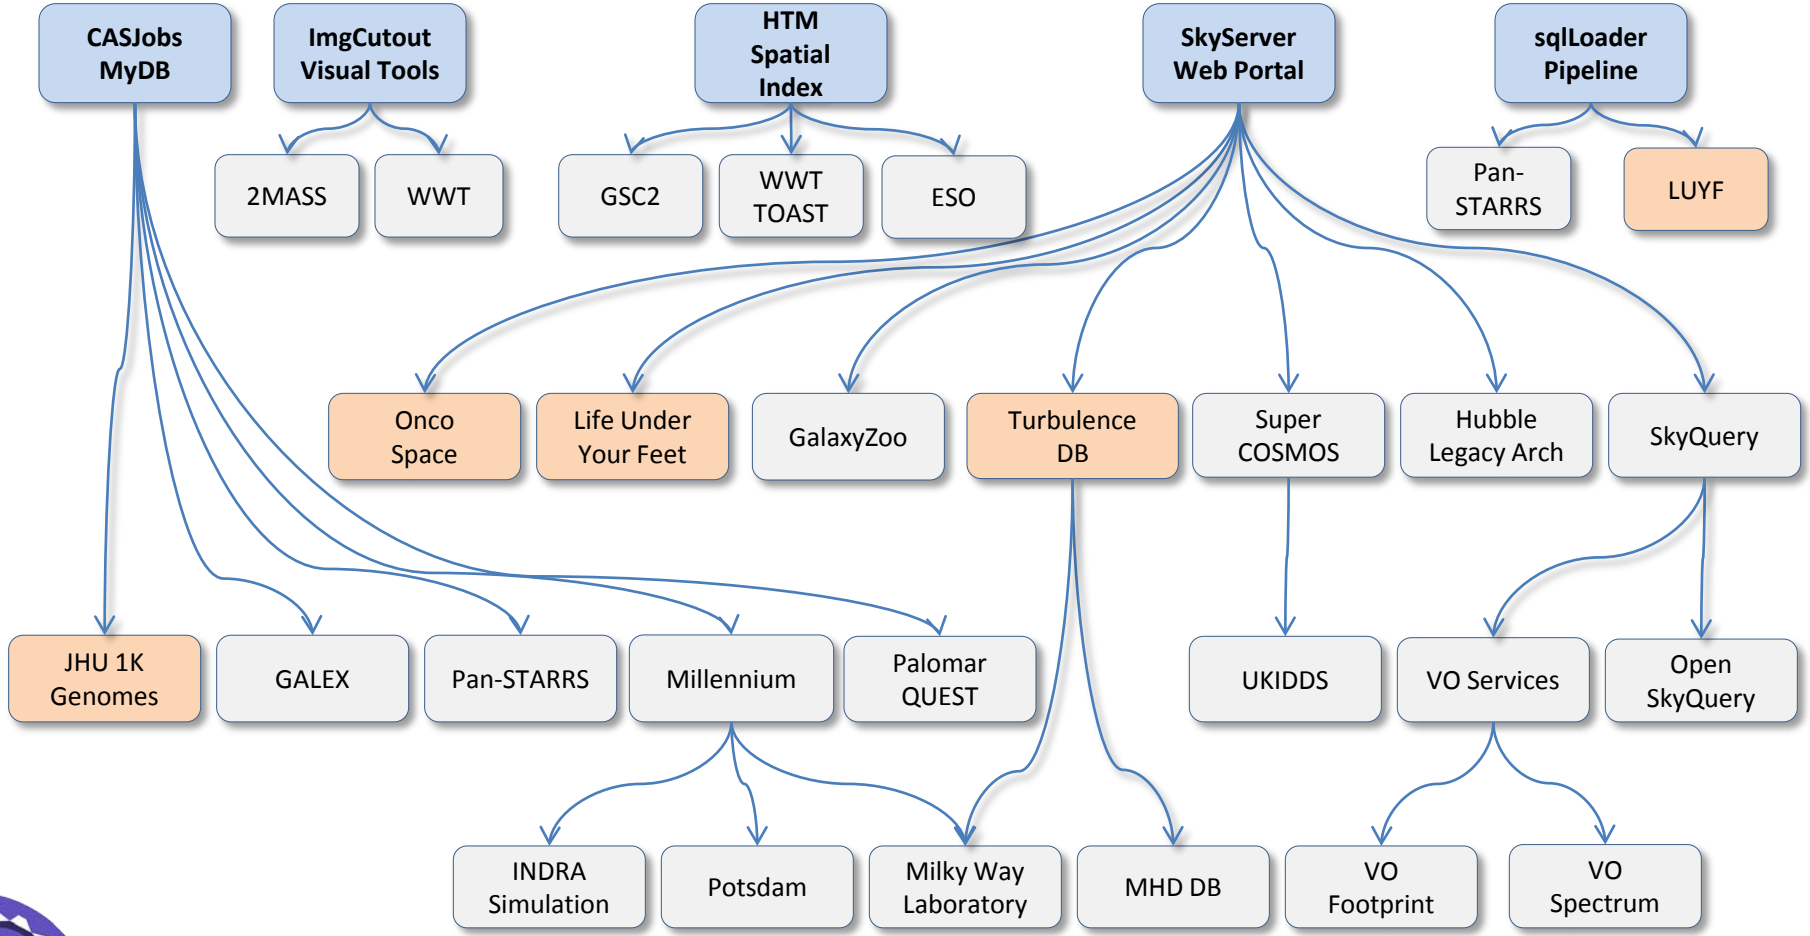

SkyServer Web Interface

- The public portal to CAS data since 2001
- Supports several levels of user acces
  - Simple to complex form queries
  - CrossID search with upload capability
  - Visual browsing of individual objects
  - Raw (filtered) SQL query
- Includes client for ImgCutout service
  - Finding Chart page
  - Google Maps-style Navigate page
  - Queryable Image List page for multiple cutouts at a time
- Schema Browser, extensive SQL help
- Virtual Observatory services (VO standards/protocols)
- Rich educational projects section (K-12+)
- Every web hit and SQL query recorded in logs

SkyServer Usage Logging

- All web hits and queries logged since day 1 (2001)
- SkyServer traffic page shows up-to-the-hour logs
  - 1.43 billion hits, 263 million SQL queries to date
  - Currently averaging 15M hits and 1.5M queries/mth
- Logging overview document at skyserver.org/doc
- 3 published papers on SkyServer traffic:
  - "SkyServer Traffic Report – The First Five Years", MS Technical Report (Singh et al. 2006)
  - "Ten Years of SkyServer – Tracking Web and SQL e-Science Usage", CiSE (Raddick et al. 2014)
  - "Ten Years of SkyServer – How Astronomers and the Public Have Embraced e-Science", CiSE (Raddick et al. 2014)

SkyServer Monthly Web Hits and SQL Queries Usage

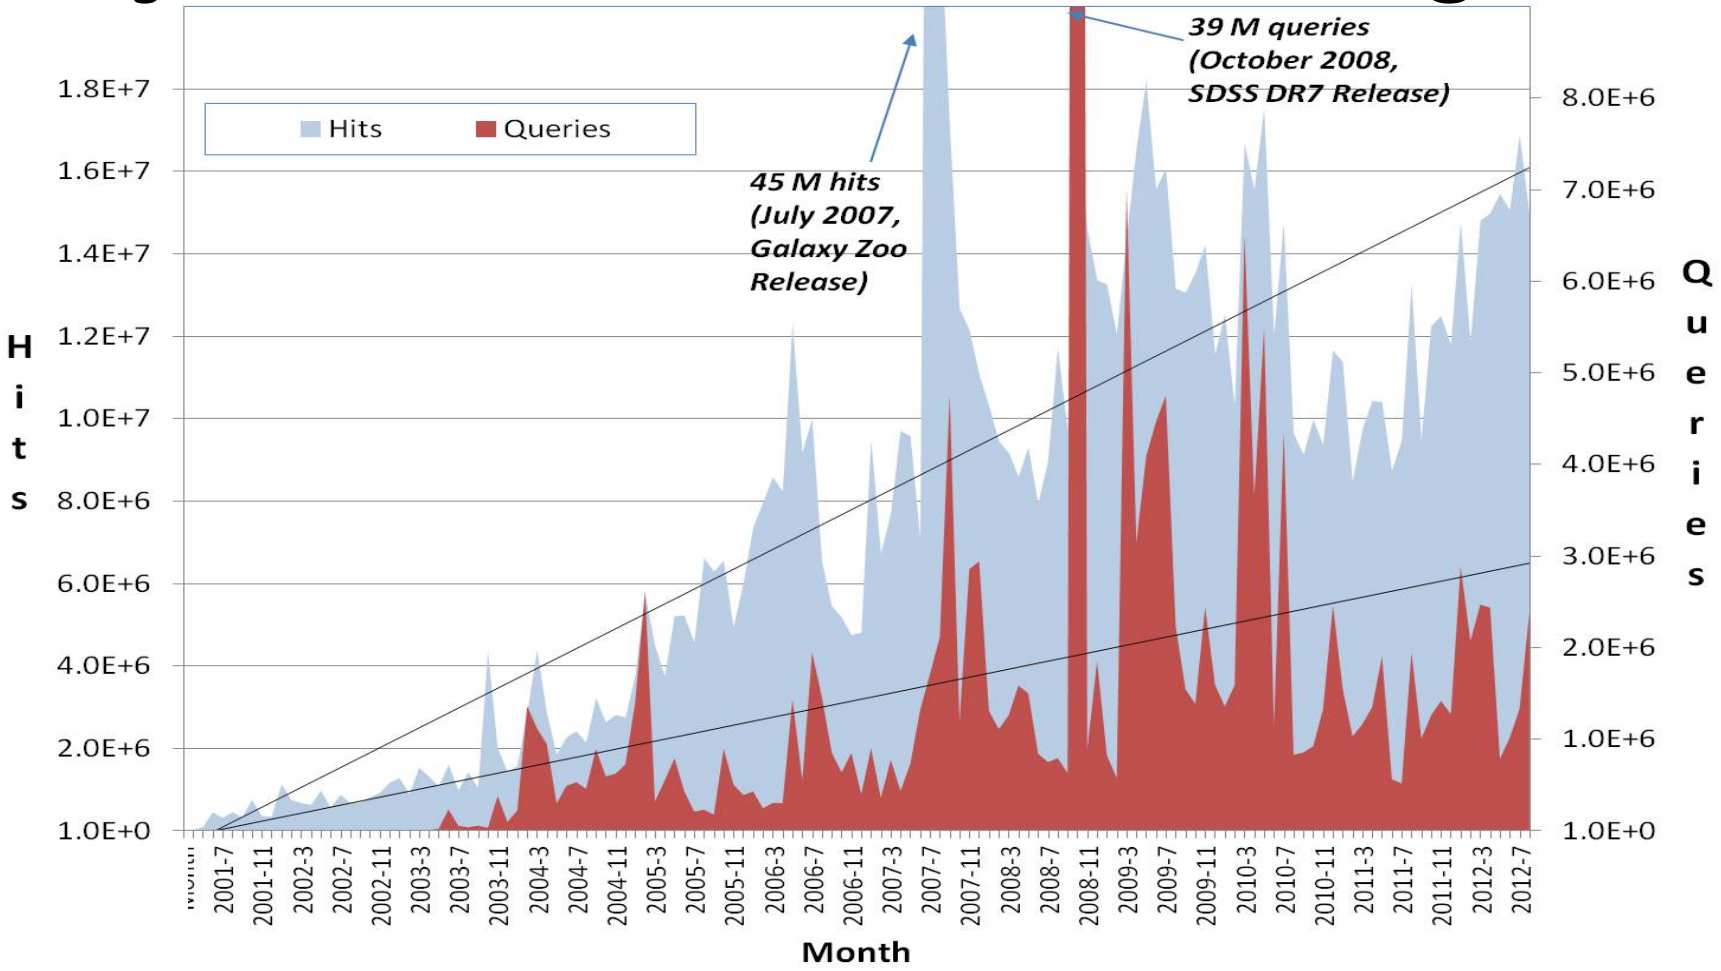

CasJobs

- Batch query workbench (launch 8/2003)
- Web application + web service backend
  - ASP.NET/C# development platform
- Workhorse of CAS data access
  - SDSS-II CasJobs: 9300 users, 6.1M jobs
  - SDSS-III CasJobs: 3800 users, 8.8M jobs
- Every user has their own SQL "MyDB"
  - Default size 0.5 GB, increased on request
- Quick (1min sync) & Long (8hr async) modes
- Complete searchable job history
- Schema browser, MyDB table browser
- Data Import, Groups feature to share data

ImgCutout

- ASP.NET/C# web service
- Client: SkyServer Visual Tools
  - Finding Chart
  - Navigate (Google Maps-style interface)
  - Image List upload with user query input
  - Explore and Quick Look object browsing
- JPEGs served from Frame table in DB
  - 3-color JPEGs generated from FITS

Data Storage

- Multiple instances of each DR
  - For redundancy, load-balancing and performance (workload segregation)
  - As many as 6 copies of most active DR!
    - For smooth ops and optimal performance
    - SkyServer, Quick CasJobs, Long Public CasJobs, Long Collab CasJobs, Imgcutout and development/backup/restore copy
- Currently ~ 120 TB of DR8-11 DBs

CAS Cluster for DR8-12

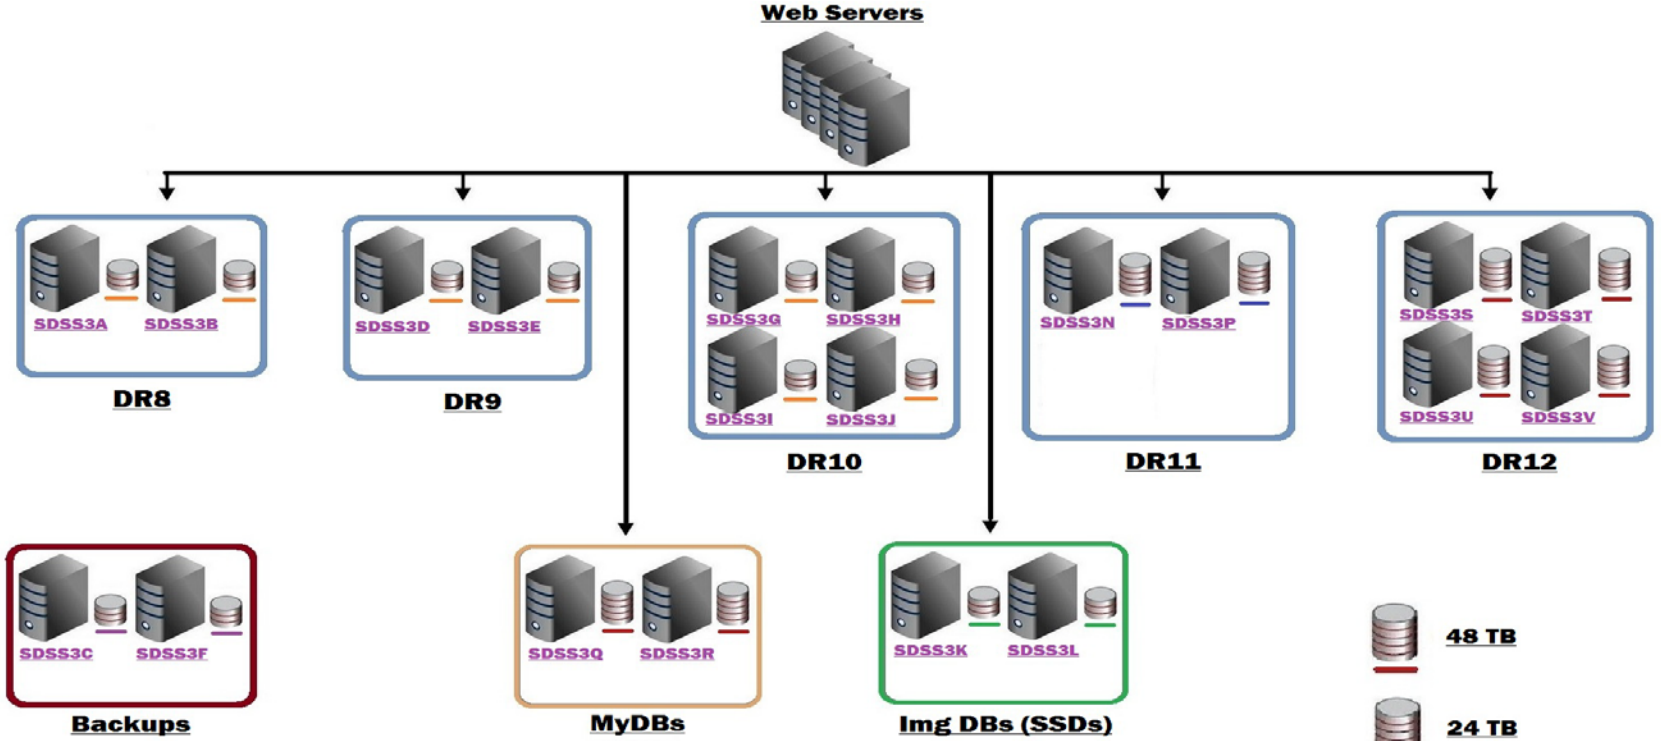

Server	CPU	RAM	Space
SDSS3A,3B,3D,3E	Intel Xeon E5440 8 core (2x4) 2.83 GHz	16 GB	16 TB
SDSS3C,3F	Intel Xeon E5440 8-Core (2x4) 2.83 GHz	16 GB	8 TB
SDSS3G,3H,3I,3J	AMD Opteron 6134 16-Core (2x8) 2.30 GHz	32 GB	16 TB
SDSS3K,3L	Intel Xeon E5-2630 24-Core (2x12) 2.30 GHz	64 GB	1 TB (SSD)
SDSS3N,3P	Intel Xeon E5-2630 24-Core (2x12) 2.30 GHz	64 GB	24 TB
SDSS3Q,3R	Intel Xeon E5-2630 v2 24-Core (2x12) 2.60 GHz	64 GB	48 TB
SDSS3S,3T,3U,3V	Intel Xeon E5-2630 v2 24-Core (2x12) 2.60 GHz	64 GB	48 TB

BestDR12 Schema

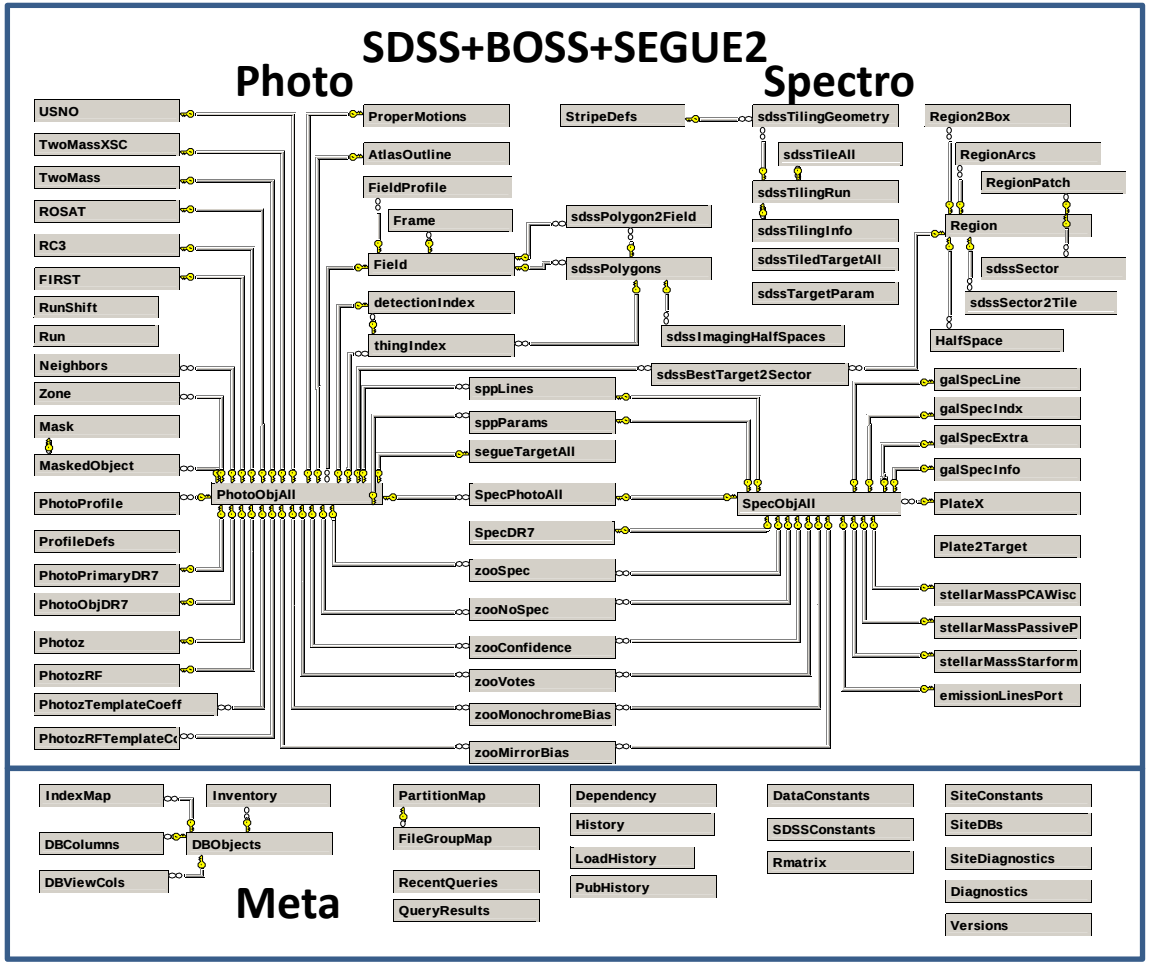

APOGEE

- apogeeField
- apogeeObject
- apogeeDesign
- apogeePlate
- apogeeStar
- apogeeVisit
- apogeeStarVisit
- apogeeStarAllVisit
- aspcapStar
- aspcapStarCovar

MARVELS

- marvelsStar
- marvelsVelocityCurveUF1D

WISE

- WISE_allsky
- WISE_xmatch

- Photo, Spectro largest group
- Meta tables for SB and xmatch
- APOGEE, WISE, MARVELS in own groups

Data Table Sizes

Name	Rows	Data GB	Index GB	Total GB
PhotoObjAll	1,231,051,050	3130.7	1579.1	4709.8
PhotoProfile	44,563,999,574	1137.1	6	1143.1
AtlasOutline	1,219,412,987	1072.1	3.6	1075.8
Frame	3,752,184	989.8	0.5	990.3
Neighbors	25,578,382,962	942.7	6.4	949.1
WISE_allsky	563,921,584	537.8	149.9	687.7
SpecObjAll	3,358,200	185.9	1.1	187
PhotoObjDR7	364,857,538	111.3	7.1	118.4
PhotoPrimaryDR7	305,789,541	93.3	0.4	93.8
SegueTargetAll	453,975,934	51.7	8.4	60.1
Zone	794,035,877	49.7	0.4	50
thingIndex	563,688,948	35.8	10.3	46.1
detectionIndex	932,891,133	26.5	16.9	43.4
FIRST	510,679,510	40.2	0.1	40.3
USNO	253,732,084	38.7	0.2	38.9
Photoz	208,478,448	31.8	0.2	32
ProperMotions	336,954,036	29.2	0.1	29.4
TwoMASS	103,577,231	14.4	11.7	26.1
WISE_xmatch	495,003,196	16.9	8.9	25.8
PhotozRF	208,478,448	25.2	0.1	25.4
PhotozTemplateCoeff	517,961,077	12.2	0.1	12.3
PhotozRFTemplateCoeff	505,398,590	11.9	0.1	12

- Size > 10 GB
- PhotoObjall largest table by far, most heavily indexed
- PhotoProfile has max rows but is thin and much less used

Data Loading

- Data products that go into CAS
  - Photo: SDSS, BOSS
  - Spectro: SDSS, BOSS, SEGUE,APOGEE
  - Window, Resolve, Region, Tiling etc.
  - Xmatches: USNOB, 2MASS, 2DF, ROSAT, FIRST, WISE All Sky
  - Galaxy Zoo 2 classifications
- Ingested in CSV format by sqlLoader

sqlLoader Data Ingest Pipeline

- System of SQL and VB scripts controlled by ASP client application (Load Monitor)
- Automates tedious data loading tasks
- Thoroughly checks data integrity
- Enables parallelism in data loading
- Provides complete history and log for each loading task
- Tracks statistics for loading performance

Loader Data Flow and Validation

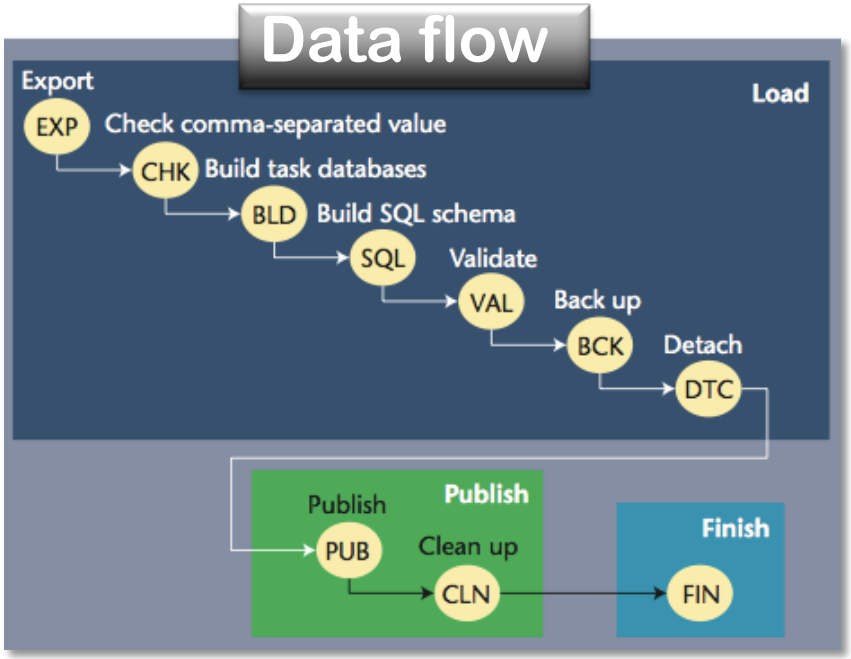

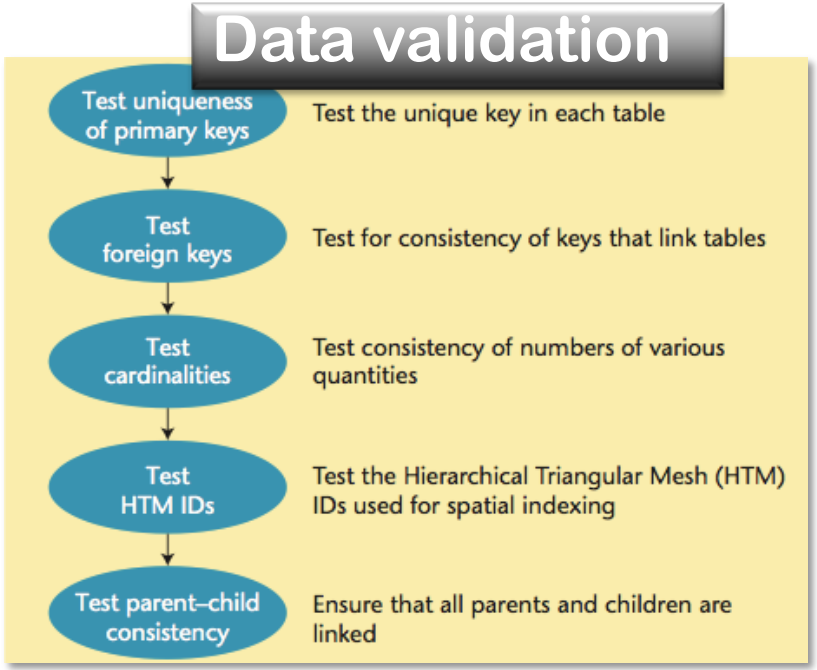

- Load-Publish-Finish main stages
- Data first loaded into temp "task" dbs
  - Can be parallelized for fast loading
- Publish stage writes data to final db
- Finish stage creates indices, computed tables

- Data correctness, consistency and integrity checks built into loading process
- Uniqueness checks, cardinality checks, relationship checks
- Have proven invaluable in the past in finding problems upstream of CAS

sqlLoader Screen Shots

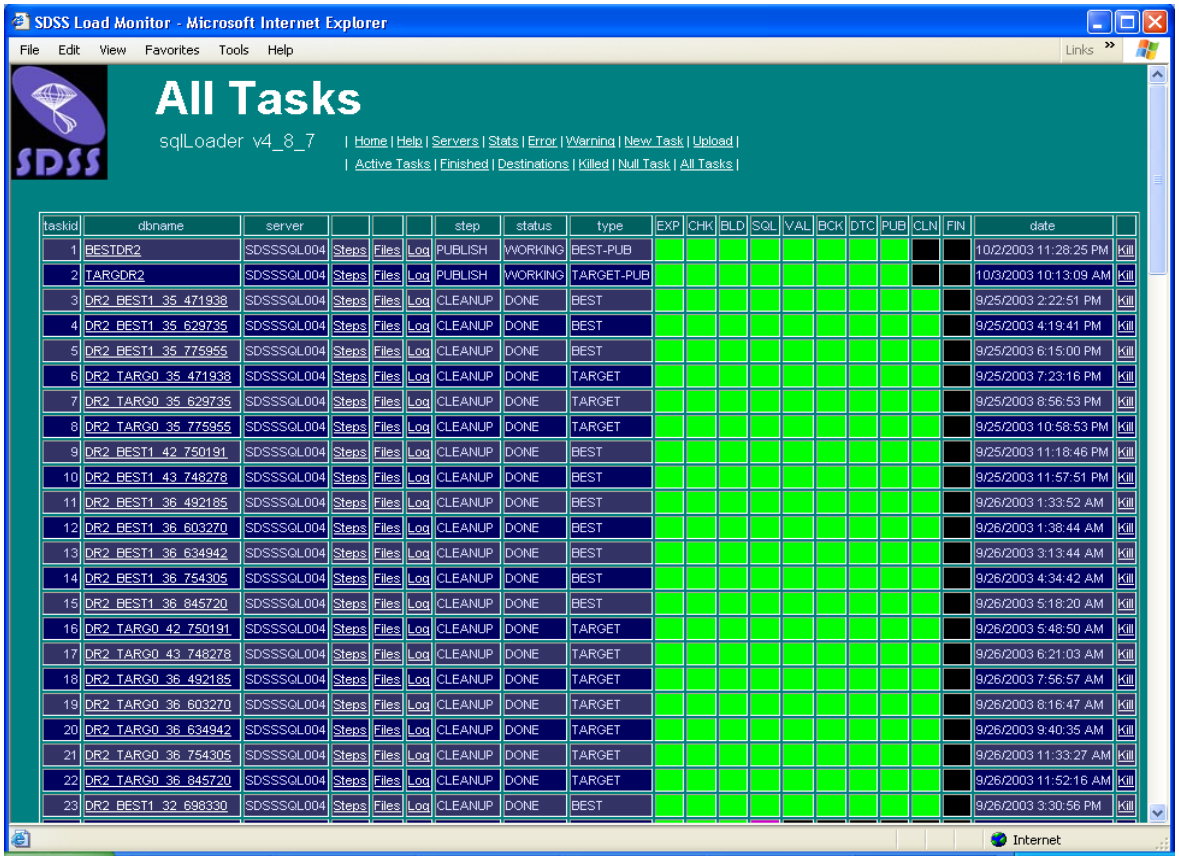

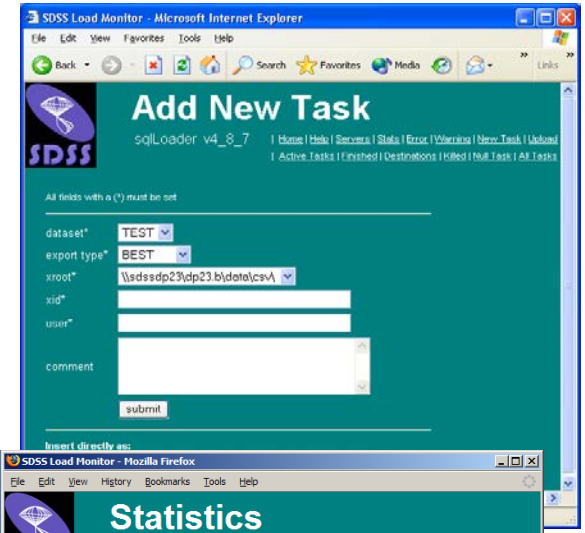

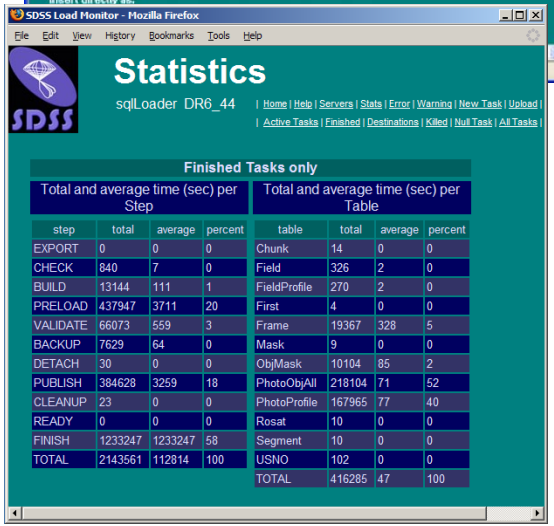

CAS Data Release Cycle

- Detailed list of steps is in overview doc
- Cycle must be repeated for each data product (photo, spectro, APOGEE etc.)
- Most steps involve multiple actions
- Testloads are usuallyy a small subset of DR data
- Entire cycle can take ~ 2 months
- Objective is to give collab a few months preview of data
- Steps in gray are only done once when all data products are in

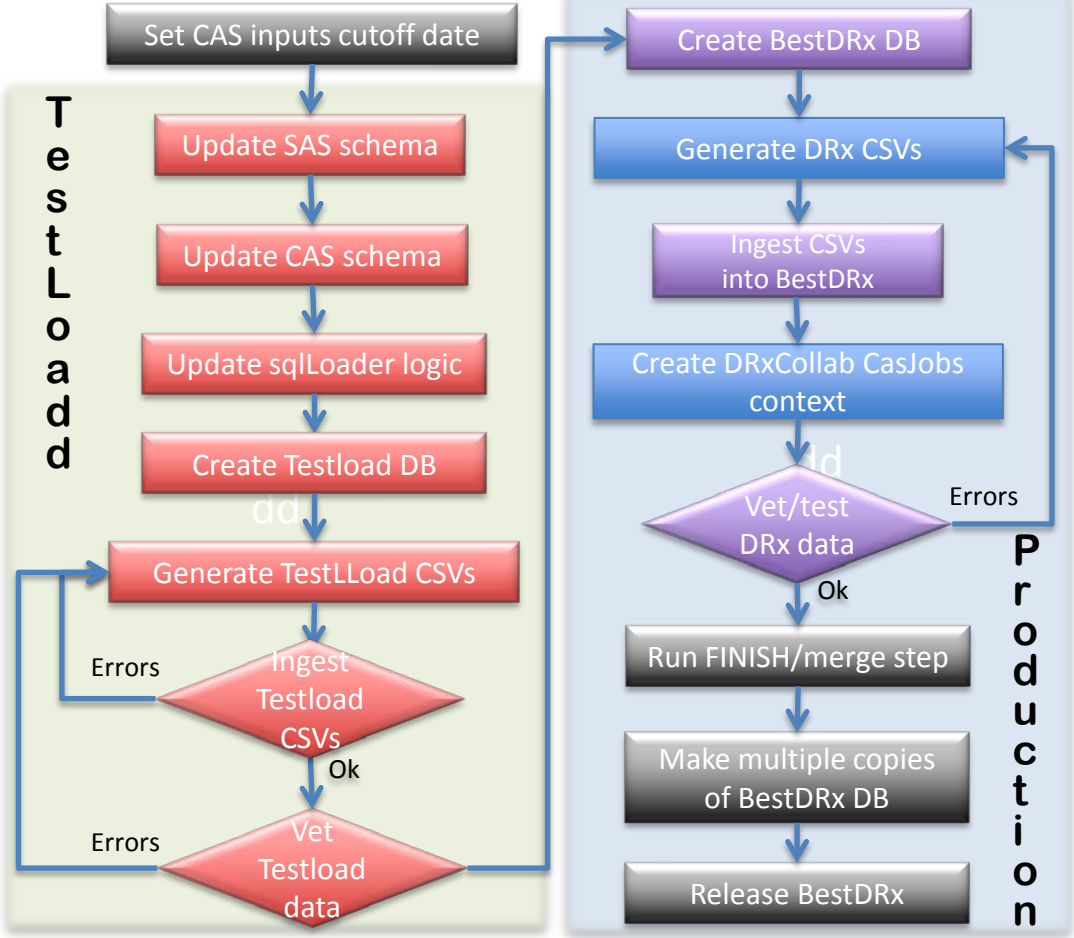

CAS Admin Site: skyserver.org

- All CAS data/software downloadable
- Data releases downloable as compressed SQL Server backups
  - Smaller subsets available for testing
- CAS mirror sites resource page
  - Mirrors can download latest updates here
- CAS/SkyServer documentation page
  - Skyserver.org/doc (data review docs here)
- SkyServer site download
- CasJobs download
- HTM spatial index doc and download

CAS Mirrors

- Official SDSS-III mirror: Brazil (LIneA)
  - Also hosts CasJobs mirror
- Other mirrors:
  - China (LAMOST)
  - Portsmouth
- SDSS-II mirrors worldwide
  - UK, Germany, Russia, China, Japan, India
- VO sites (NED, CDS)

CAS for SDSS-IV

- Continue current procedures, best practices
- Data volume increases manageable
  - Expect 2-4 TB per DR through 2020
- APOGEE-2, eBOSS data not radically different
  - Apply current tools, processes with tweaks
- MaNGA data will be biggest challenge
  - Both for complexity/novelty and size of data
  - Will need to develop new tools from scratch
- MaNGA data visualization in SkyServer
  - Tier 1 targeted for DR13 (mid-2016)
  - 0.5 FTE additional dev effort through 2018

SciServer

- JHU DIBBs project (sciserver.org)
- Generalize and extend SkyServer/VO framework
  - Reengineer/refactor SDSS data access tools for maximum reuse and extensibility
  - Port ASP SkyServer code to ASP.NET/C#
  - Convert CasJobs web services from SOAP to RESTful
  - Reengineer SkyQuery (VO service) on GrayWulf platform
  - Integrate CasJobs and SkyQuery
  - Integrate SciDrive DropBox-like (VO) service with CasJobs/MyDB
  - Keystone SSO for access to all services
- Take ownership of legacy SDSS datasets
  - Integrate multiple SDSS phases: new sdss.org website
- Additional 1 FTE development on SDSS software (thru 2015)
- Additional dev hardware and storage
- Project management 1 FTE (M Rippin)
  - Formalized software development processes

Thank you!

- All documents and links to published articles about the SDSS CAS can be found at http://skyserver.org/doc/
- Happy to take questions after Jordan Raddick's SDSS website overview

SDSS.org Integration and Rebranding

Jordan Raddick, JHU