



Catalog Archive Server

Design and Performance Overview

Ani Thakar (JHU)

CAS team at JHU



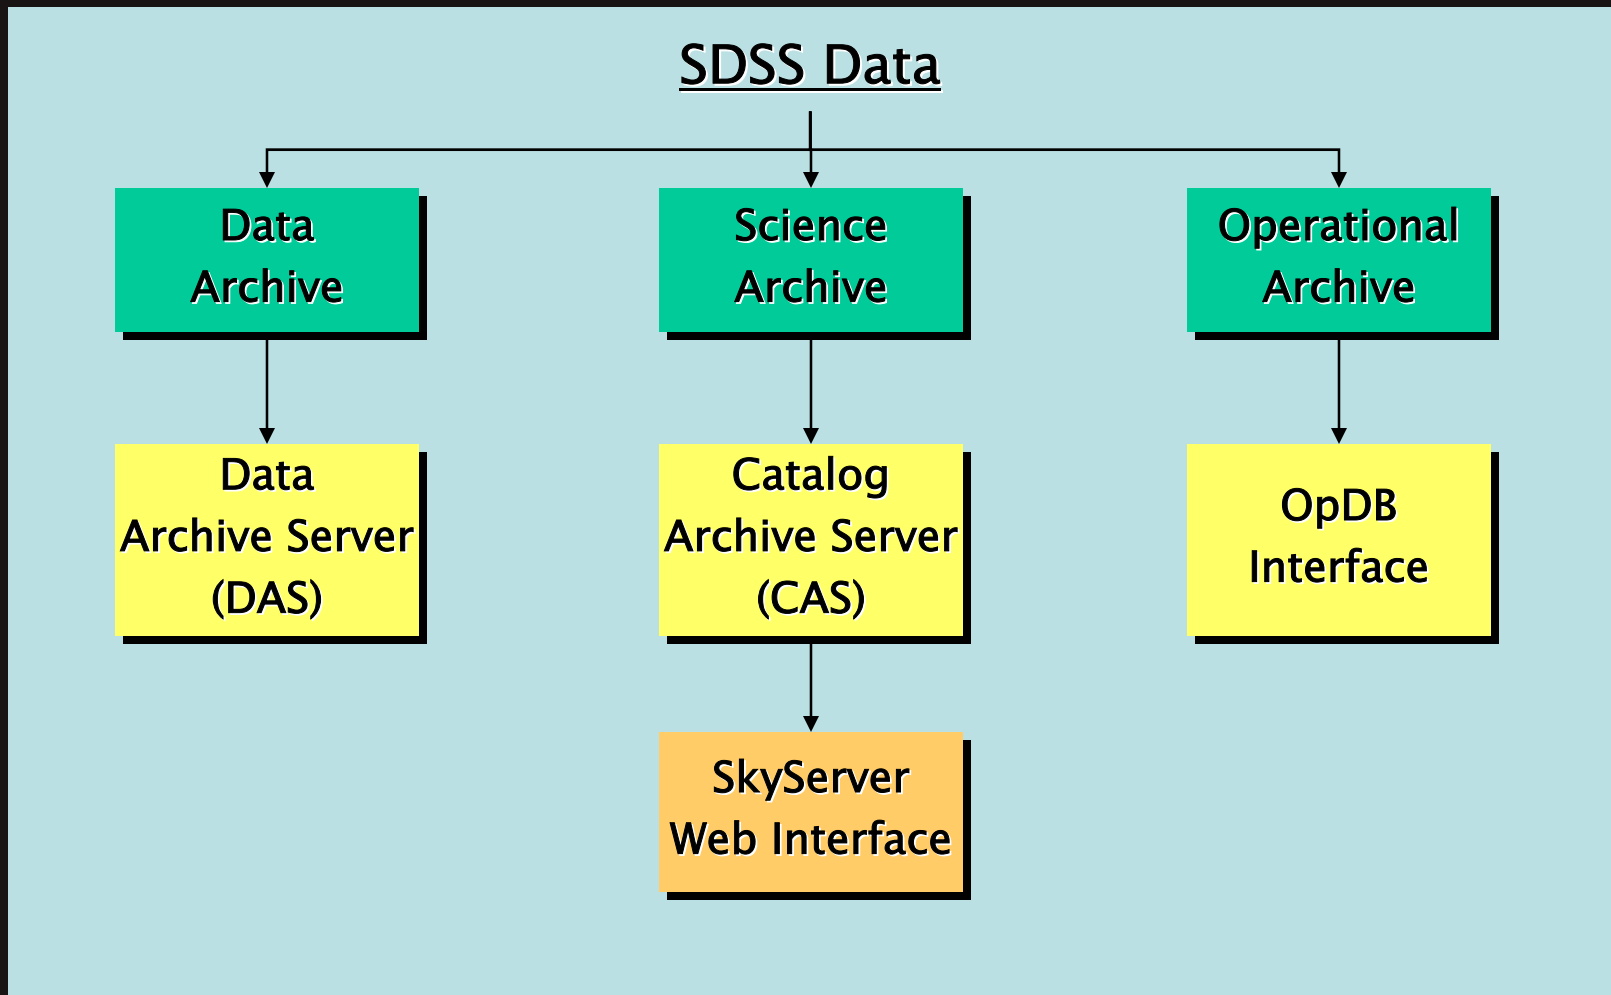
Tamas Budavári	sqlCI, Photoz
George Fekete	HTM, JPEGs
Jim Gray (MS)	SkyServer, sqlLoader
Nolan Li	sdssQA, CasJobs
Tanu Malik	sdssQA
Maria Nieto-Santisteban	ImgCutout, MySkyServer
Wil O'Mullane	CasJobs, JPEGs, SkyServer
Alex Szalay	SkyServer, sqlLoader
Ani Thakar	sqlFits2Csv, sqlLoader, SkyServer
Jan vandenBerg	Cluster maintenance, DBA

Outline



- SDSS Data Overview
- CAS Design Overview
- Basic Functional Requirements
- Data Loading and Replacement Strategy
- SkyServer Front End
- Documentation
- DR1 CAS Experience

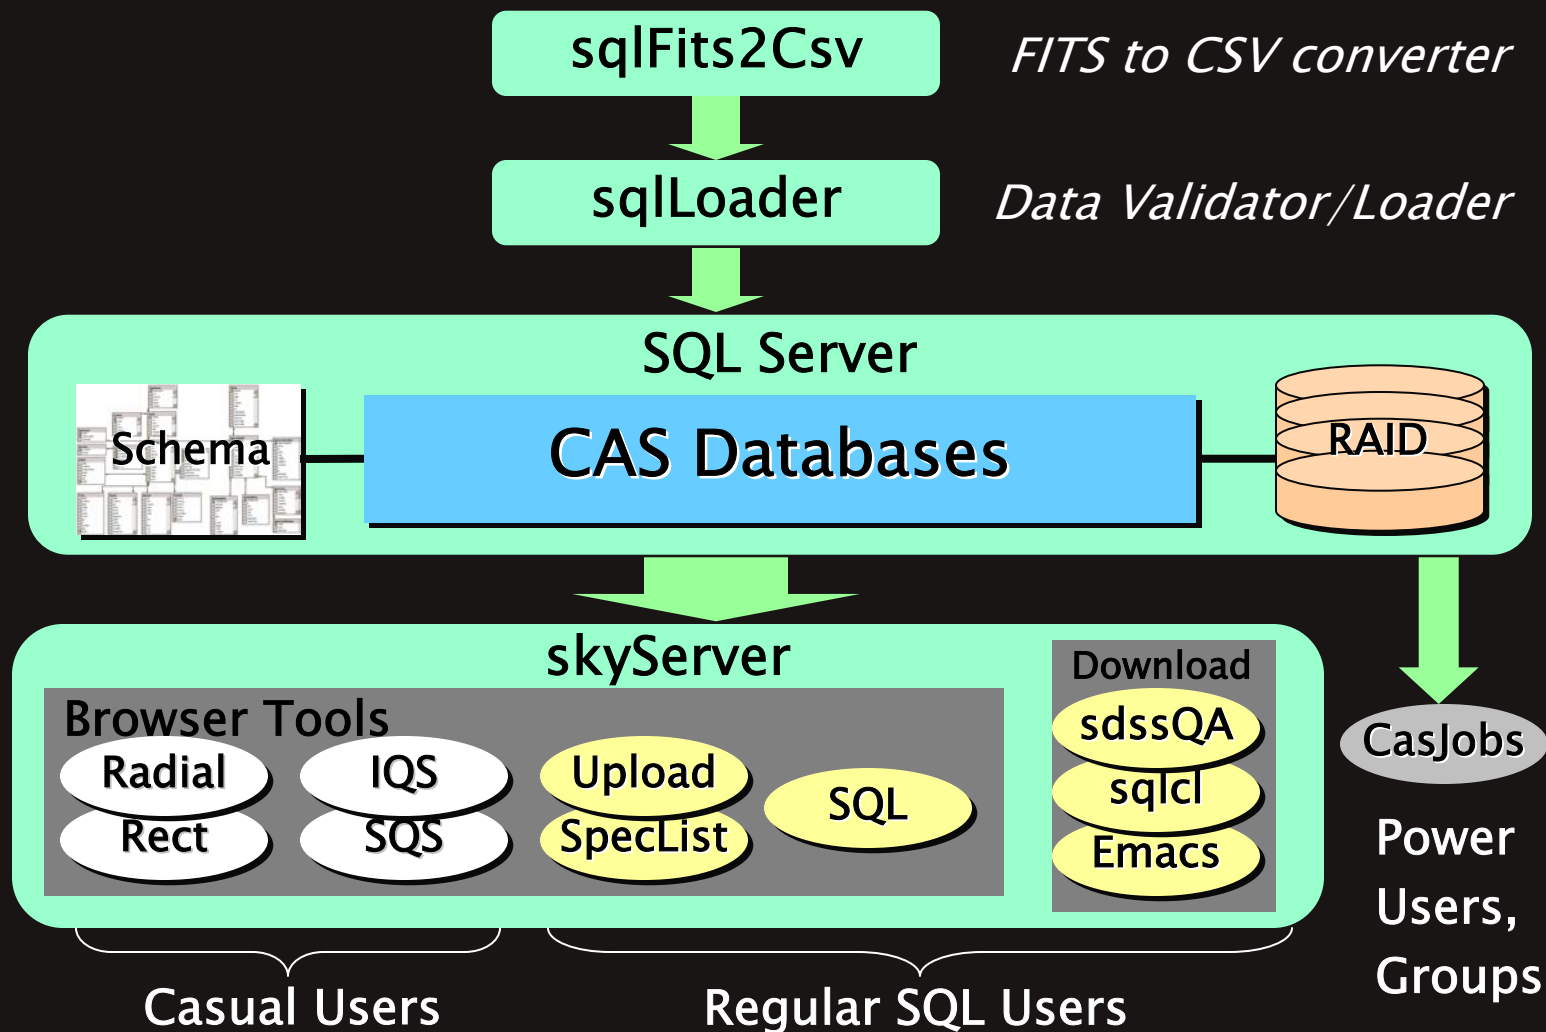
SDSS Data Overview



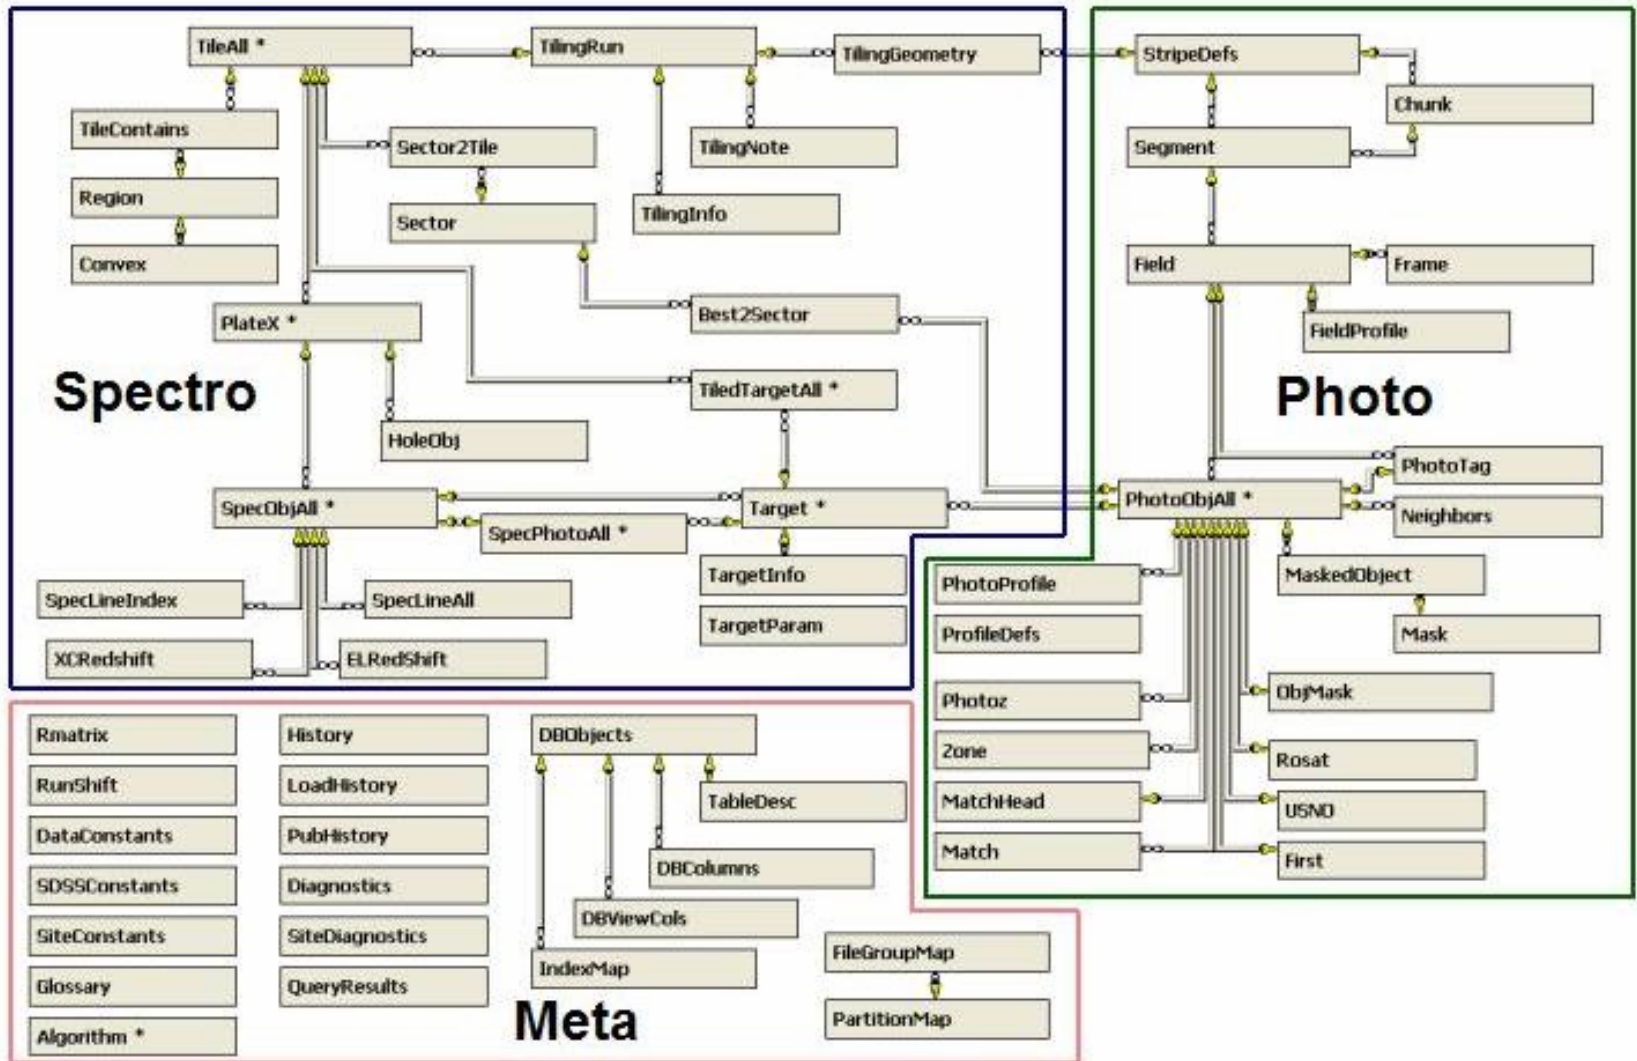
Science Archive Requirements

- Calibrated object catalogs
 - PhotoObj and associated tables
- Spectroscopic parameters
 - SpecObj and associated tables
- References to atlas images and extracted spectra
 - Atlas outlines and specline information in catalog
- Ability for manual target selection
- Ability to extract subsets of data
 - Size of subsets limited by network bandwidth
- Smooth transition to public distribution
 - Online access, mirror sites

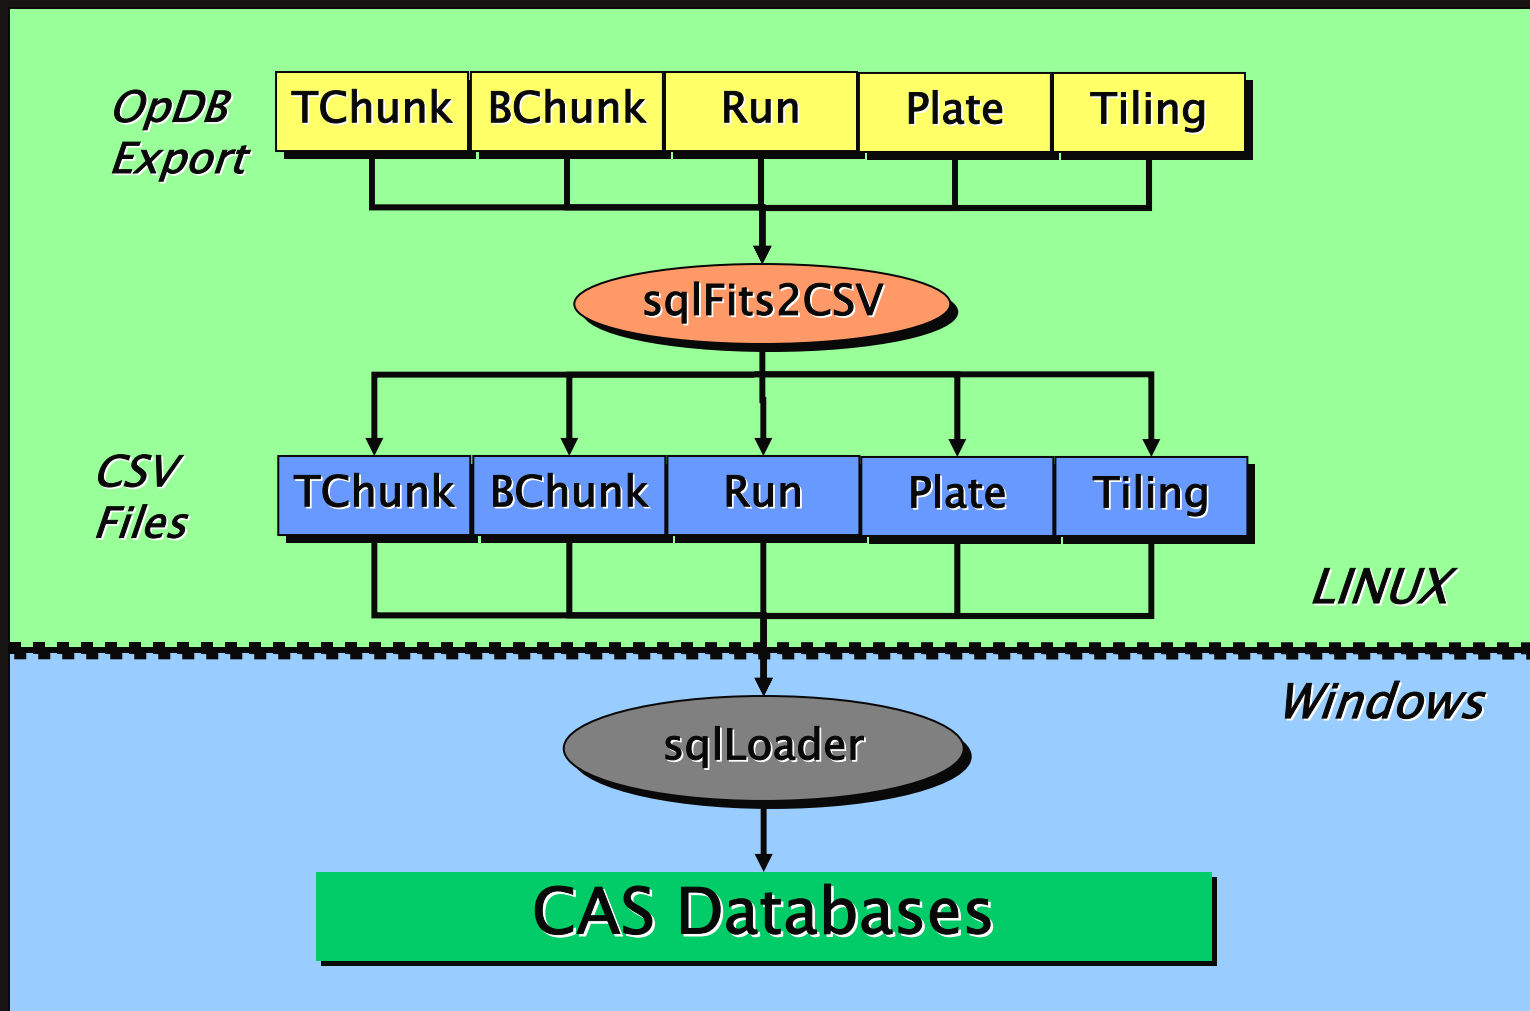
Design Overview



CAS Data Model



Loading Data Flow

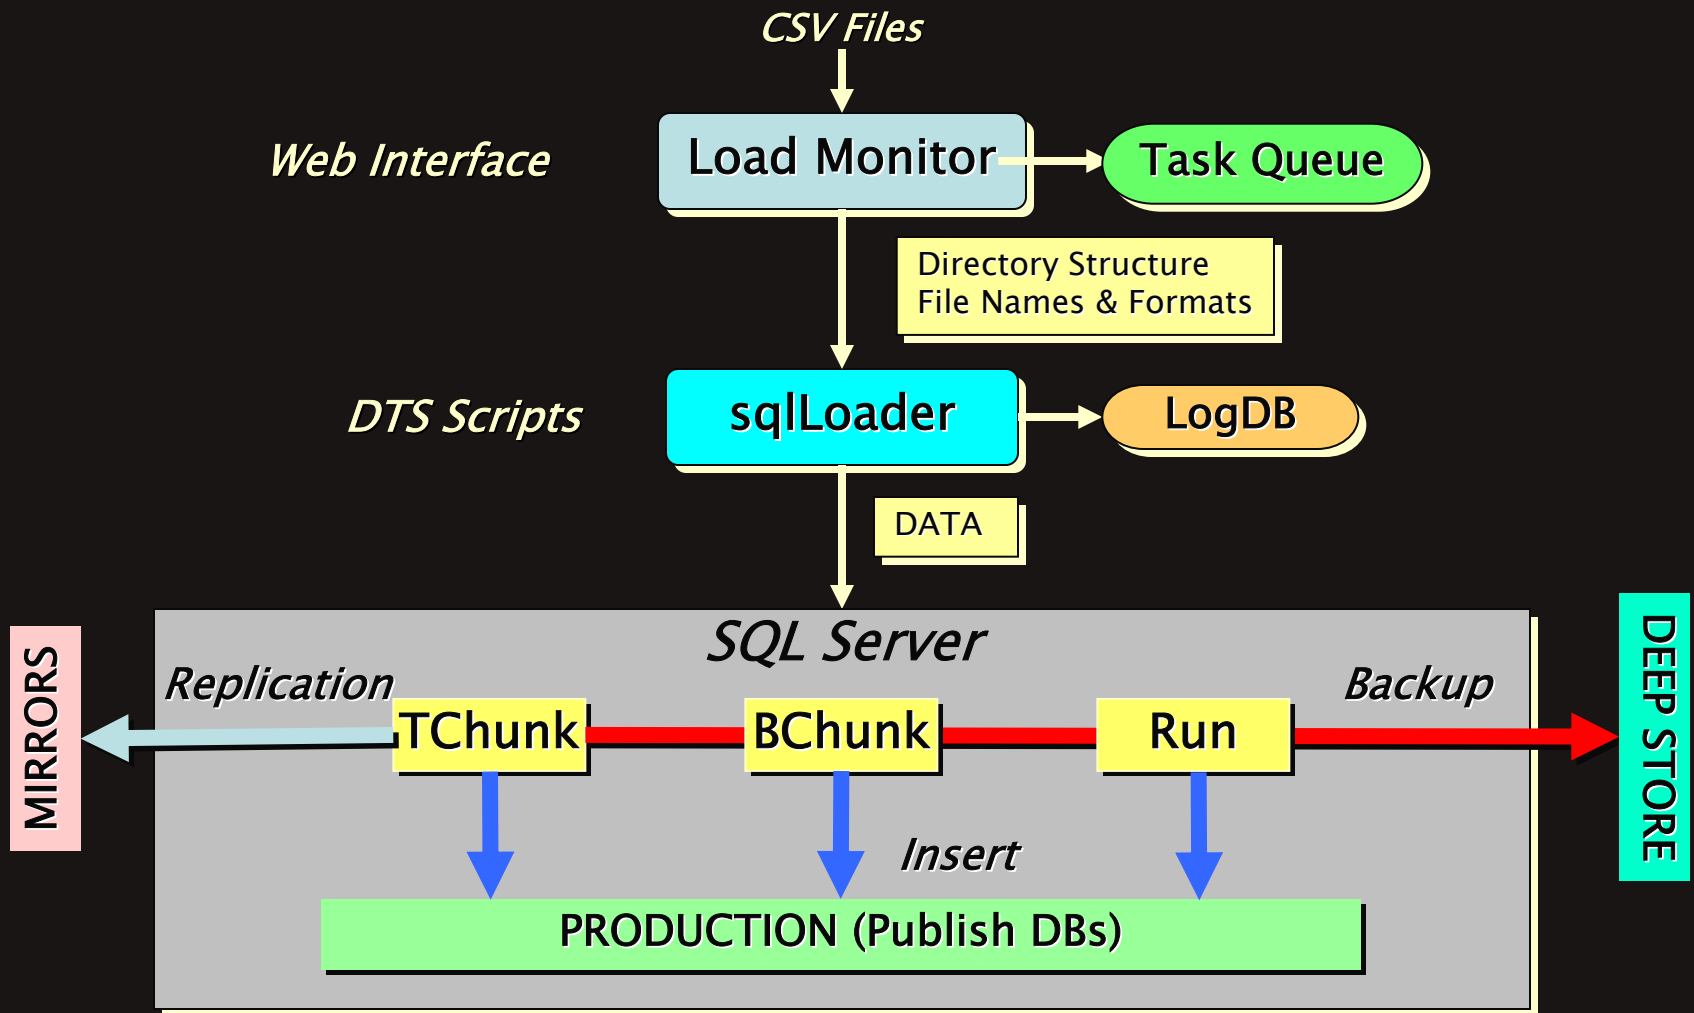


sqlLoader

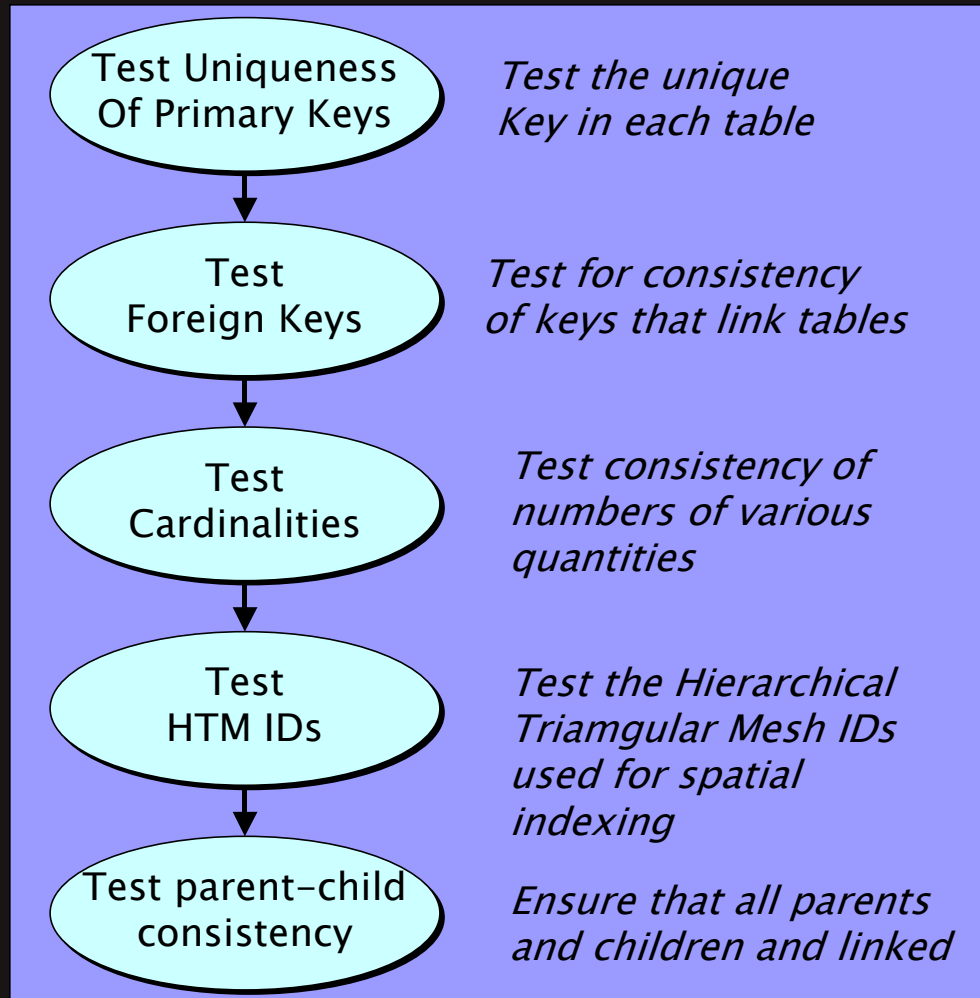


- Distributed workflow system
- System of scripts & procedures for data loader
 - SQL scripts, stored procedures
 - VBS scripts, DTS packages
 - Used linked servers, SQL agent jobs to automate multi-server loading
- Load Monitor (admin) web interface
 - Submit and monitor loading tasks
 - Cancel (kill) tasks
 - Monitor loadserver cluster
- Thorough data validation scripts
 - Most errors are in input data
 - This is our one chance to get the data right

sqlLoader Flow

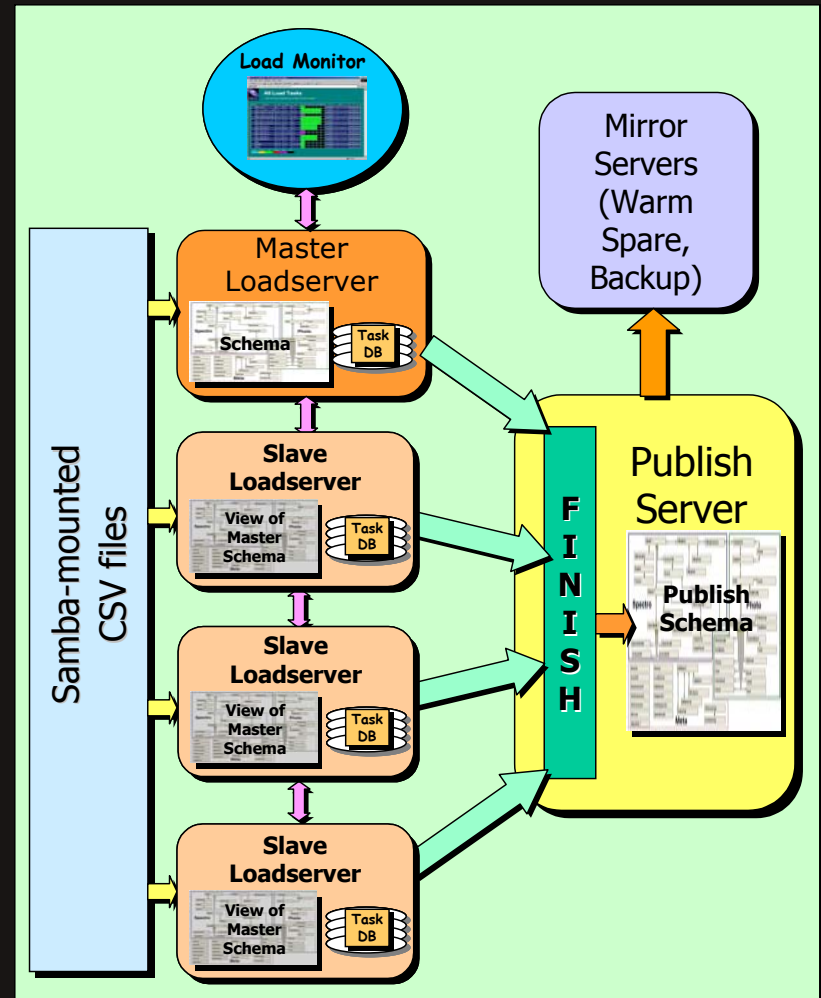


sqlLoader Data Validation

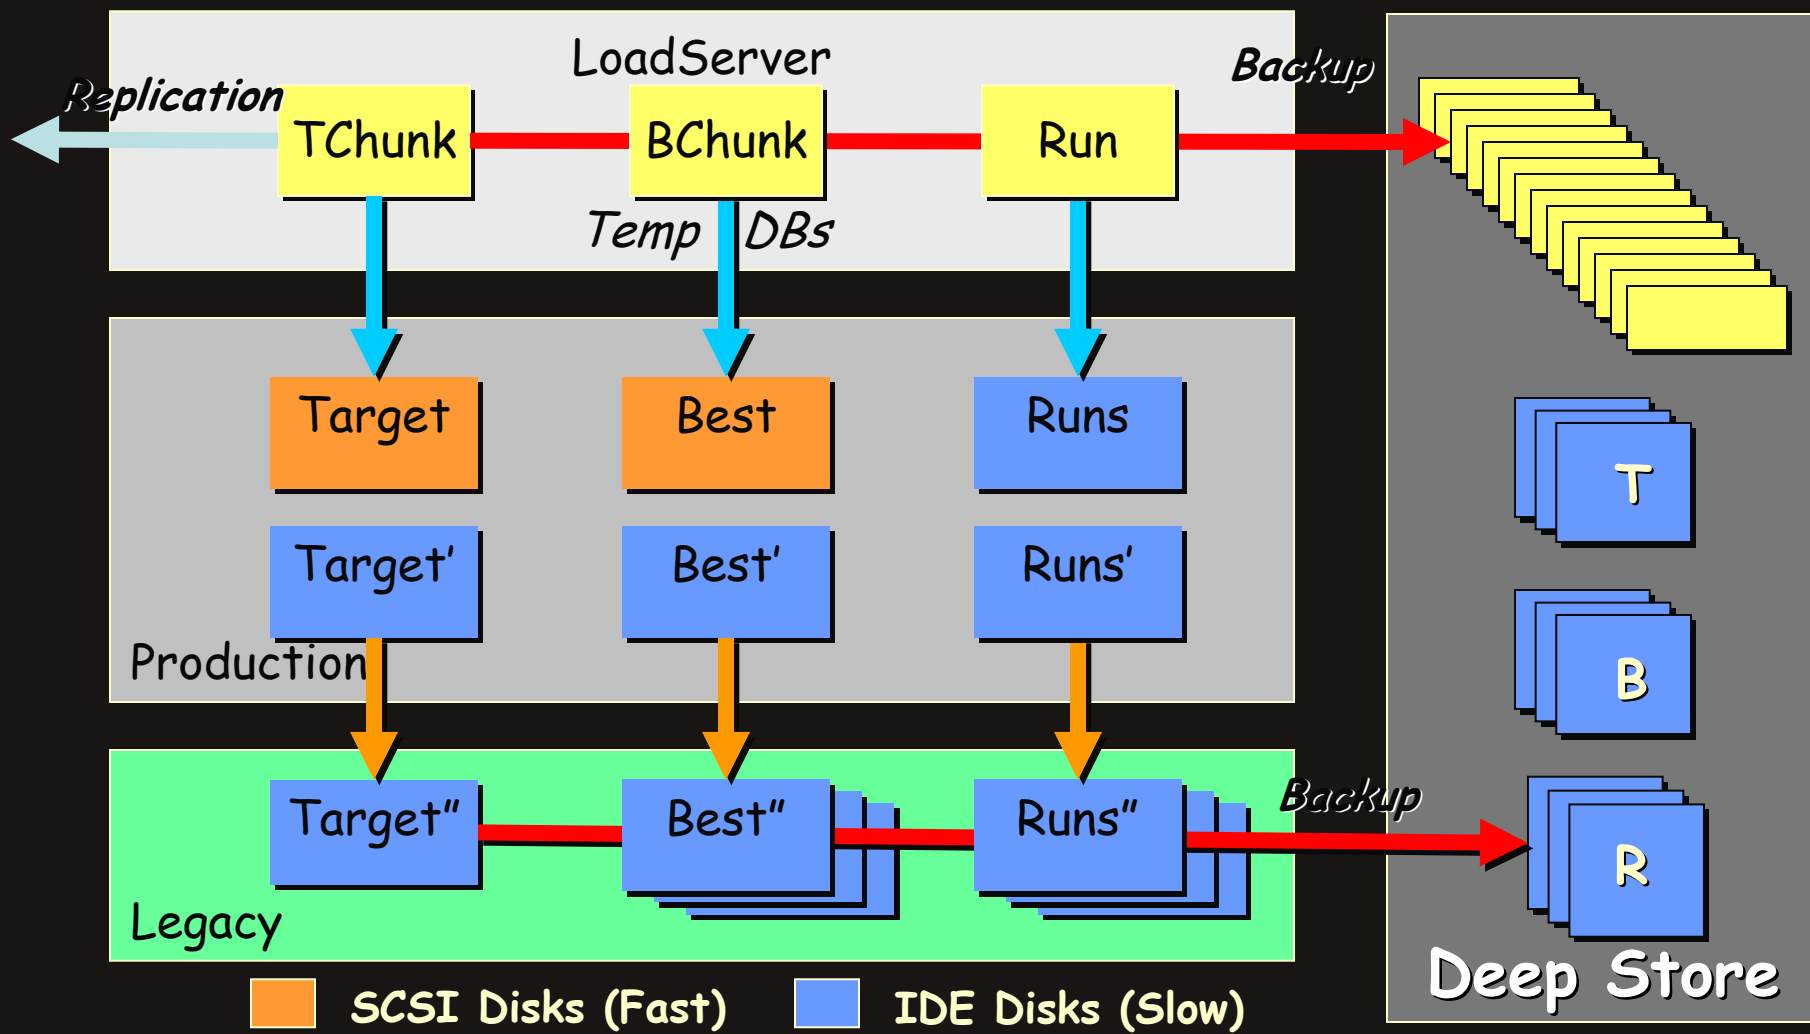


sqlLoader Distributed Loading

- Master/Slave servers
- Loading in parallel
 - Each slave server has remote *view* of schema
- Publishing sequential
 - Finish/Merge step
 - Pull data from loaders
- Configuration can be changed
 - Additional loaders can be added later



Archive Backup Strategy

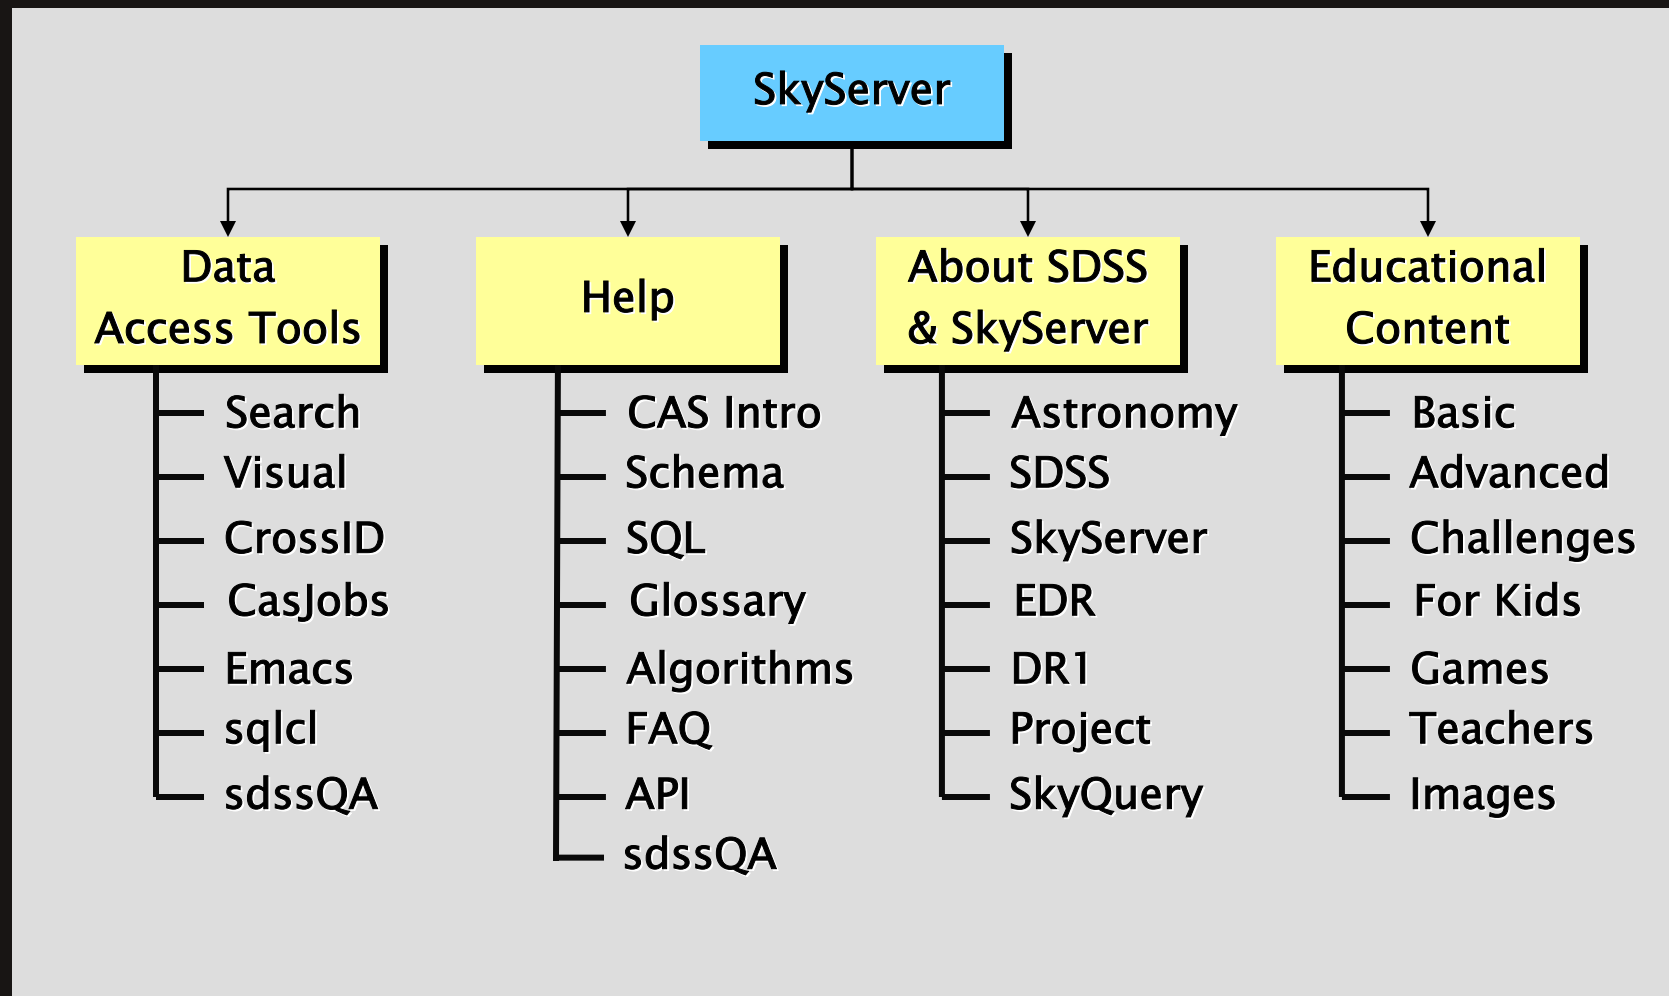


Loading Strategy



- Public data releases reloaded from scratch
 - Usually involves reprocessing of data
 - 2 backups of master archive maintained
 - Storage requirement: $6 \times$ size of archive
 - Leaves enough room for CSV files and task DBs
- Collaboration data incrementally loaded
 - Starting point is backup of public release
- Ability to back out changes
 - Delete and reload a chunk
 - Currently must be done manually
 - Needs to be integrated into loader scripts

SkyServer Front End



CAS Documentation



- Lot of it checked into DB, auto-generated
 - Table, view, function, column, indices
 - Auto-generated Schema Browser
 - Glossary and Algorithms
- CAS Intro
 - Data model and data access tools
- FAQ
- API – direct access to skyserver tools
- SQL intro
- sdssQA
- Sample Queries
- [sqlLoader User Guide](#), [loader workflow steps](#)

DR1 Experience




- Public Release June, 2003
- Site upgraded October, 2003
 - New look & feel
 - Enhanced documentation and tools
- Overview of 2003 usage
 - Over million hits/month average on website
 - Vast majority of queries fast (under 1 minute)
 - Very small fraction of slow queries
 - Significantly slows down server and affects everyone


DR1 Usage



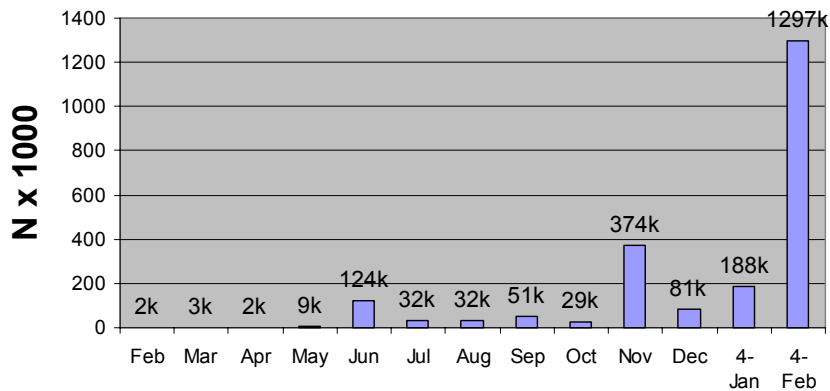
DR1 Hits per month



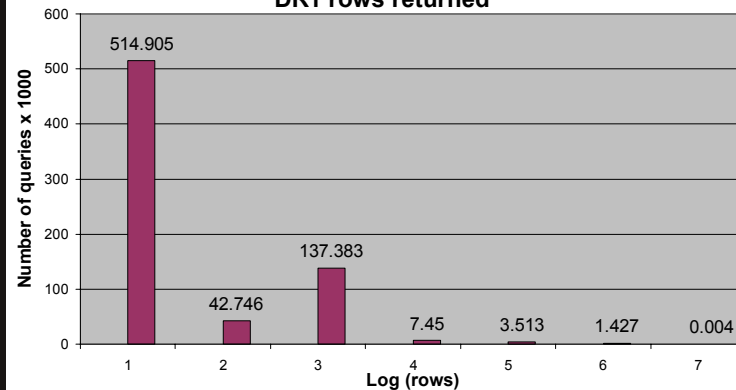
DR1 Query Execution Times





DR1 Queries



DR1 rows returned



DR2 Usage



- Same usage patterns as DR1
- Many short queries, few long ones

