The Sloan Digital Sky Survey

Migrating a (Large) Science Database to the Cloud

Ani Thakar

Alex Szalay

Center for Astrophysical Sciences and
Institute for Data Intensive Engineering and Science (IDIES)
The Johns Hopkins University

Motivation

- Exponential growth in scientific data size
  - Astronomy no exception, SDSS leading the way
- Large scientific databases in the cloud
  - "large" = multi-TB , "database" = SQL Server
  - On-demand resources very attractive to scientists
    - Don't want to become FT data managers
    - Most projects don't have have deep IT pockets
- Focus on migrating data in this talk
  - Is it easy? Is it even possible?
  - Performance evaluation premature at best
  - Economics even more so

Data Avalanche in Astronomy

- Moore's Law squared
- 2009 Physics Nobels:
  - CCD technology
  - Optical fiber technology
- Both were critical for SDSS, new generation of astronomy surveys
- SDSS now few TB
- Pan-STARRS next 50x SDSS, 100 TB
- LSST later One SDSS every 2-3 nights, PBs

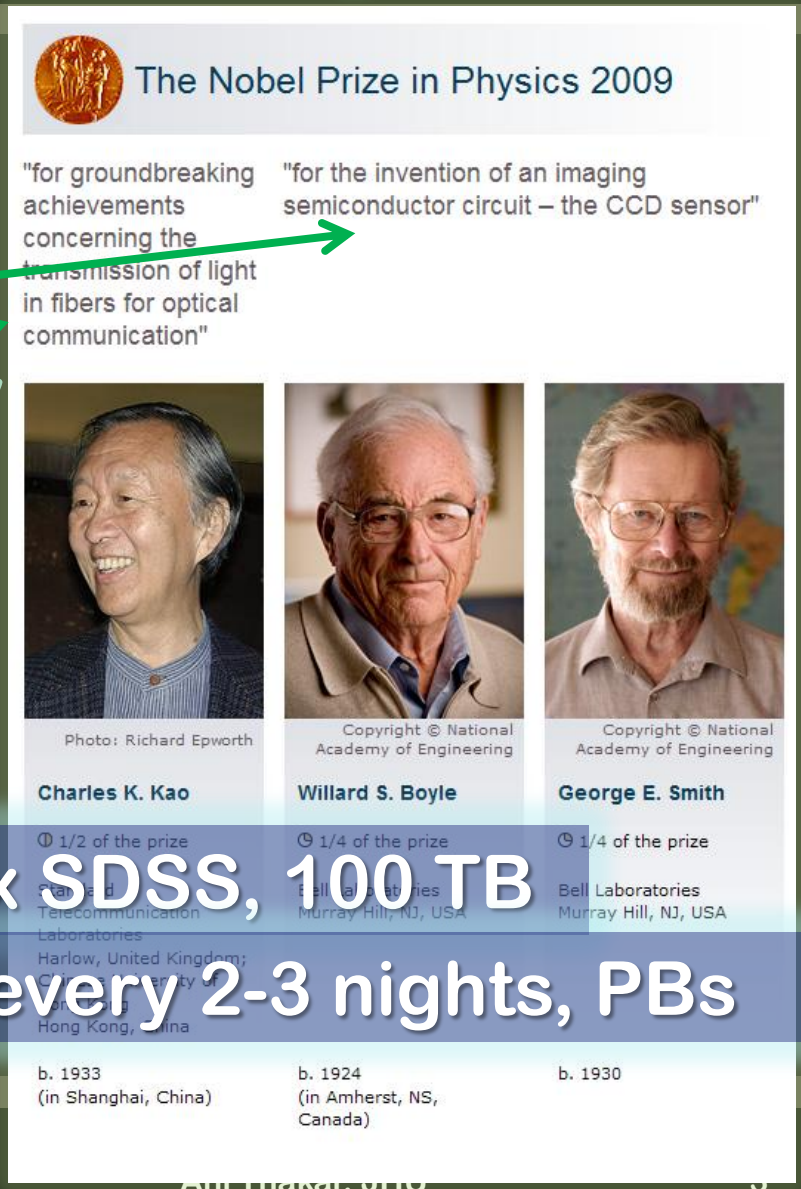

Cloudy with a chance of … pain?

- Can entire DB be moved to the cloud as is?
- Restrictions on data size
  - For now, use subset of original database
  - In future, will have to partition the data first
- Migrating the data to the cloud
  - How to copy data into the cloud?
  - Will it need to be extracted from the DBMS?
  - Will schema have to be modified?
  - Will science functionality be compromised?
  - What about query performance?

Cloud Experiments

- Dataset
  - Sloan Digital Sky Survey (SDSS)
  - Public dataset, no restrictions
  - Reasonably complex schema, usage model
  - Easy to generate subset of arbitrary size
  - Only available as SQL Server database
- Clouds
  - Amazon Elastic Cloud 2 (EC2)
    - AWS wanted to host public SDSS dataset
  - Microsoft Azure
    - Natural fit for SQL Server databases

SDSS SkyServer DR7 - Mozilla Firefox

File Edit View History Bookmarks Tools Help

SDSS SkyServer DR7

Sloan Digital Sky Survey / SkyServer

Home | Tools | Schema | Projects | Astronomy | SDSS | Contact Us | Download | Site Search | Help

Due to system maintenance this site will be unavailable Thursday October, 22nd from 7AM central until 7:30AM central. We apologize for the inconvenience.

Welcome to the **DR7 site**!!!

This website presents data from the Sloan Digital Sky Survey, a project to make a map of a large part of the universe. We would like to show you the beauty of the universe, and share with you our excitement as we build the largest map in the history of the world.

News

The site hosts data from **Data Release 7 (DR7)**, **What's new in DR7**, **what's new on this site**, and **known problems**. **More...**

For Astronomers

A separate branch of this website for professional astronomers (English)

More...

SDSS is supported by

SkyServer Tools

- Famous places
- Get images
- Visual Tools
- Explore
- Search
- Object Cross-ID
- CasJobs

Science Projects

- Basic
- Advanced
- Challenges
- For Kids
- Games and Contests
- Teachers
- Links to other projects

Info Links

- About Astronomy
- About the SDSS
- About the SkyServer
- SDSS Data Release 7
- SDSS Project Website
- Open SkyQuery
- Images of RC3 Galaxies

Help

- Getting Started
- FAQ
- How To
- Glossary
- Schema Browser
- Sample SQL Queries
- Details of SDSS Data

Powered by Microsoft

Site Traffic
Privacy Policy




Contact Us

SDSS Catalog DB Schema

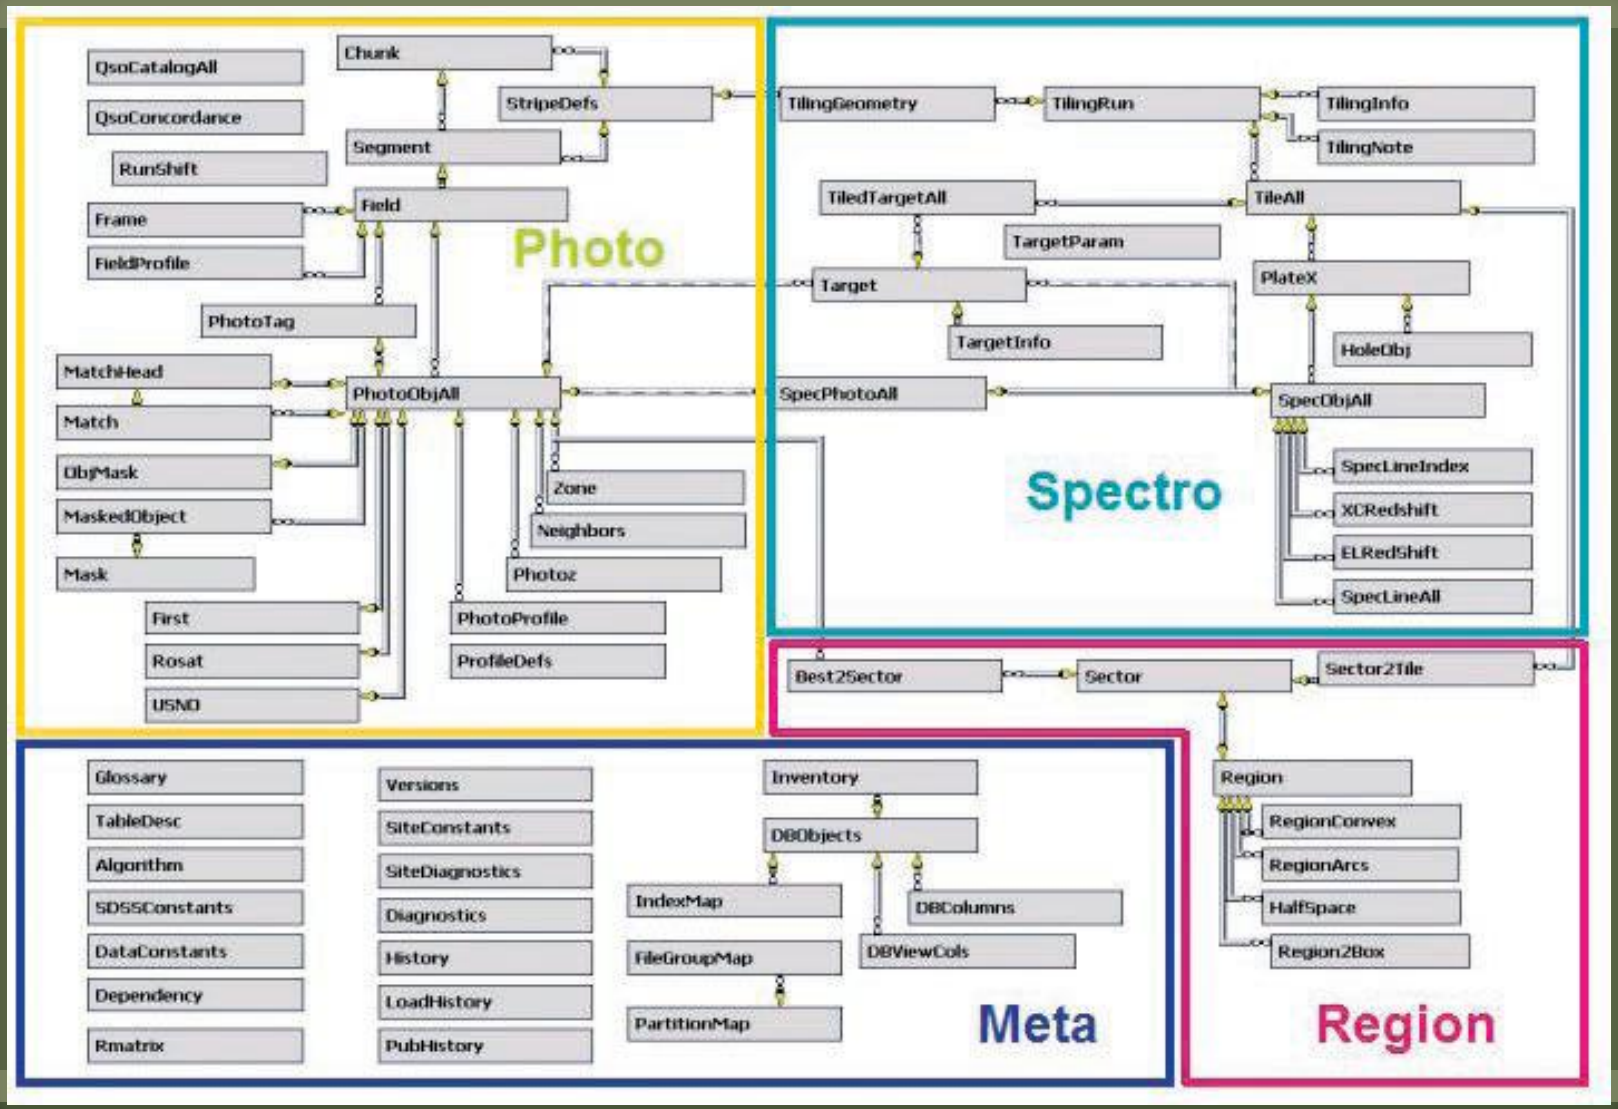

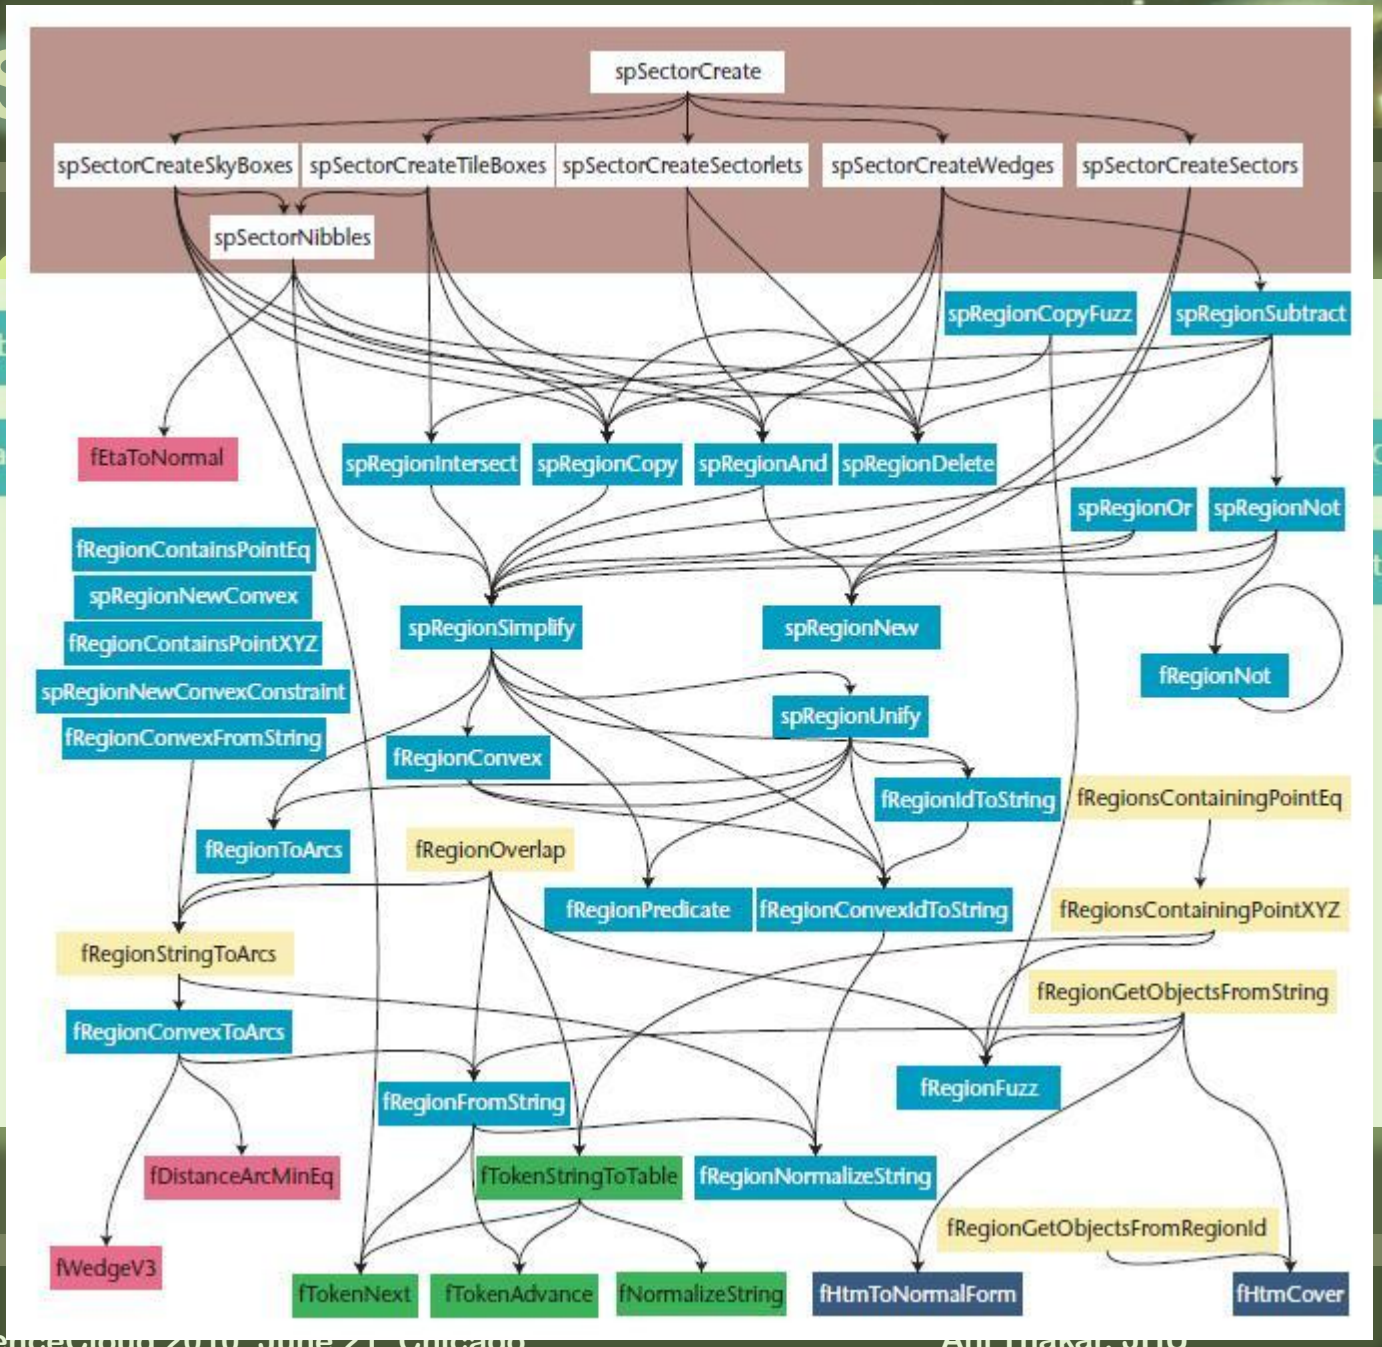
spSectorCreate
spSectorCreateSkyBoxes
spSectorCreateTileBoxes
spSectorCreateSectorlets
spSectorCreateWedges
spSectorCreateSectors
spSectorNibbles
spRegionCopyFuzz
spRegionSubtract
fEtaToNormal
spRegionIntersect
spRegionCopy
spRegionAnd
spRegionDelete
spRegionOr
spRegionNot
fRegionContainsPointEq
spRegionNewConvex
fRegionContainsPointXYZ
spRegionNewConvexConstraint
fRegionConvexFromString
spRegionSimplify
spRegionNew
fRegionNot
spRegionUnify
fRegionConvex
fRegionIdToString
fRegionsContainingPointEq
fRegionToArcs
fRegionOverlap
fRegionPredicate
fRegionConvexIdToString
fRegionsContainingPointXYZ
fRegionStringToArcs
fRegionGetObjectsFromString
fRegionConvexToArcs
fRegionFuzz
fRegionFromString
fDistanceArcMinEq
fTokenStringToTable
fRegionNormalizeString
fRegionGetObjectsFromRegionId
fWedgeV3
fTokenNext
fTokenAdvance
fNormalizeString
fHtmToNormalForm
fHtmCover

Migrating data to Amazon EC2

- 1 TB data size limit (per instance)
- SDSS DR6 100 GB subset: BestDR6_100GB
  - Actually more like 150 GB
  - Large enough for performance tests, small enough to be migrated in a few days/weeks
- Several manual steps to create DB instance
- Could not connect to DB from outside
- Preliminary performance test, but without optimizing within cloud

SDSS public dataset on EC2

- Public datasets on Amazon
  - Stored as snapshot available on AWS
  - Advertised on AWS blog
  - Anyone can pull into their account
    - Data is free, but not the usage
  - Create a running instance
  - Multiple instances deployed manually
- First SQL Server dataset on EC2 (?)
  - AWS also created a LINUX snapshot of SQL Server DB!

AWS Blog Entry

http://aws.typepad.com/

New Public Data Set: Sloan Digital Sky Survey DR6 Subset

The Sloan Digital Sky Survey, or SDSS, is now available as a Public Data Set.

Weighing in at 180 GB, the SDSS is the most ambitious astronomical survey ever undertaken. The researchers have used a 2.5 meter, 120 megapixel telescope located in Apache Point, New Mexico to capture images of over one quarter of the sky, or about 230 million celestial objects. They have also created 3-dimensional maps containing more than 930,000 galaxies and 120,000 quasars.

This new public data set (which is a subset of the entire SDSS) will be of interest to students, educators, hobby astronomers, and researchers. From a standing start, it is possible to launch an EC2 instance, create an Elastic Block Store volume with this data, attach the volume to the instance and start examining and processing the data in less than ten minutes.

The data set takes the form of a Microsoft SQL Server MDF file. Once you have created your EBS volume and attached it to your Windows EC2 instance, you can access the data using SQL Server Enterprise Manager or SQL Server Management Studio. The SDSS makes use of stored procedures, user defined functions, and a spatial indexing library, so porting it to another database would be a fairly complex undertaking.

I know from experience (my son Andy is studying Astronomy at the University of Washington and is always showing me the "please delete your unnecessary files" emails from the department's administrator) that storage space is always at a premium in academic settings, due in part to the existence of large scale data sets like this. The combination of EC2, EBS, this public data set, and our AWS in Education program should enable students and educators to analyze, process, display, and study the universe in revolutionary ways.

-- Jeff;

September 28, 2009 in Amazon EC2 | Permalink | Comments (1) | TrackBack (0)

How it's supposed to work

- With other db dumps, assumption is that users will set up their own DB and import the data
- Public SDSS dataset to serve two kinds of users:
  a) People who currently access the SDSS CAS
  b) General AWS users who are interested in the data
- For a), should be able to replicate the same services that SDSS currently has, but using a SQL Server instance on EC2
- For b), users should have everything they need on the AWS public dataset page for SDSS

Steps to create DB in EC2

- Create “snapshot” of database first
- Create storage for DB: 200 GB EBS volume
  - Instantiate snapshot as volume of required size
- Create SQL Server 2005 Amazon Machine Image (AMI)
  - AMI instance from snapshot
- Attach AMI instance to EBS volume
  - Creates running instance of DB
- Get Elastic IP to point to instance

Steps to create Web interface

- Create volume (small) with win2003
- Create instance with Win 2003 AMI (only IIS, no SQL Server)
- Attach volume to instance
- Get public DNS, admin account
- BUT … couldn't connect to SQL Server IP
- So outside world cannot connect to the data as yet

Steps to Create DB Server

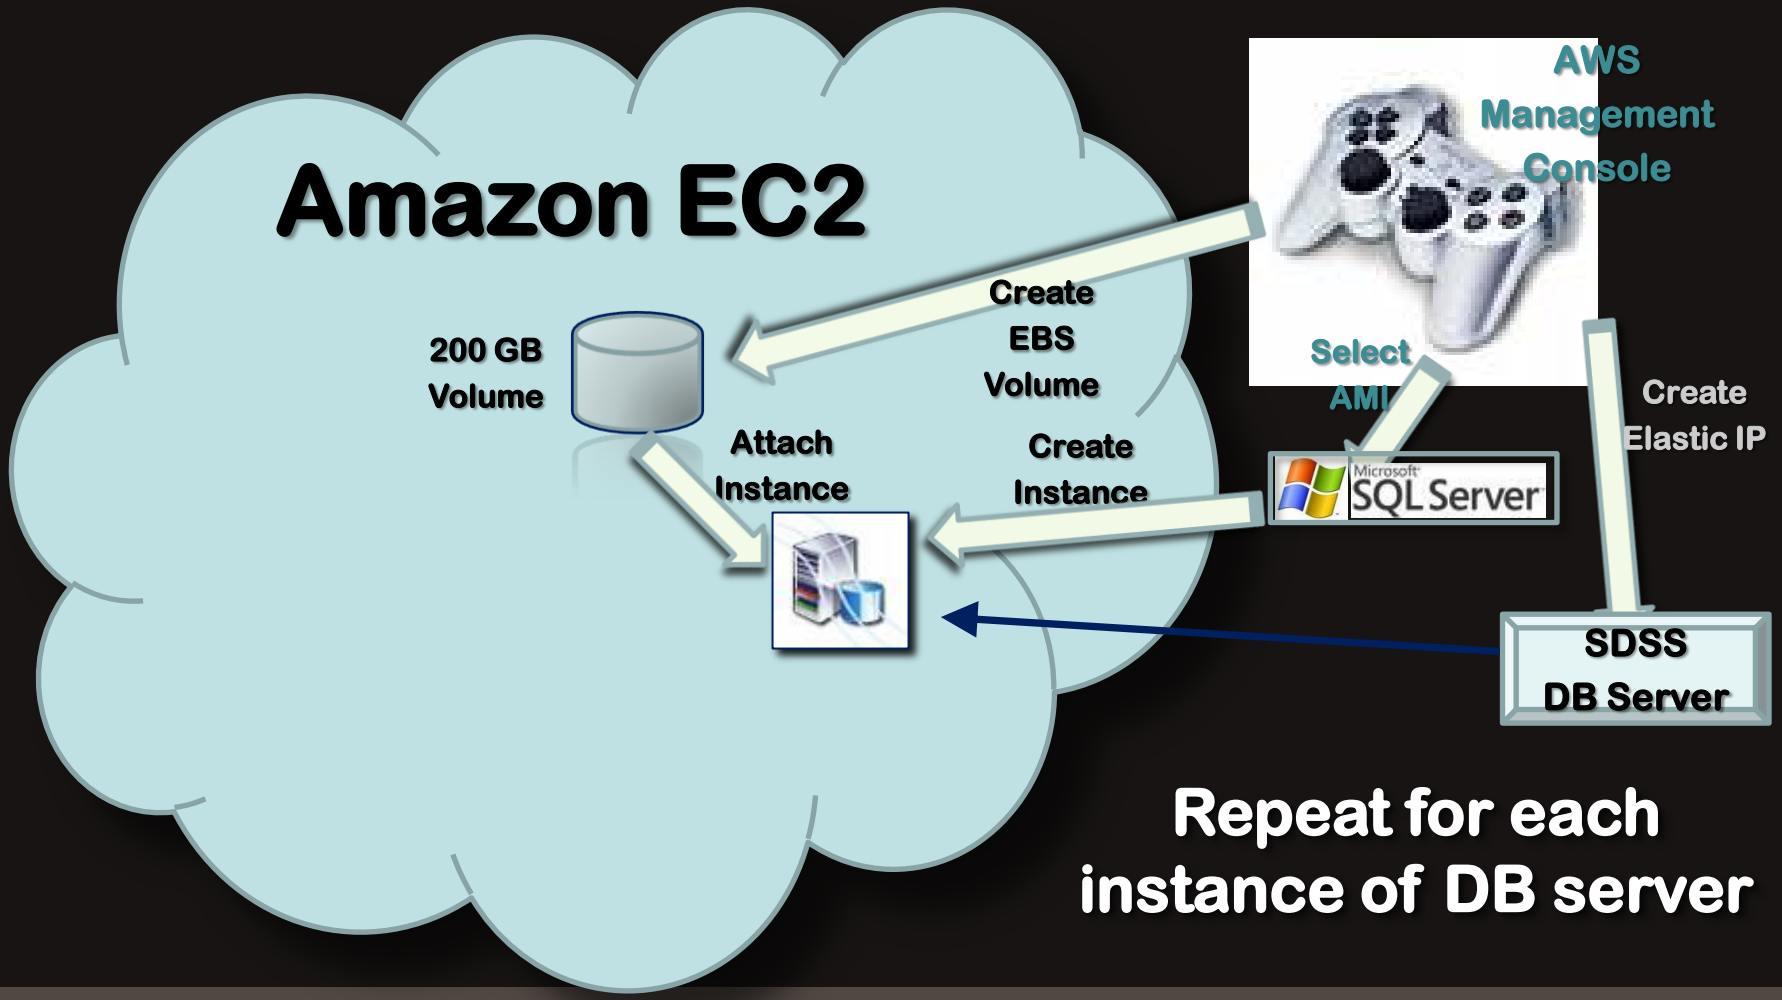

Comparing Cost & Performance

Query	GW_CPU	GW_Elapsed	log(GW_IO)	EC2_CPU	EC2_Elapsed	log(EC2_IO)	GW_IO	EC2_IO
QSX11	4.3125	2.766667	3.502199	4.135417	22.124	3.520047	3178.333	3311.667
QSX12	29.57292	84.28067	4.974367	25.05208	660.9524	4.966658	94268.67	92610
QSX13	24.54167	80.62667	4.972407	19.44792	648.5096	4.965402	93844	92342.67
QSX15	4.90625	2.889667	3.546378	4.604167	23.022	3.552181	3518.667	3566

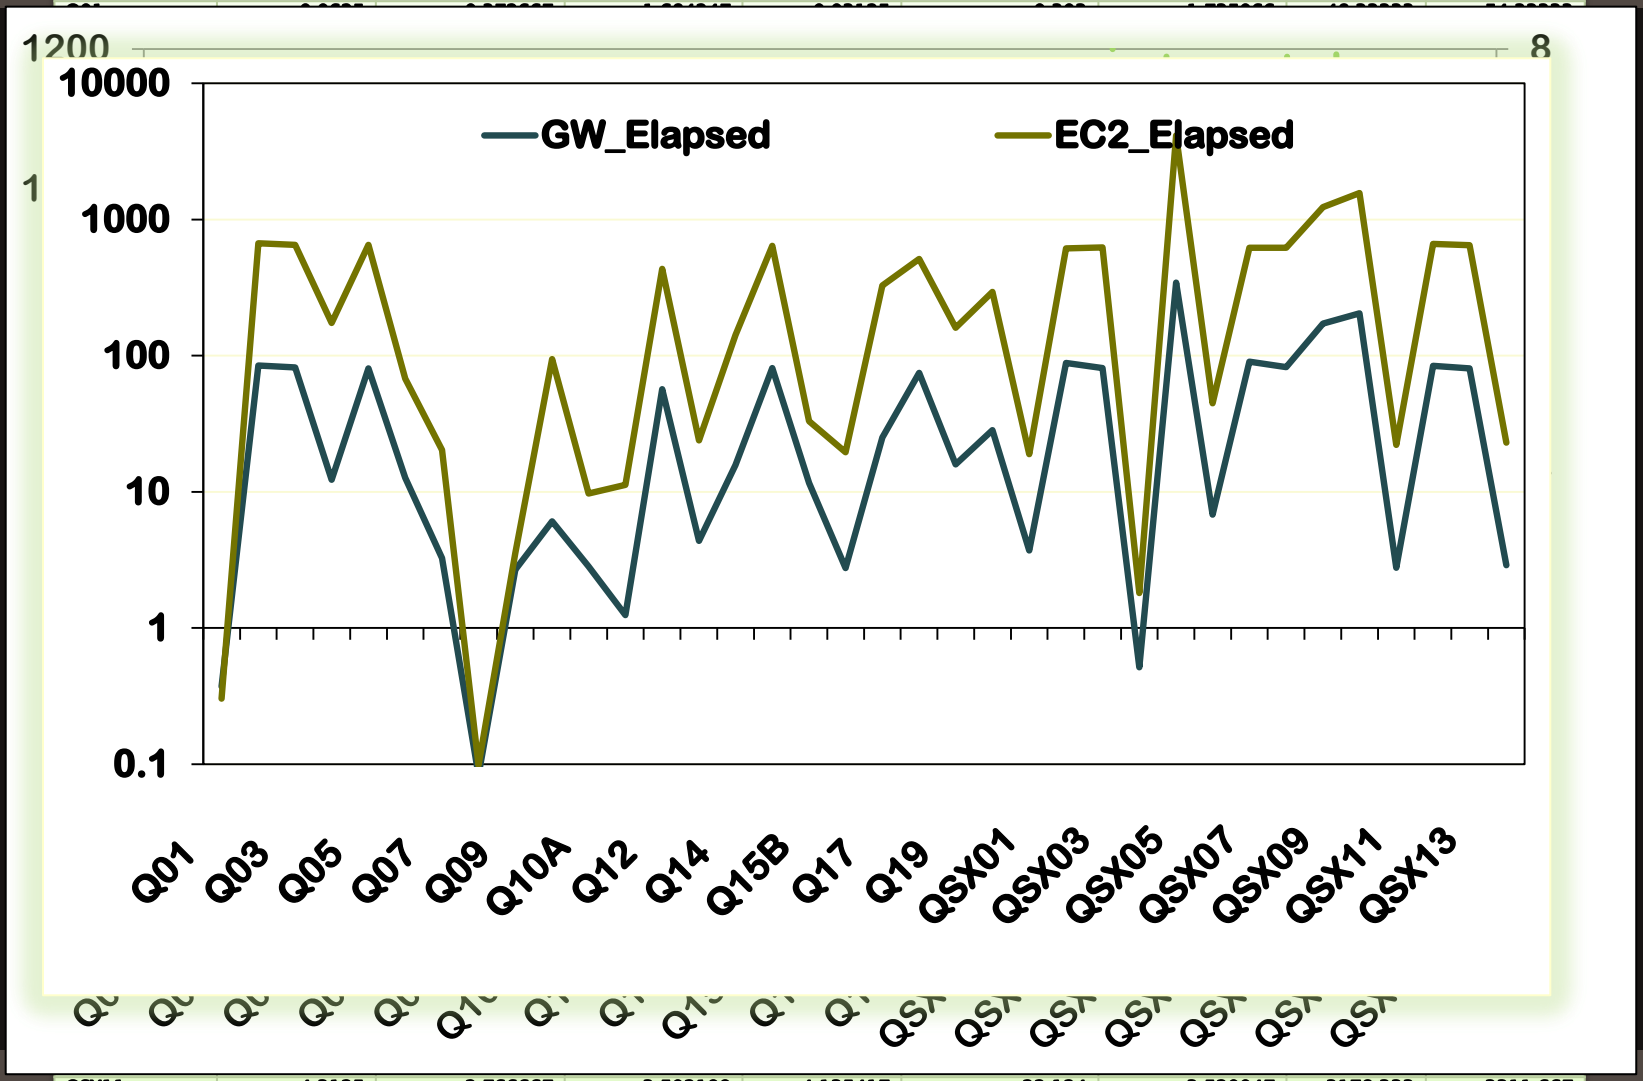

SDSS Production Cluster at FNAL

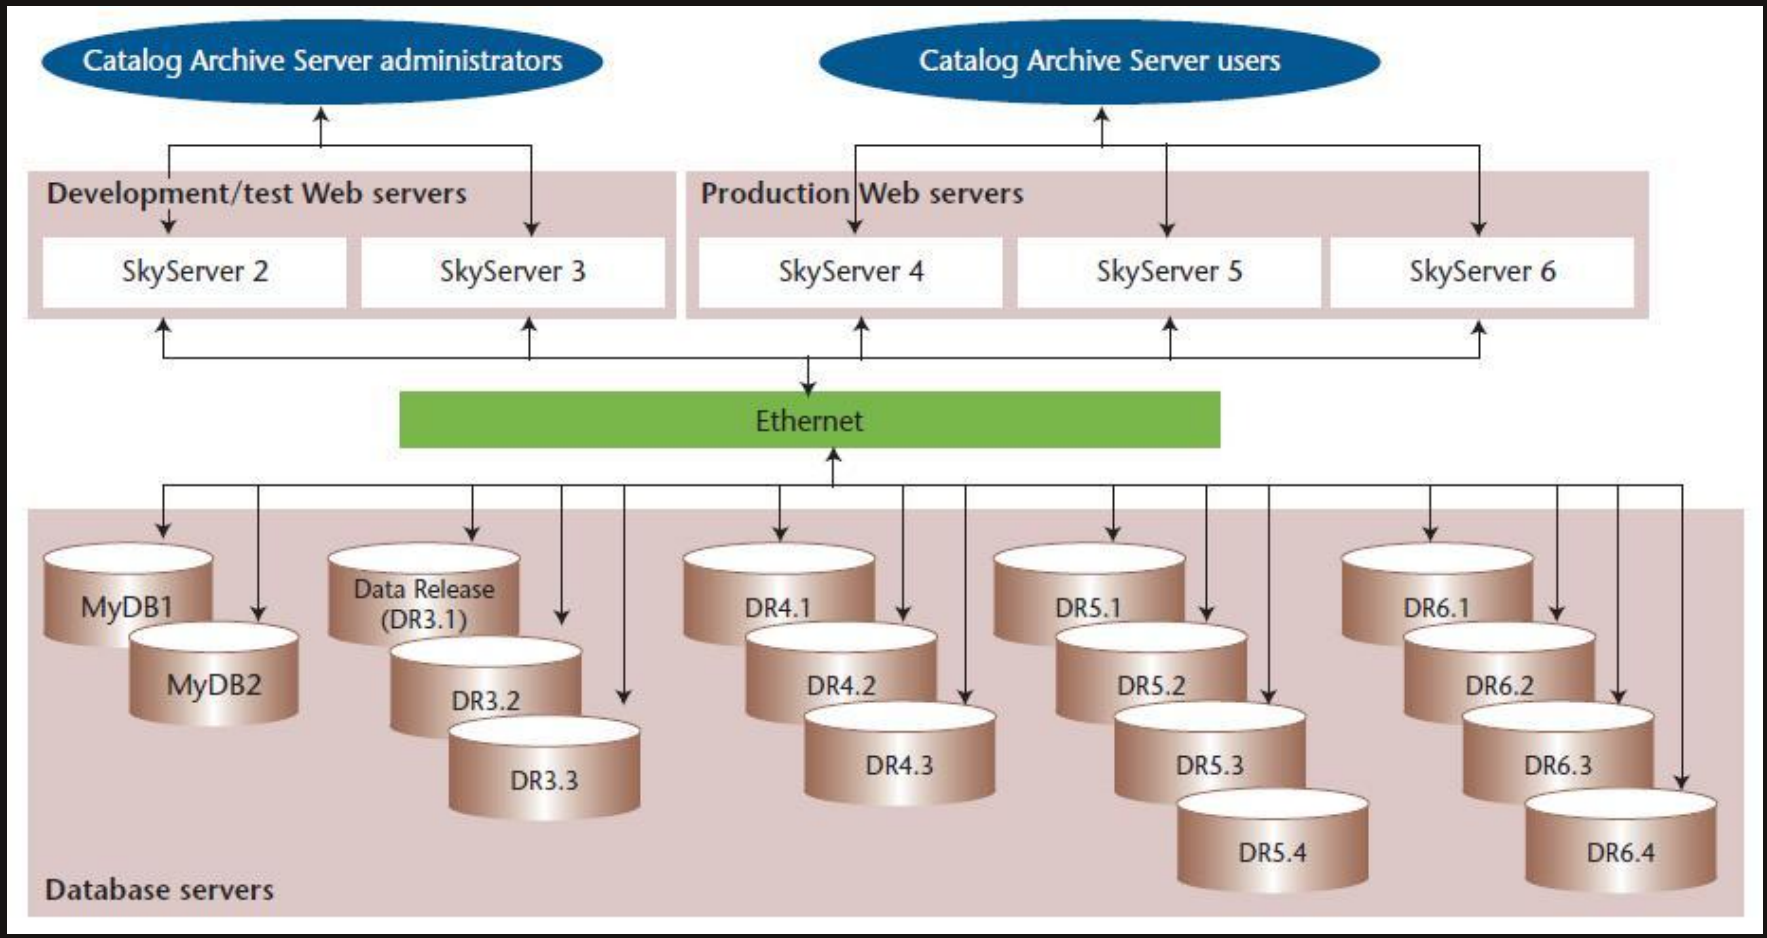

Migrating data to Microsoft Azure

- 10 GB data size limit (⇑ 50 GB last week!)
- SDSS DR6 10 GB subset: BestDR6_10GB
- Two ways to migrate database
  - Script it out of source db (very painful)
    - Many options to select
    - Runs out of memory on server!
  - Use SQL Azure Migration Wizard (much better!)
    - Runs for few hours
    - Produces huge trace of errors, many items skipped
    - But does produce a working db in Azure

SQL Azure Migration Wizard

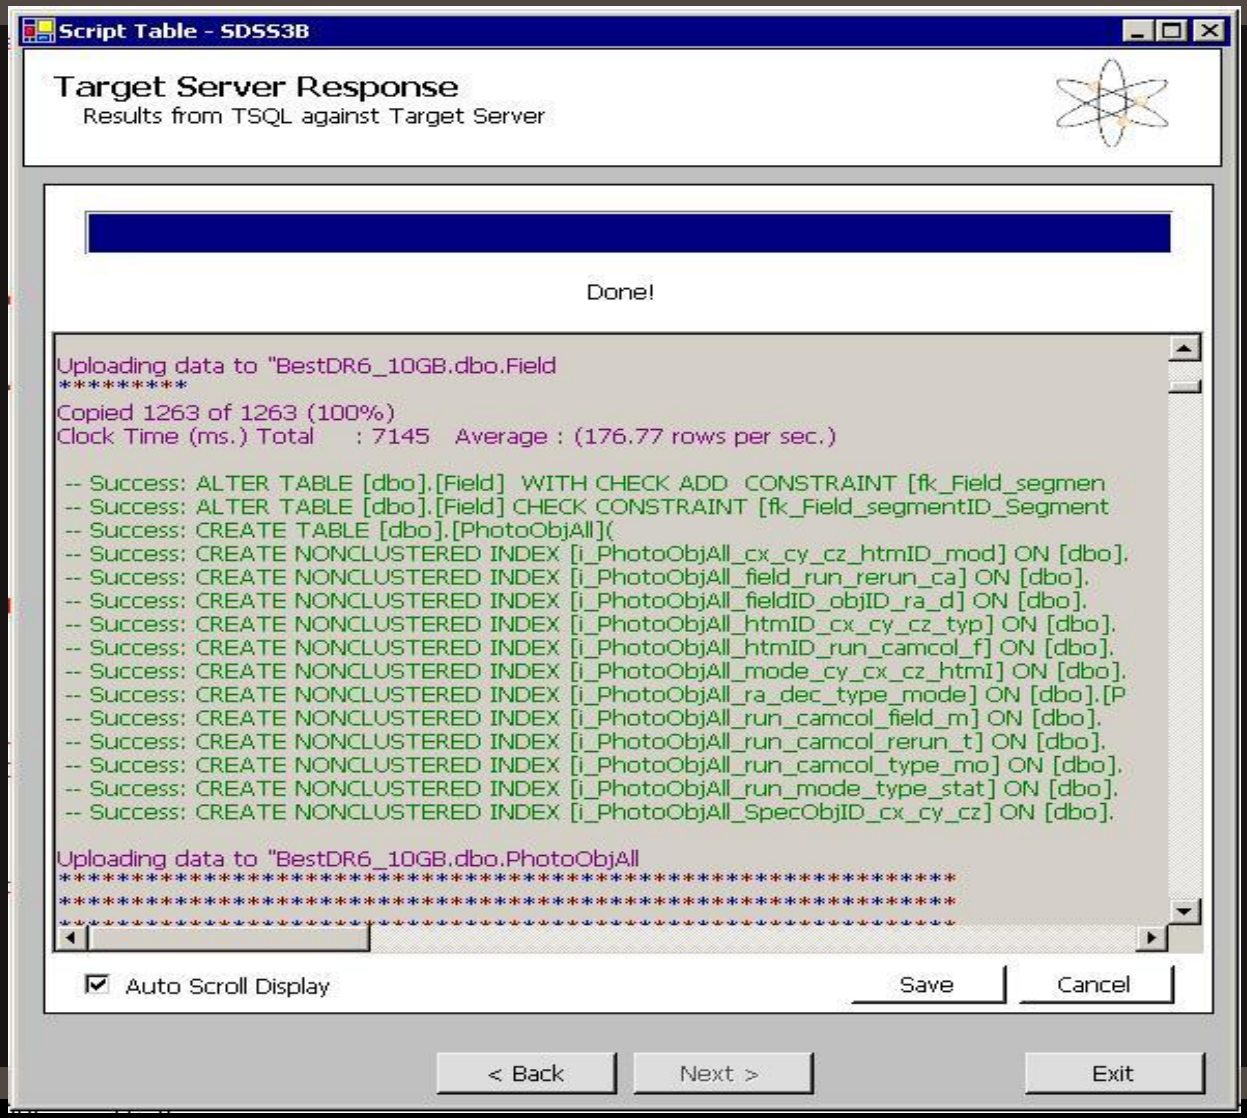

Unsupported MS-SQL features

- No references to other databases
  - Can't run command (including SQL) scripts from ***Master*** DB
- Global temp objects disallowed
  - Can't use performance counters in test query suite
  - Difficult to benchmark and compare performance
- SQL-CLR function bindings not supported
  - Can't use our HTM spatial indexing library
  - Makes spatial searches much slower

Unsupported MS-SQL features

- T-SQL directives
  - e.g., to set the level of parallelism
- Built-in T-SQL functions
  - Probably can do without these for now
  - More for convenience
- Deprecated features
  - Lose the SQL Server 2000 baggage

Azure sign of success?

- Ok, so data in Azure, but at what cost?!
  - Meaningful performance comparison not possible
  - Dataset too small
  - Schema features stripped
  - No spatial index, so many queries crippled
- Connected to Azure DB from outside cloud
  - Hooked up SkyServer WI
  - Connected with SQL Server client (SSMS)
  - Ran simple queries from both

Conclusions

- Migrating scientific databases to the cloud not really possible at the moment
- Even migrating smaller DBs can be painful
  - Several steps to deploy each copy
  - Cloud may not support full functionality
  - Problems limited to SQL Server DBs?
- Large DBs will have to be partitioned
  - Set up distributed databases in the cloud
  - Query distributed databases from outside
- Haven't even talked about economics yet

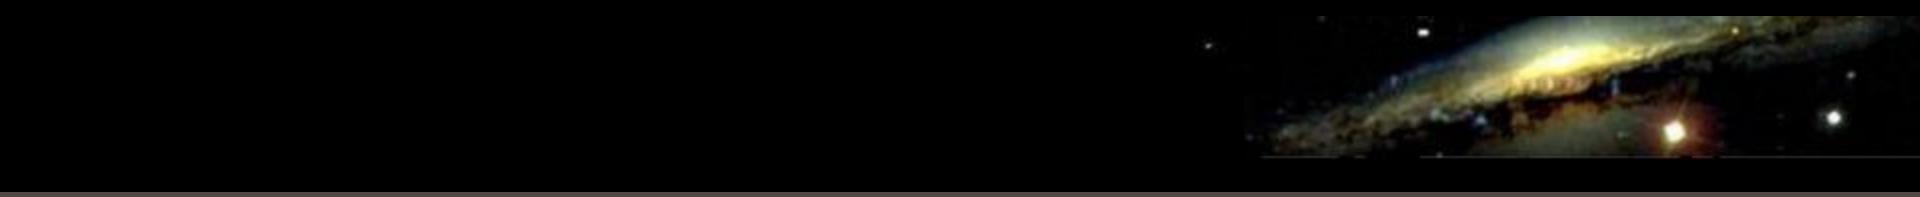

"The one who says it can't be done should never stop the one who's doing it."

Anon.

Hope to hear more positive experiences of others!

Thank you!

A Flight Through 400,000 SDSS Galaxies (1/500th of total number of galaxies in SDSS DR7)

SDSS DR7

Miguel A. Aragon (JHU)
Mark Subbarao (Adler P.)
Alex Szalay (JHU)