Scientific Data Generators: Plague or Panacea?

INRIA-Illinois Workshop
Nov 19, 2012

Rob Pennington
CTO and Deputy Director
NCSA

National Center for Supercomputing Applications
University of Illinois at Urbana-Champaign

NCSA – High Level View

- Established in 1986 by NSF as a supercomputing center
  - National, University and Industry resource
- Four major areas at NCSA
  - Extreme Scale Computing
  - Distributed Cyberinfrastructure
  - Data Cyberinfrastructure
  - Private Sector Program
- Open research - access to the NSF funded resources
  - Allocated by national review panels, no cost to the researchers
  - Primarily academic researchers
    - Most researchers are <u>not</u> associated with UIUC
  - Open research by industry users also by allocation
- Industrial proprietary research on a cost recovery basis

Big Data

- High Performance Computing – started from the high end and spread downwards
  - Originally considered to be the $10M system that only a few needed, understood and could afford to purchase
  - Slow dissemination beyond a dedicated and committed cadre
  - Now ubiquitous because the usefulness has been demonstrated
- Big Data is different – it comes from all directions
  - Science, Business, Government … Personal Digital Archiving
- Most data is now born digital:
  - "Big Data" … does not refer just to the volume of data, but also to its variety and velocity. Big data includes large, diverse, complex, longitudinal and/or distributed data sets generated from instruments, sensors, Internet transactions, email, video, click streams, and/or all other digital sources. NSF 12-499

Four Examples of Big Data Generators

- XSEDE
  - Cyberinfrastructure framework
  - >$100M, data volume TBD over 5 years
- Dark Energy Survey
  - Astronomical survey instrument
  - >$40M, ~2 PB over 5 years
- Large Synoptic Survey Telescope
  - Next generation astronomical survey instrument
  - ~$400M, >100PB over 10 years
- Blue Waters
  - Simulation instrument
  - >$300M and >300 PB over 5 years

Dark Energy Survey

- DES combines four probes of Dark Energy:
  - Type Ia Supernovae (SN)
  - Baryon Acoustic Oscillations (BAO)
  - Galaxy clusters (GC)
  - Weak Gravitational Lensing (WL)
- Observing Program
  - Five years, total of 2 PB
  - 525 nights of observation
  - 300 million galaxies
  - Survey will image 5000 square degrees
  - 5 optical filters to obtain crude spectra
- Also observe smaller patches to discover and study thousands of supernovae.

CTIO 4m Telescope

Dark Energy Camera
570 Mpixels
74 CCDs

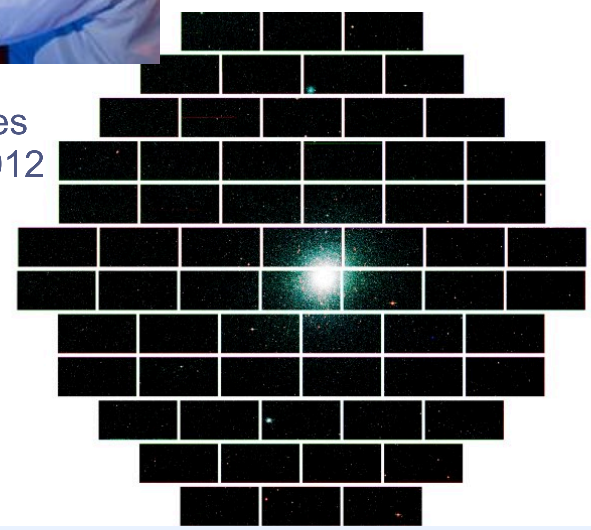

First Images
17 Sept 2012

Instrument Dedication: 9 Nov 2012, Chile

Src: D. Petravick/NCSA, DES project

DES Scientists and Science Facilities

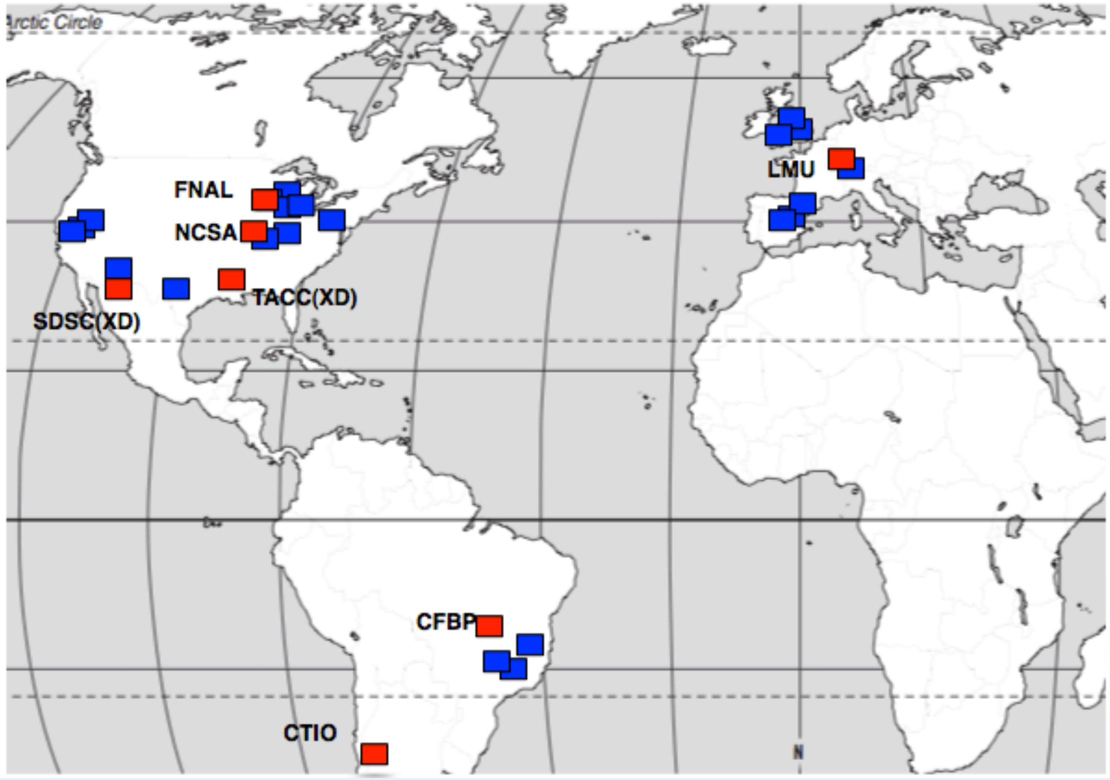

120 Scientist
From 6
Countries.

NCSA's role in the collaboration is to produce and serve data used for dark energy science analysis.

Src: D. Petravick/NCSA, DES project

LSST Science Parameters Summary

- Camera
  - 3.2 Gpix, 189 CCDs, 9.6 deg^2,
  - 6 filters, single-frame magnitude limits:
    - u=24, g=25.1, r=24.9, i=24.1, z=23.4, y=22.2

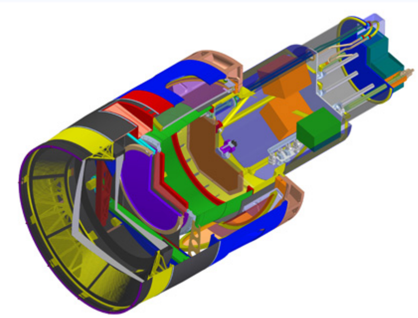

- Cadence
  - Survey area: 28,000 deg^2 including Galactic Confusion Zone
  - Total visits: 2,500,000
  - 298.35 nights per year; 10 hr/night mean
    - 19.5 sec exposure time, 2 exposures per visit
  - Epochs (visits to any given location):
    - u=60, g=83, r=187, i=187, z=170, y=170; total=857

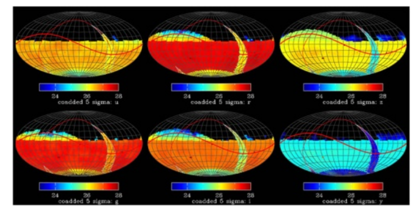

- Science
  - 1.97×10^6 Sources/visit
  - DR1: 7.33×10^9 galaxies, 5.90×10^9 stars (incl. visibility, uncertainty)
  - DR11: 2.34×10^{10} galaxies, 1.41×10^{10} stars

Src: M. Freemon/NCSA, LSST project

XSEDE Designed for Innovation & Evolution

- Coordination, support and creating a national/international cyberinfrastructure, led by NCSA
  - No funding for major hardware
  - ~120 FTE funded across 17 partner institutions
- An environment in which all resources, data and services relevant to a researcher can be embedded and shared
  - Campus bridging creating a single virtual system with interactive data transfer and resource sharing capabilities
    - "make my data accessible everywhere I want to be"
  - Coordinated archival approach to ensure persistence of important datasets beyond the lifetime of particular service providers
- An underlying infrastructure to support this
  - Open architecture with judicious use of standards designed to evolve in a non-disruptive way
  - Interoperability of XSEDE with other cyberinfrastructures

Src: J. Towns/NCSA, XSEDE project

XSEDE Systems Architecture

- Access Layer:
  - provides user-oriented interfaces to services
  - APIs, CLIs, filesystems, GUIs
- Services Layer:
  - protocols that XSEDE users can use to invoke service layer functions
  - execution management, discovery, information services, identity, accounting, allocation, data management,…
  - quite literally, the core of the architecture
- Resources Layer:
  - compute servers, filesystems, databases, instruments, networks, etc.

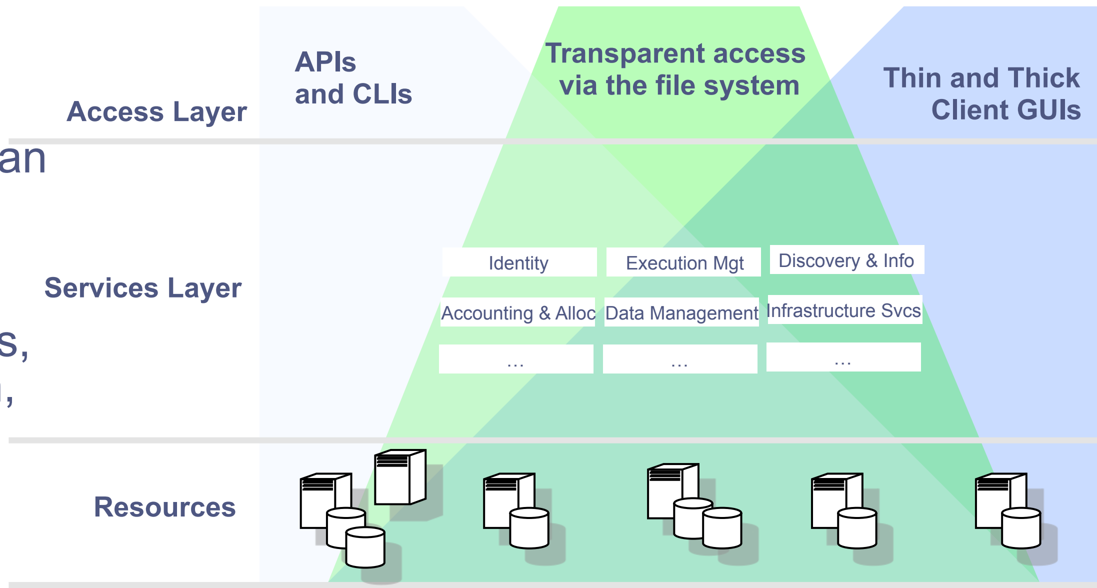

Src: J. Towns/NCSA, XSEDE project

Blue Waters Computing Super-system

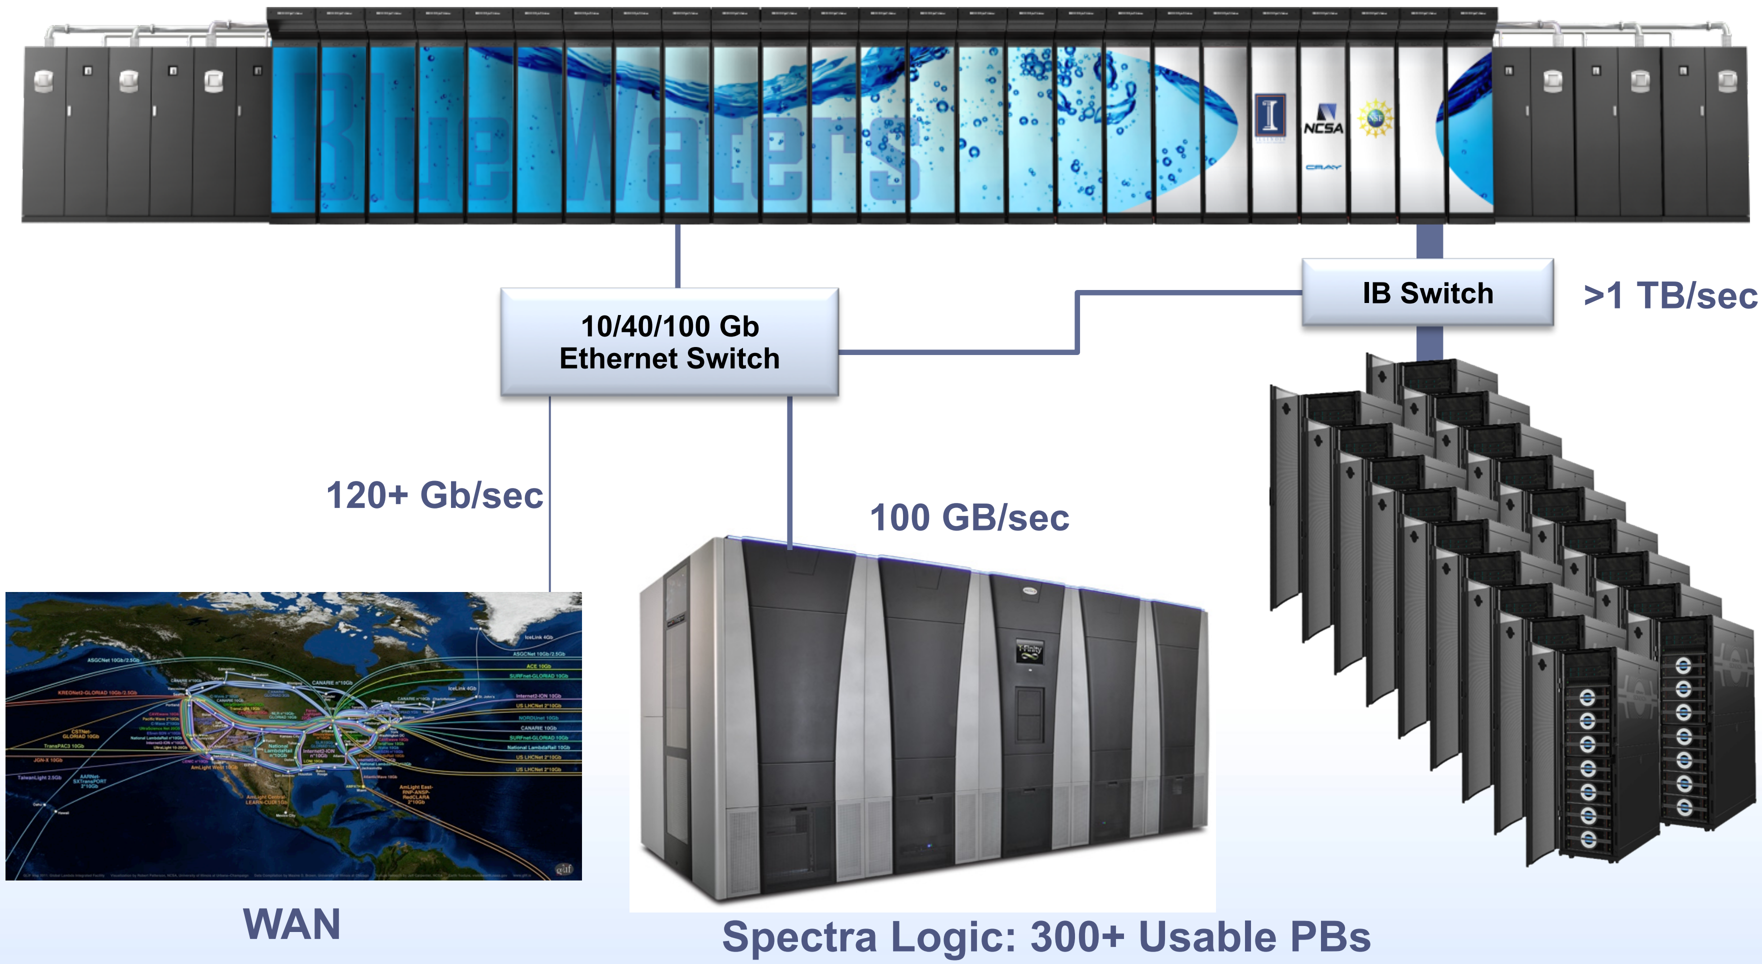

Blue Waters System Statistics

Total Cabinets	**> 300**
Total XE6 Cabinets	> 235
Total XK6 Cabinets	> 30
Compute Nodes	> 25,000
Usable Storage Bandwidth	> 1TB/s
Aggregate System Memory	> 1.5 Petabytes
# Spinning Disks	> 17,000
# Memory DIMMs	> 190,000
# AMD processors	> 49,000
Integrated Near-Line Storage	Up to 500 PB
Bandwidth to Near-Line Storage	100 Gb/s
External Network Bandwidth	Up to 300 Gb/s

300+PB HPSS Archive & Data Movement

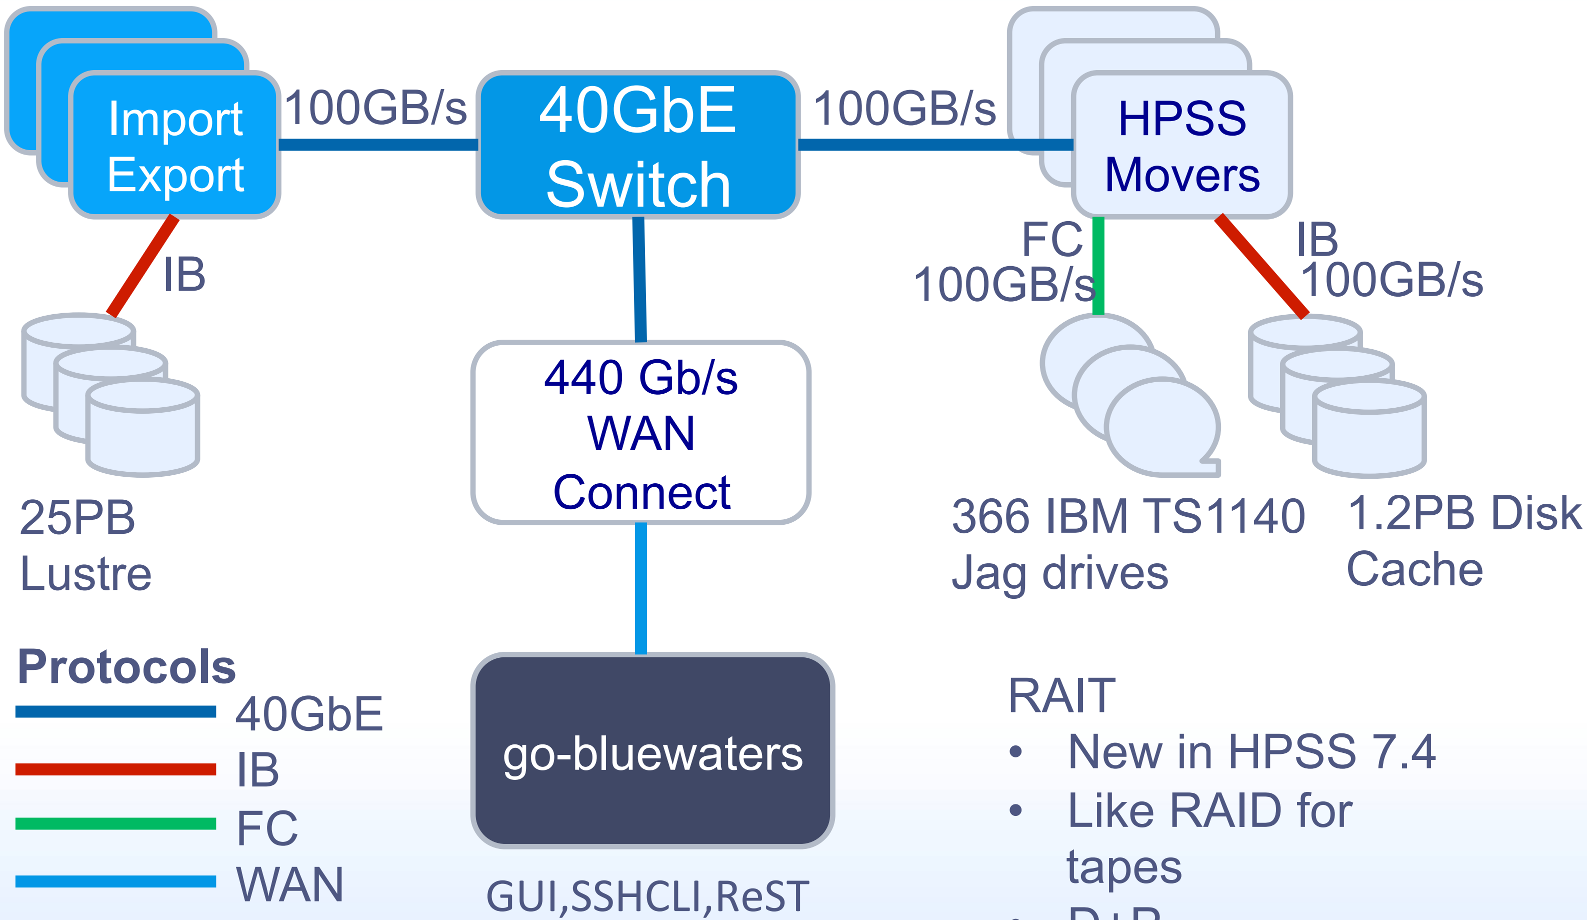

RAIT

- New in HPSS 7.4
- Like RAID for tapes
- D+P configurations

Funding Agencies Viewpoints

- NSB Task Force on Data Policy, Dec 2011:
  - "To address the challenges associated with increasing scale, scope, and complexity of data, each science and engineering research community should take the responsibility for determining its own standards and conventions for data stewardship and for coordination across the research enterprise. Funding agencies and stakeholder communities must partner together during data policy development so that recommendations can be implemented by each science and engineering research community."
- NSF Data Management Plans for awards
- Interagency Public Access Coordination
  - March 2012 report to Congress
  - Interagency Working Group on Digital Data (IWGDD)
  - Task Force on Public Access to Scholarly Publications

Viewpoint Shift…

- Dark Energy Camera
  - $20M/PB
  - Data, catalogs, data rights/access for extending research
  - Probable long term interest and support
- Large Synoptic Survey Telescope
  - ~$4M/PB
  - Data, catalogs, shared access for extending research
  - Probable long term interest and support
- Large scale simulation instruments
  - ~$1M/PB
  - Data, limited to life of allocation
- What about shared reusable simulation data …

NRC Viewpoint in Climate Modeling

- "A National Strategy for Advancing Climate Modeling"
  - Recommendation 10.3: The United States should support transformational research to **bring analysis to data** rather than the other way around in order to make the projected data volumes useful.
  - Recommendation 10.4: The **data-sharing infrastructure** for supporting international and national model intercomparisons and other simulations of broad interest—including archiving and distributing model outputs to the research and user communities—**is essential** for the U.S. climate modeling enterprise and should be supported as an operational backbone for climate research and serving the user community.

G8 Initiative ExArch

- Climate Analytics on Distributed Exascale Data Archives
  - ANR, DFG, NSERC/CRSNG, NSF, RCUK
  - “This international project targets the rapidly growing demands of climate science data management as models increase the precision with which they depict spatial structure and the completeness with which they describe a vast range of physical processes. The ExArch project is principally a framework for the scientific interpretation of multi-model ensembles at the peta- and exa- scale. It applies a strategy, a prototype infrastructure and demonstration usage examples in the context of the imminent CMIP5 archive, which will be the largest of its kind ever assembled in this domain. It will attach the ExArch framework to the CORDEX experiment, pushing beyond CMIP5 in resolution, albeit at regional scale.”

NSF Award Abstract 1119308

Areas for Investigation/Next Steps

- Resource management
  - Focus on data products as an alternative to cycles
  - Volume, velocity and variety
  - Usability, maintainability and usefulness of the data products
- Distributed data and tools
  - Campus <=> mid-range environments <=> high end environments
- Data rights and access
  - Recognition and attribution
- Longevity and sustainability
  - Metrics for evaluation of dataset value
  - Sustainable, long term storage and access
- Reproducibility of results
- Integration with observational/experimental data

Large Scale Collaborations with Shared Data

"It really IS possible for hundreds of people at tens of institutions to work together in a non-cat-herded manner (few if any constraints on subjects or projects) to produce excellent science."

James Gunn,

SDSS From Asteroids to Cosmology,

Aug 2008