

Scientific computing at BNL

Eric Lançon

Brookhaven National Laboratory

April 4, 2017

70 YEARS OF
DISCOVERY

A CENTURY OF SERVICE



BROOKHAVEN
NATIONAL LABORATORY

LHCOPN-LHCONE meeting - BNL

BNL is a Data Driven Science Laboratory

- **BNL is dominated by Data-rich Experimental Facilities:**
 - **RHIC** - Relativistic Heavy Ion Collider - High Energy Physics Experiment supporting over 1000 scientists world wide
 - **NSLS II** - Newest and Brightest Synchrotron in the world opened in the world, supporting a multitude of scientific research in academia, industry and national security
 - **CFN** - Center for Functional Nanomaterials, combines theory and experiment to probe materials
- **BNL supports large scale Experimental Facilities:**
 - **STAR/PHENIX** - the RHIC experiments
 - **LHC Atlas** - Largest Tier One Center outside CERN, data storage and processing
 - **ARM** - Atmospheric Radiation Measurement Program - Partner in multi-side facility, operating its external data Center

RHIC



NSLS II



CFN



LISF




SDCC





Some of the facilities on BNL site




RHIC: 2 main experiments PHENIX & STAR




Some of the current and future experiments of the physics department



Dark Energy Survey



Large Synoptic Survey Telescope (LSST)



Scientific Data and Computing Center




- **Service Operation For:**
RHIC, LHC Atlas, BER ARM, LQCD, RIKEN, BES Center for Functional Nano Materials, National Synchrotron Light Source II, National Nuclear Data Center, Simons Foundation
- **17 Data Services, 1000's of users**

SDCC end of 2016

BNL Cloud

Used by OSG, ATLAS and various groups from BNL outside physics department



KNL cluster



HTC CPU = 15,000 cores
Disk storage = 15 PB



(2017 : USQCD cluster)



Big Data Network Fabric

Current fabric interswitch links = 3 Tbps.
Scalable to Pbps.



Tape Mass Storage




HPC Institutional Cluster




CFN TEM

Institutional Cluster




HTC CPU = 35,000 cores
Disk storage = 25 PB





RHIC Facilities

SDCC in numbers

- **70k CPU cores (~100k in 2017)**
 - 2017: HPC **1/3** of capacity



- **~50 PB of disk storage**
 - of various technologies
- **~100 PB of tape storage**
 - Largest HPSS tape library in the US, 4th worldwide ⁽¹⁾
- **2x100 Gbps connection to ESNet**
 - Onsite ESNet support



(1) http://www.hpss-collaboration.org/learn_who_petabyte_data.shtml

High Throughput Parallel Archiving

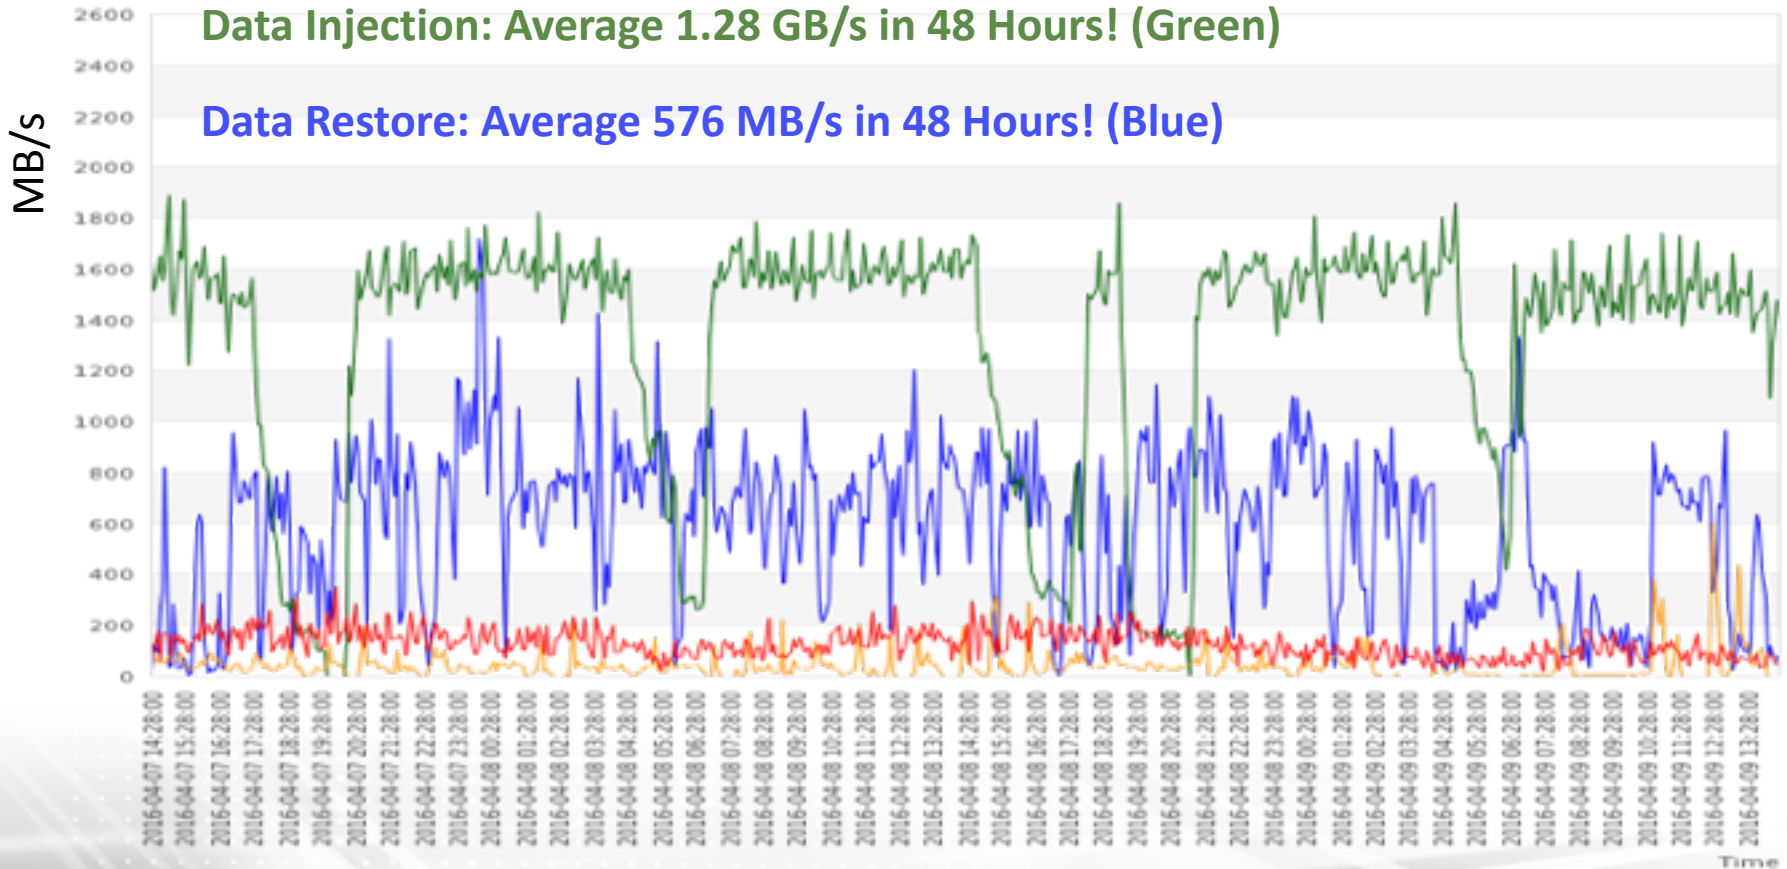
RHIC RUN 16 - STAR

STAR Data Transfer View
Range: 2016-04-07 14:28:00 - 2016-04-09 14:23:00
RAW Write: 215.3 TB, 62303 files, avg size: 3.54 GB, avg rate: 1.28 GB/s
DST Write: 20.51 TB, 7218 files, avg size: 2.91 GB, avg rate: 124.44 MB/s
RAW Read: 95.03 TB, 23407 files, avg size: 4.16 GB, avg rate: 576.66 MB/s
DST Read: 7.24 TB, 64652 files, avg size: 117.45 MB, avg rate: 43.94 MB/s

— RAW Staging
— RAW Write
— DST Staging
— DST Write

Data Injection: Average 1.28 GB/s in 48 Hours! (Green)

Data Restore: Average 576 MB/s in 48 Hours! (Blue)



SDCC and RHIC and ATLAS experiments

- **The RHIC Tier 0**

- Store and process data from RHIC experiments
- Provide analysis means for 1'200 users
- Long term data preservation
- Simulation resources for future programs (sPHENIX & EIC)


- **The US ATLAS Tier 1**

- ~25% of ATLAS Tier 1 computing capacity worldwide
- Store RAW data from LHC and from simulation
- Distribute data to the 4 **US Tier 2** sites + analysis site (**SLAC**)
- Analysis center for US physicists
 - From 41 institutes (incl. 4 Nat. Labs)
 - 600 physicists, 190 PhDs




CPU delivered to ATLAS by sites in 2016

WallClock consumption in seconds





The first 15

WallClock consumption in seconds



ATLAS: Data transfers (FTS)

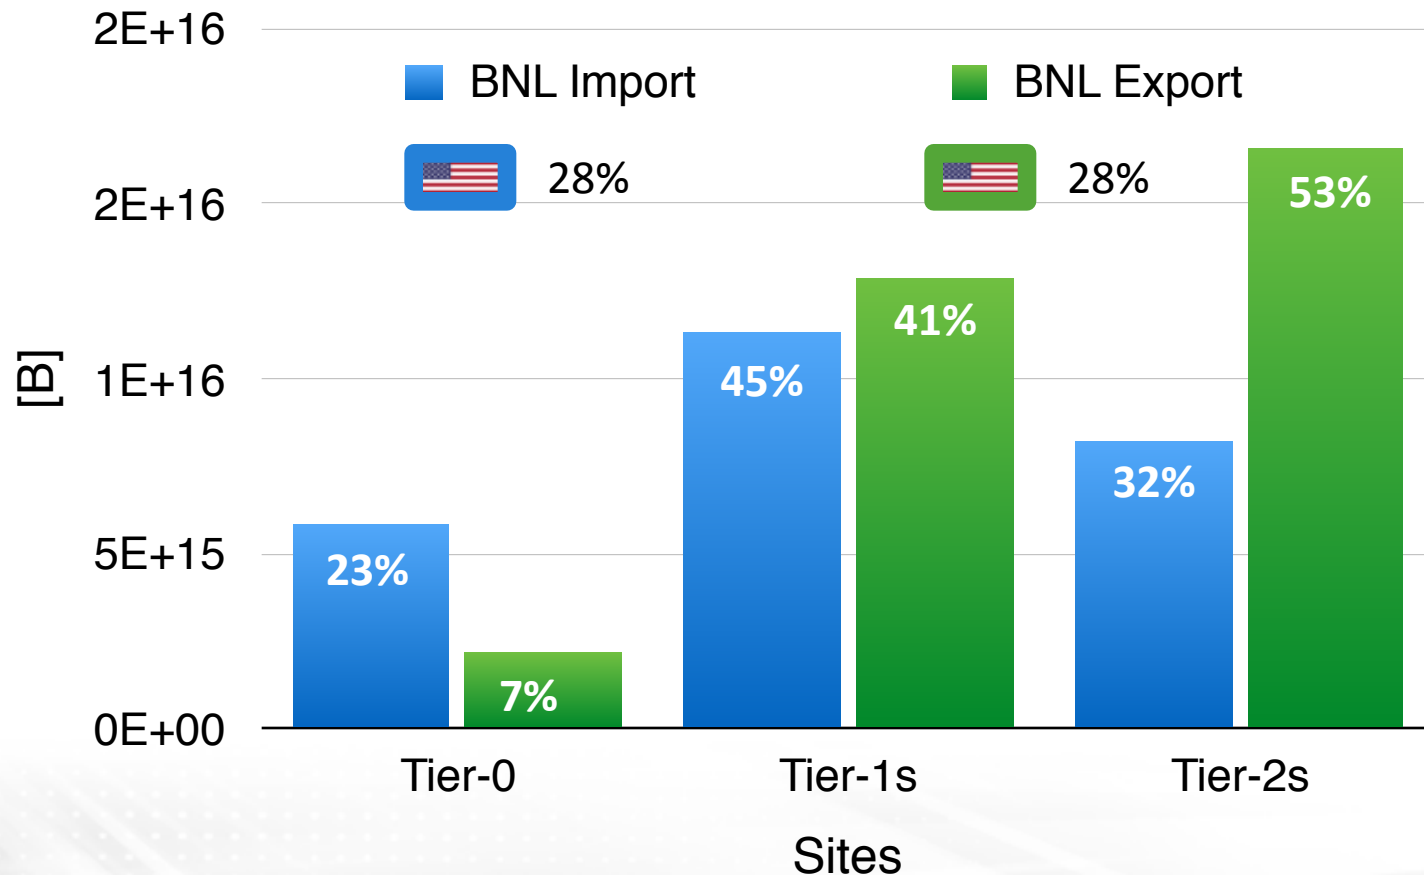
- Over a year (since Apr. 2016)
 - Data import : 25 PB
 - Data export : 35 PB
 - 175 PB processed



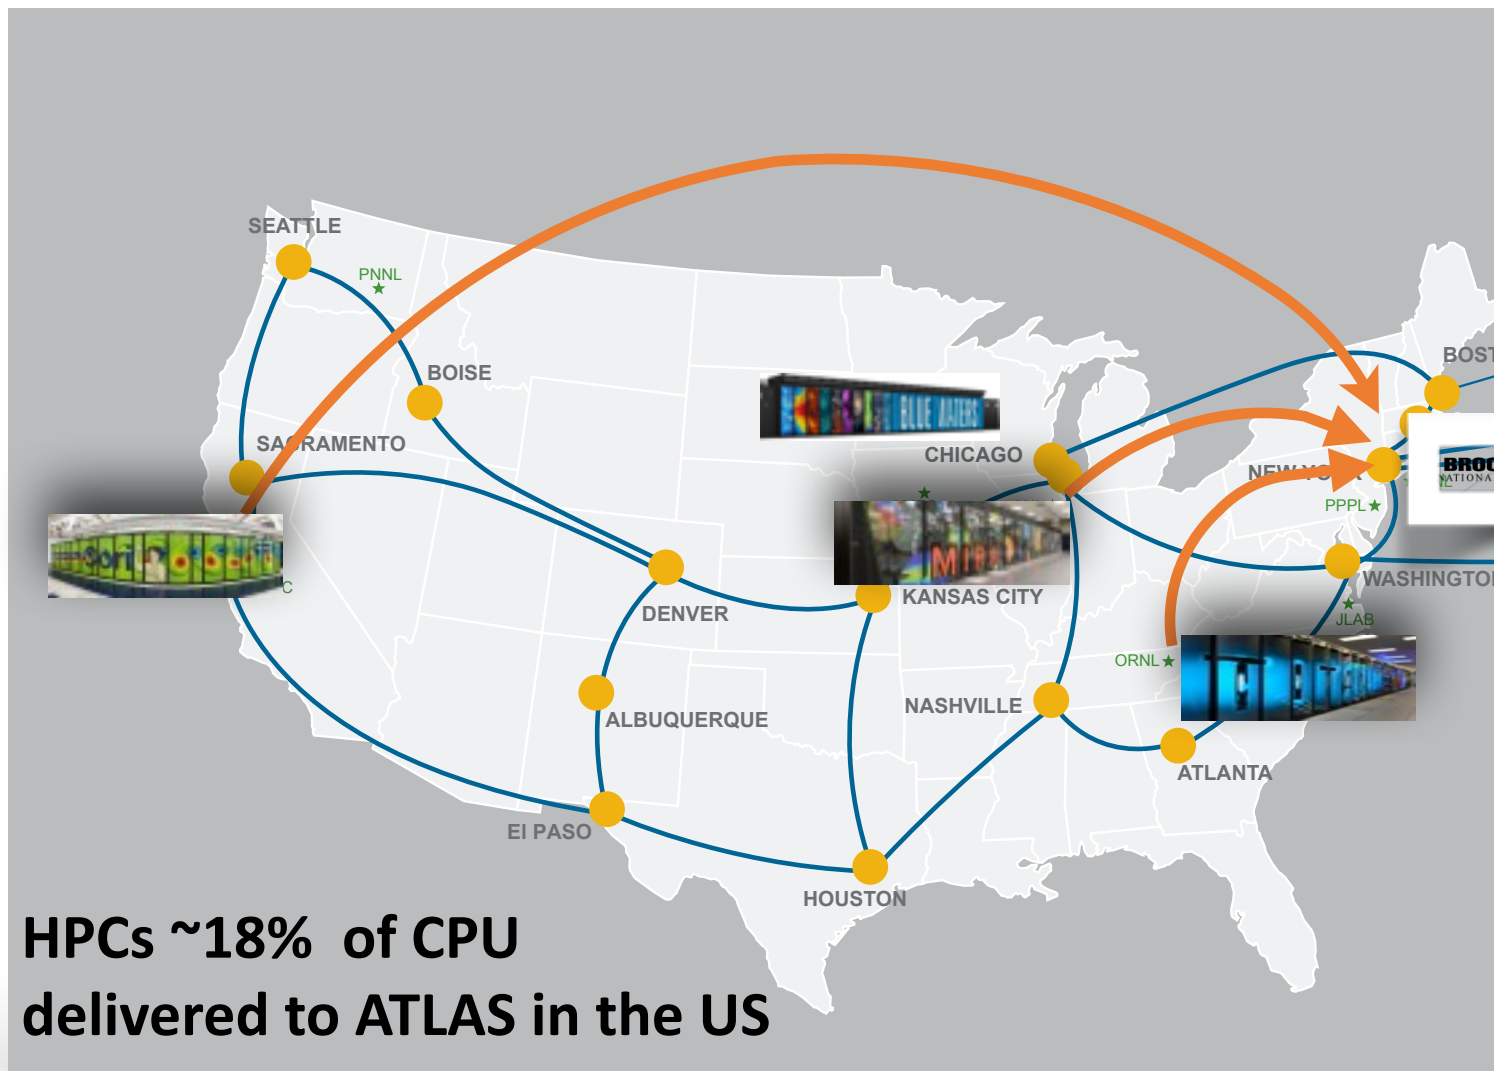
- Department of Energy Office of Science National Labs
- Argonne National Laboratory (Chicago, IL)
- Brookhaven National Laboratory (Upton, NY)
- Lawrence Berkeley National Laboratory (Berkeley, CA)
- Los Alamos National Laboratory (Los Alamos, NM)
- Pacific Northwest National Laboratory (Richland, WA)
- SLAC National Accelerator Laboratory (Menlo Park, CA)
- Thomas Jefferson National Accelerator Facility (Charlottesville, VA)
- Lawrence Livermore National Laboratory (Livermore, CA)
- ORNL (Oak Ridge National Laboratory) (Knoxville, TN)
- PNNL (Pacific Northwest National Laboratory) (Richland, WA)
- PSF (Purdue Supercomputing Facility) (West Lafayette, IN)
- SDSC (San Diego Supercomputer Center) (San Diego, CA)
- SNL (Sandia National Laboratories) (Livermore, CA)
- UCSD (University of California, San Diego) (San Diego, CA)
- U.S. National Science Foundation (Washington, DC)

ATLAS: Data transfers (FTS)

- Over a year (since Apr. 2016)



ATLAS: BNL and HPCs



**HPCs ~18% of CPU
delivered to ATLAS in the US**

Data produced at HPCs stored mostly at BNL

and small fraction through FTS

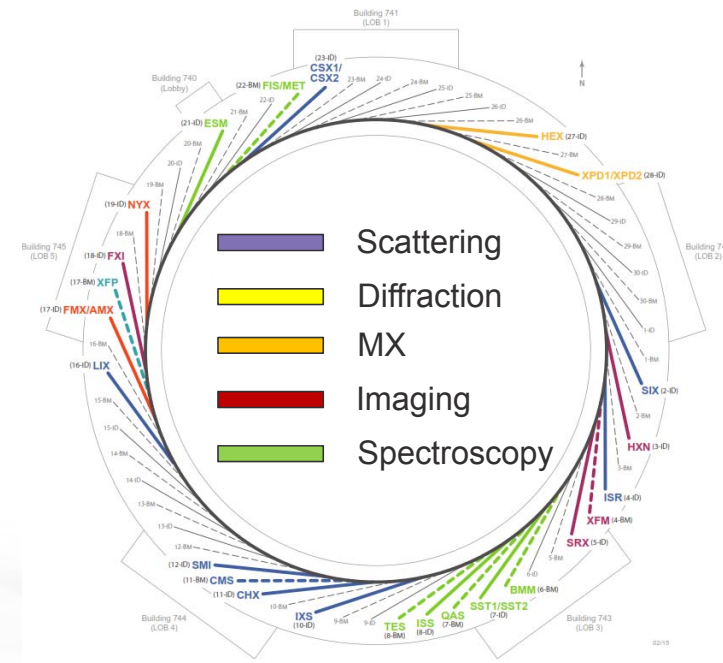
National Synchrotron Light Source II (NSLS-II)




Visit tomorrow morning at coffee break

NSLS-II: Science Driven Data Challenges

- **Brightest synchrotron in the world**
 - 2,000 users per year (4,000 in FY17)
 - 2/3 from Universities
 - 16 beamlines operational at present, full build out 60
 - Supports Multitude of Sciences - Physics, Chemistry, Biology, Climate, Bio-Medical ...
- **Data Challenges:**
 - NSLS II - Coherent Hard X-Ray (CHX) Beamline 4.5 GB/s sustained data rates
 - NSLS II - Hard X-Ray Nanoprobes (HXN) 50 GB/s sustained, 1 – 5 TB/s in burst
 - High Data Velocity, Real Time, Correlated Events, Reliability Requirements
- **Needed data driven real time steering of experiments**




Estimated NSLS-II Storage Needs



- Estimated for the 28 NSLS-II beamlines under development
- Based on current information on operating beamlines
- Takes into account estimated increases in both detector and beamline operational reliability.

New Data Center (FY22)

NSLS-I Building



SPACE LEGEND

	PROGRAM NET AREA (SQF)	ESTIMATED AREA (SQF)
DATA HALL A	600	60%
DATA HALL B	300	200
DATA HALL C	100	100
DATA HALL D	200	400
NETWORK EQUIPMENT	100	110
OFFICES	100	110
MECH/ELEC	200	200
IT TECH AREA	50	50
STORAGE	100	100
LOADING / UNCRATING	50	200
OPERATIONS	10	10
TOTAL NET AREAS	600	900

Phase I : 60% Space Plan, 3.6 MW

~3.5 x today's capacity

DATE:	01 AUGUST 2016
PROJECT:	CORE FACILITY REVITALIZATION (CFR) BUILDING 725
DISCIPLINE:	ARCHITECTURE
SHEET:	2.4.1.2 FIRST FLOOR PLAN

Summary and Outlook

- Successfully meet computational needs of US ATLAS and RHIC programs
 - a world class big-data processing facility
 - close to the users and responsive to their needs
- Expanding capacities for
 - integration of National Nuclear Data Center, Atmospheric Radiation Measurement, NSLS-II,...
- Rapid growth adding HPC capabilities
- Challenging and exciting years in front of us!