Open Science Grid

UCSD

State of OSG

Frank Würthwein
OSG Executive Director
UCSD/SDSC

Open Science Grid

UCSD

Advance Open Science through distributed High Throughput Computing

Let's parse this OSG statement of purpose a little ….

Open Science

- All of open science irrespective of discipline
- Advance the maximum possible dynamic range of science, groups, and institutions
  - From individual undergraduates to international collaborations with thousands of members.
  - From small colleges, museums, zoos, to national scale centers of open science.
- Advancing this entire spectrum requires us to have a diversified portfolio of services

OSG Accounting by Field of Science

Order here is from most active to least active.

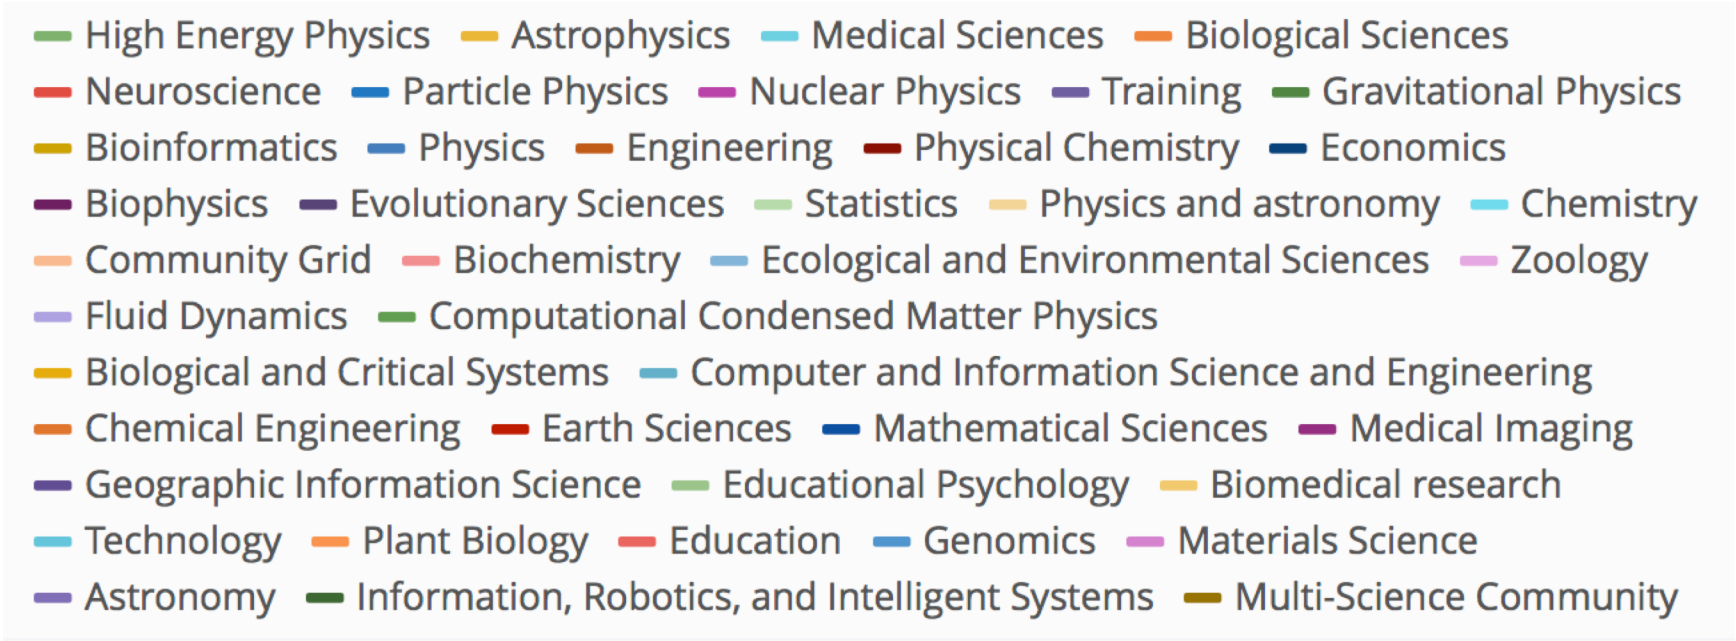

Category is based on self-identification … i.e. people say what they do.

OSG Accounting by Organization

People self-declare their organization

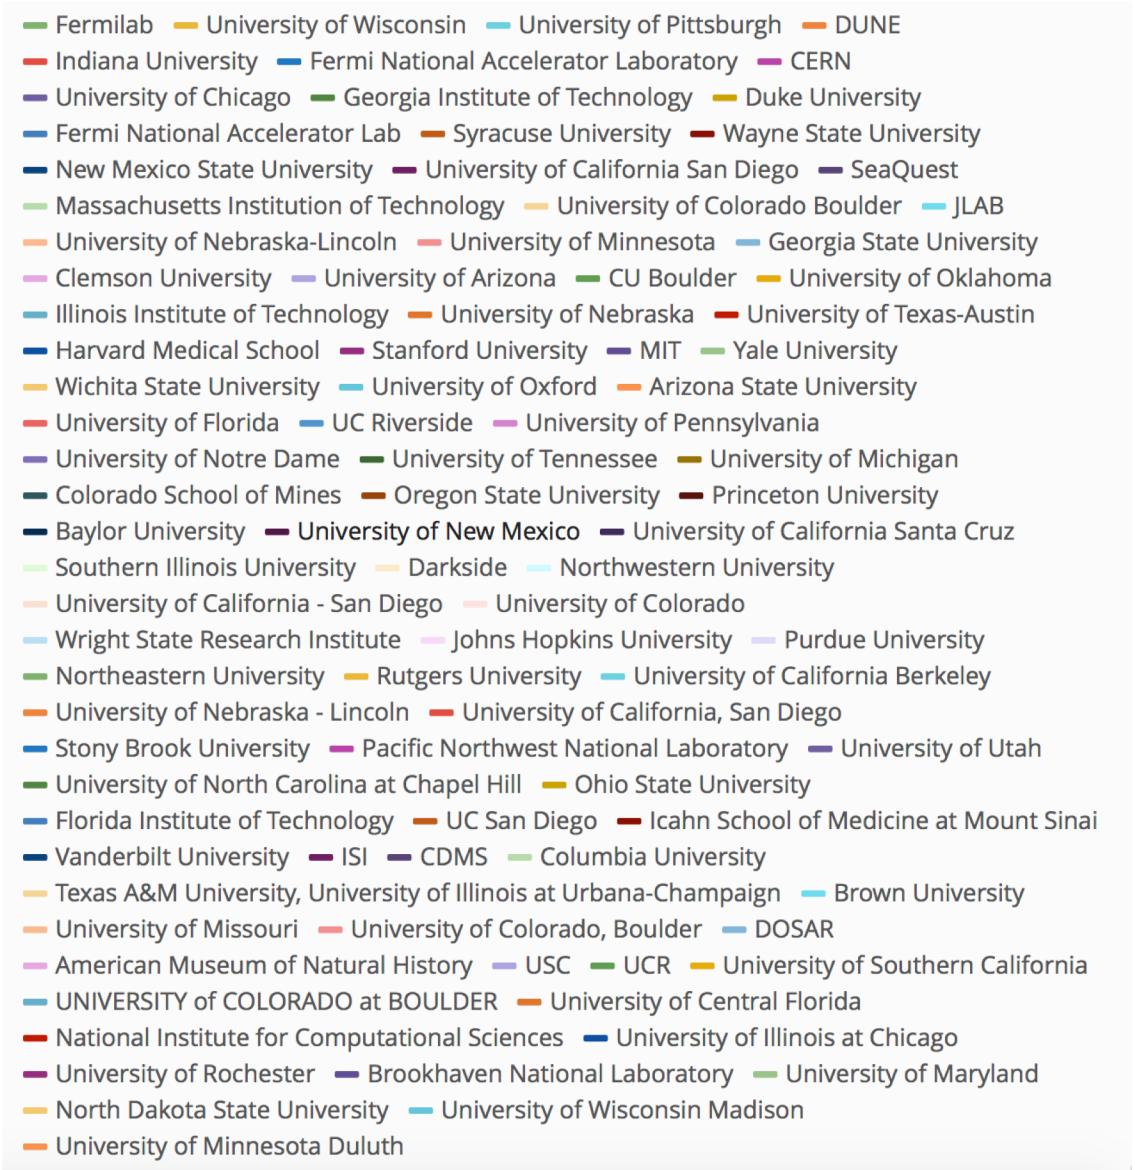

OSG serves 4 distinct groups

- The individual researchers and small groups on OSG-Connect
- The campus Research Support Organizations
  - Teach IT organizations & support services so they can integrate with OSG
  - Train the Trainers (to support their researchers)
- Multi-institutional Science Teams
  - XENON, GlueX, SPT, Simons, and many many more
  - Collaborations between multiple campuses
- The 4 "big science" projects:
  - US-ATLAS, US-CMS, LIGO, IceCube

distributed High Throughput Computing (dHTC)

- The challenge in successful dHTC is two-fold:
  - Separate a big computing problem in many individually schedulable small problems.
  - Minimize your requirements in order to maximize the raw capacity that you can effectively use.
- We teach researchers how to meet these challenges.
- We take your many small problems, and schedule their successful execution.
- We allow you to curate & publish your software & data and deliver them to you at runtime.

Ingenious Parallelism

GPU computing is part of dHTC

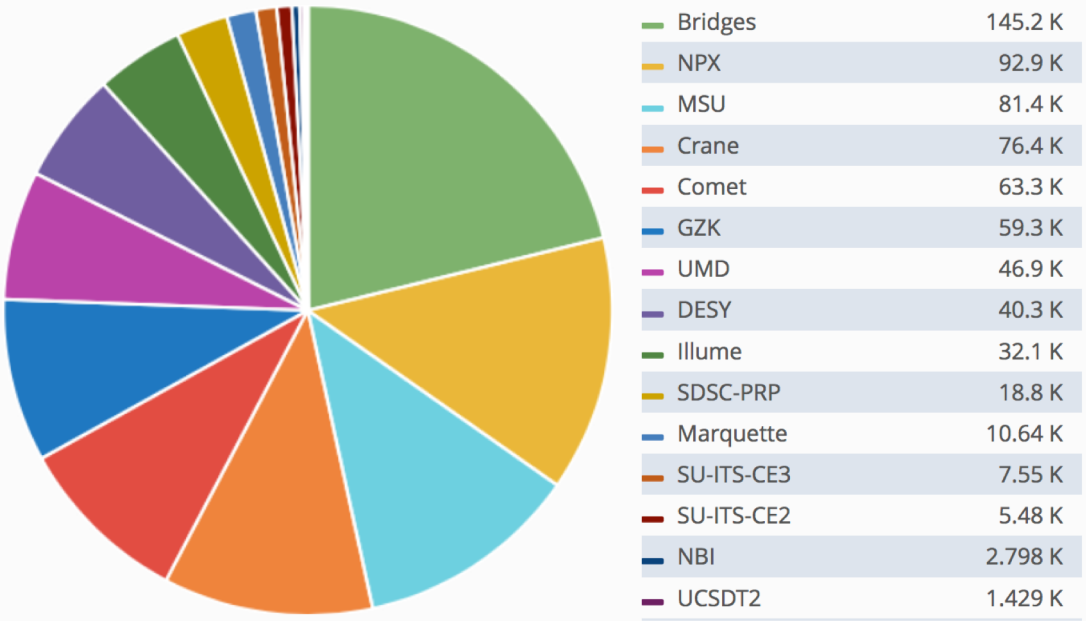

Almost all of it consumed by IceCube

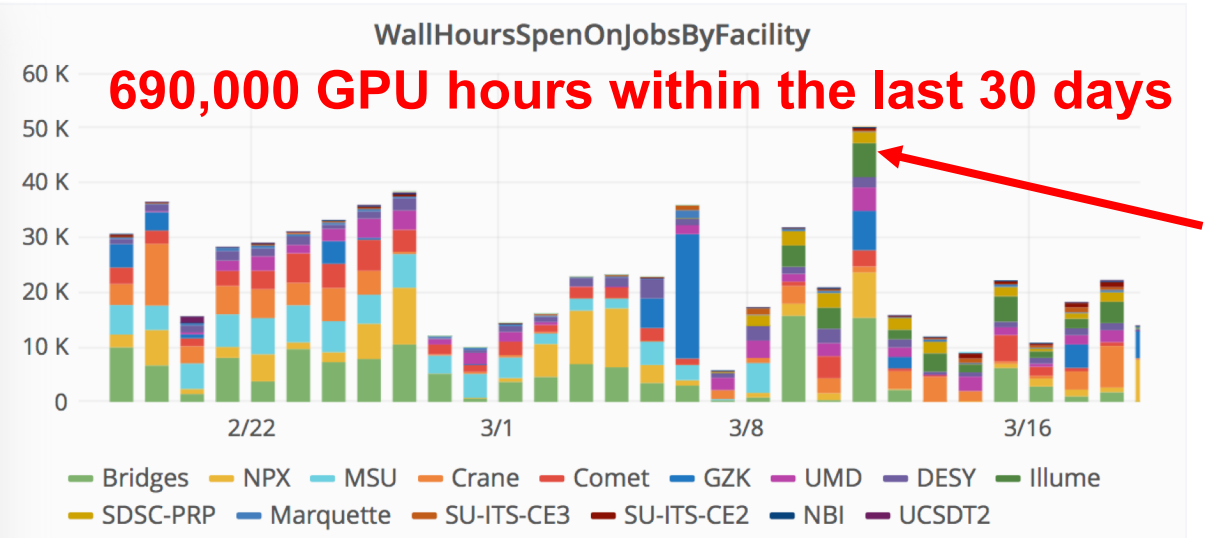

Peaking at 2,000 GPUs concurrently.

Advance = Growing effective Capacity

- 10 years ago, OSG provided an integrated software stack that organizations had to deploy in order for their researchers to benefit from OSG.
  - We learned the hard way that this is too high a bar to jump across for most researchers and institutions.
- Today, we offer to operate services ourselves in order to minimize the threshold to entry for everybody else.

Our Aspirational Goal:

A campus Research IT Organization should not have to learn anything ”non-standard” in order to have their researchers benefit from OSG, or have their resources be available via OSG.

Well, …. we got some ways to go before we reach that goal …

Services OSG Operates (I)

- OSG-Connect, a submission host for individual researchers.
  - You get an account, and we teach you how to use OSG.
- A Compute Federation
  - gWMS system that submits glideins to 300+ CEs globally.
  - HTCondor pools for multiple VOs
  - HTCondor Annex to integrate commercial cloud resources into your HTCondor pool.
  - CVMFS infrastructure to support software publication, including a module environment supporting >200 software modules.
  - Compute Elements to integrate clusters into OSG

OSG Compute Federation

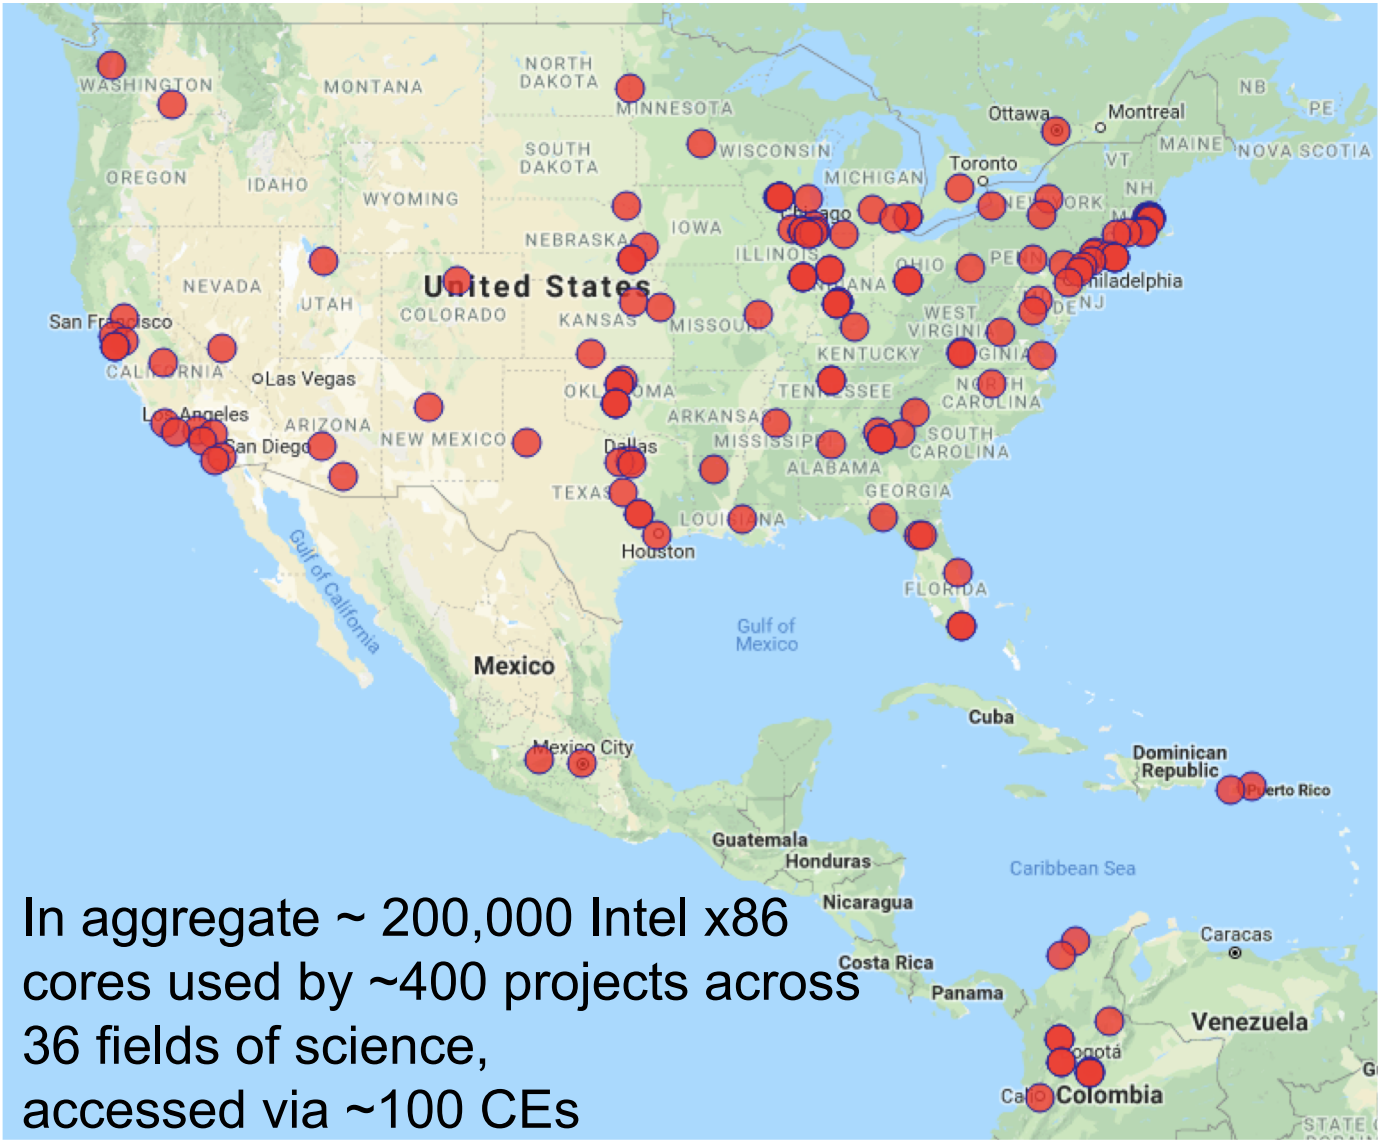

In aggregate ~ 200,000 Intel x86 cores used by ~400 projects across 36 fields of science, accessed via ~100 CEs

Services OSG Operates (II)

- A Data Federation
  - 1.5 PB disk space on our submission host that is exported into OSG
    - Individual undergraduates have storage allocations
  - An xrootd infrastructure that anybody can add their own data origin(s) to.
  - Data publication in our federated namespace using CVMFS.
  - Xrootd caches to hide data access latencies and reduce network traffic.
- Policy-based data placement (Rucio)

We have created the illusion of a read-only global file system for all of open science.

OSG Data Federation Origins

We want your Data in our Federation !!!

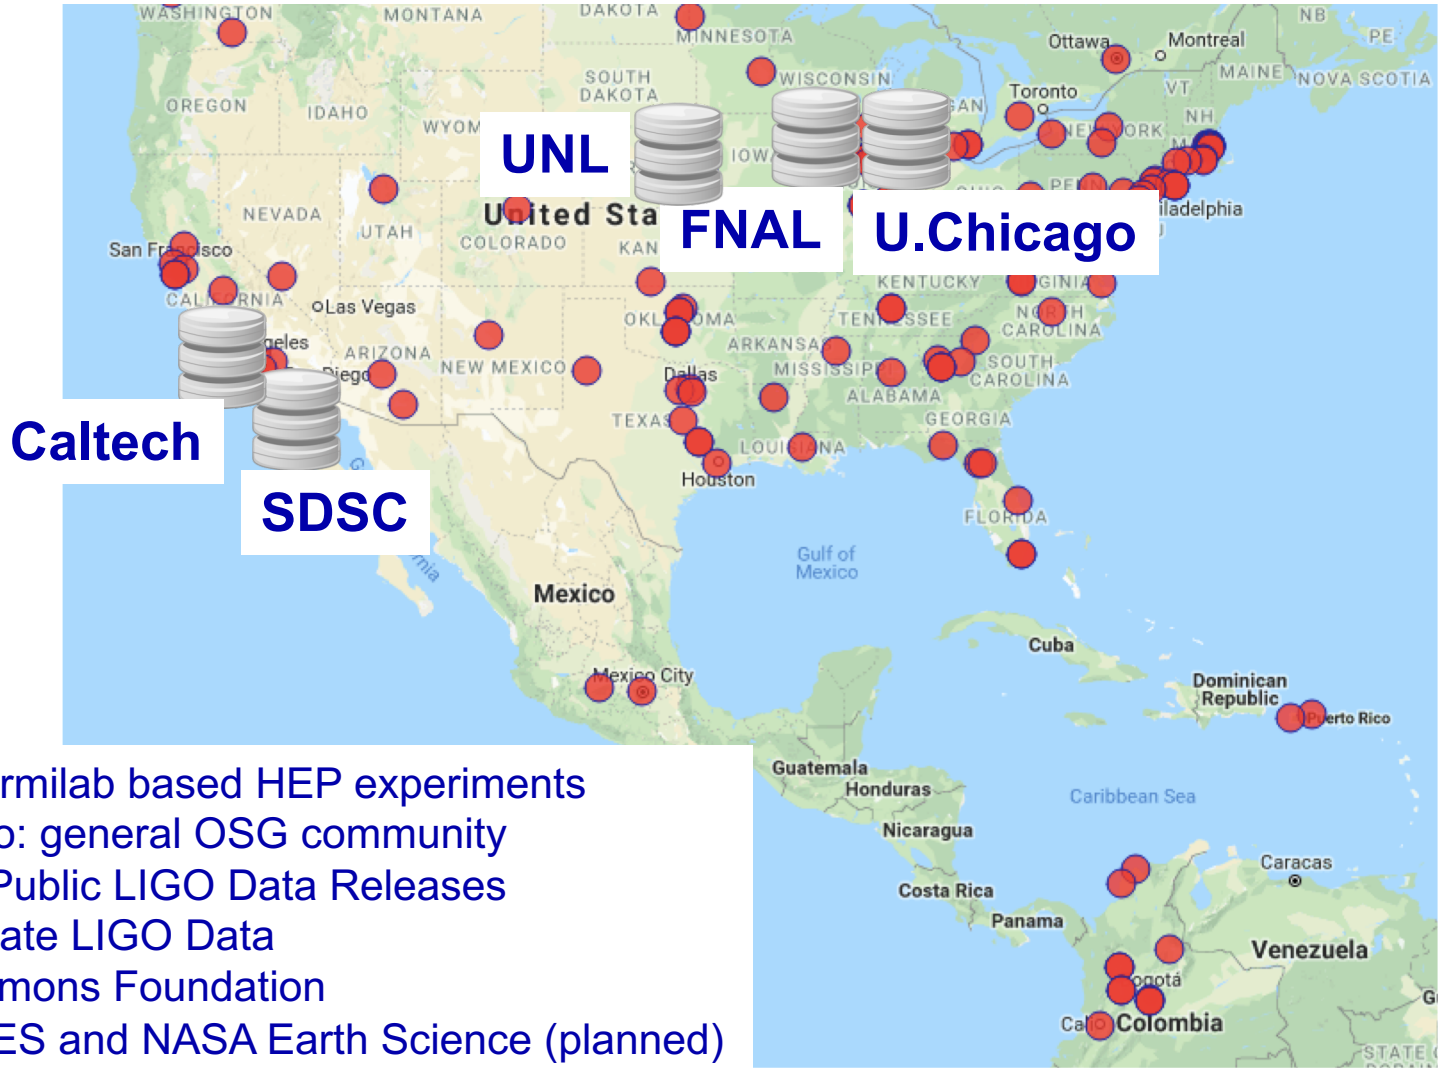

FNAL: Fermilab based HEP experiments
U.Chicago: general OSG community
Caltech: Public LIGO Data Releases
UNL: Private LIGO Data
SDSC: Simons Foundation
NCSA: DES and NASA Earth Science (planned)

Caches in network & at endpoints

Open Science Grid

UCSD

6 Data Origins and 9 Data Caches

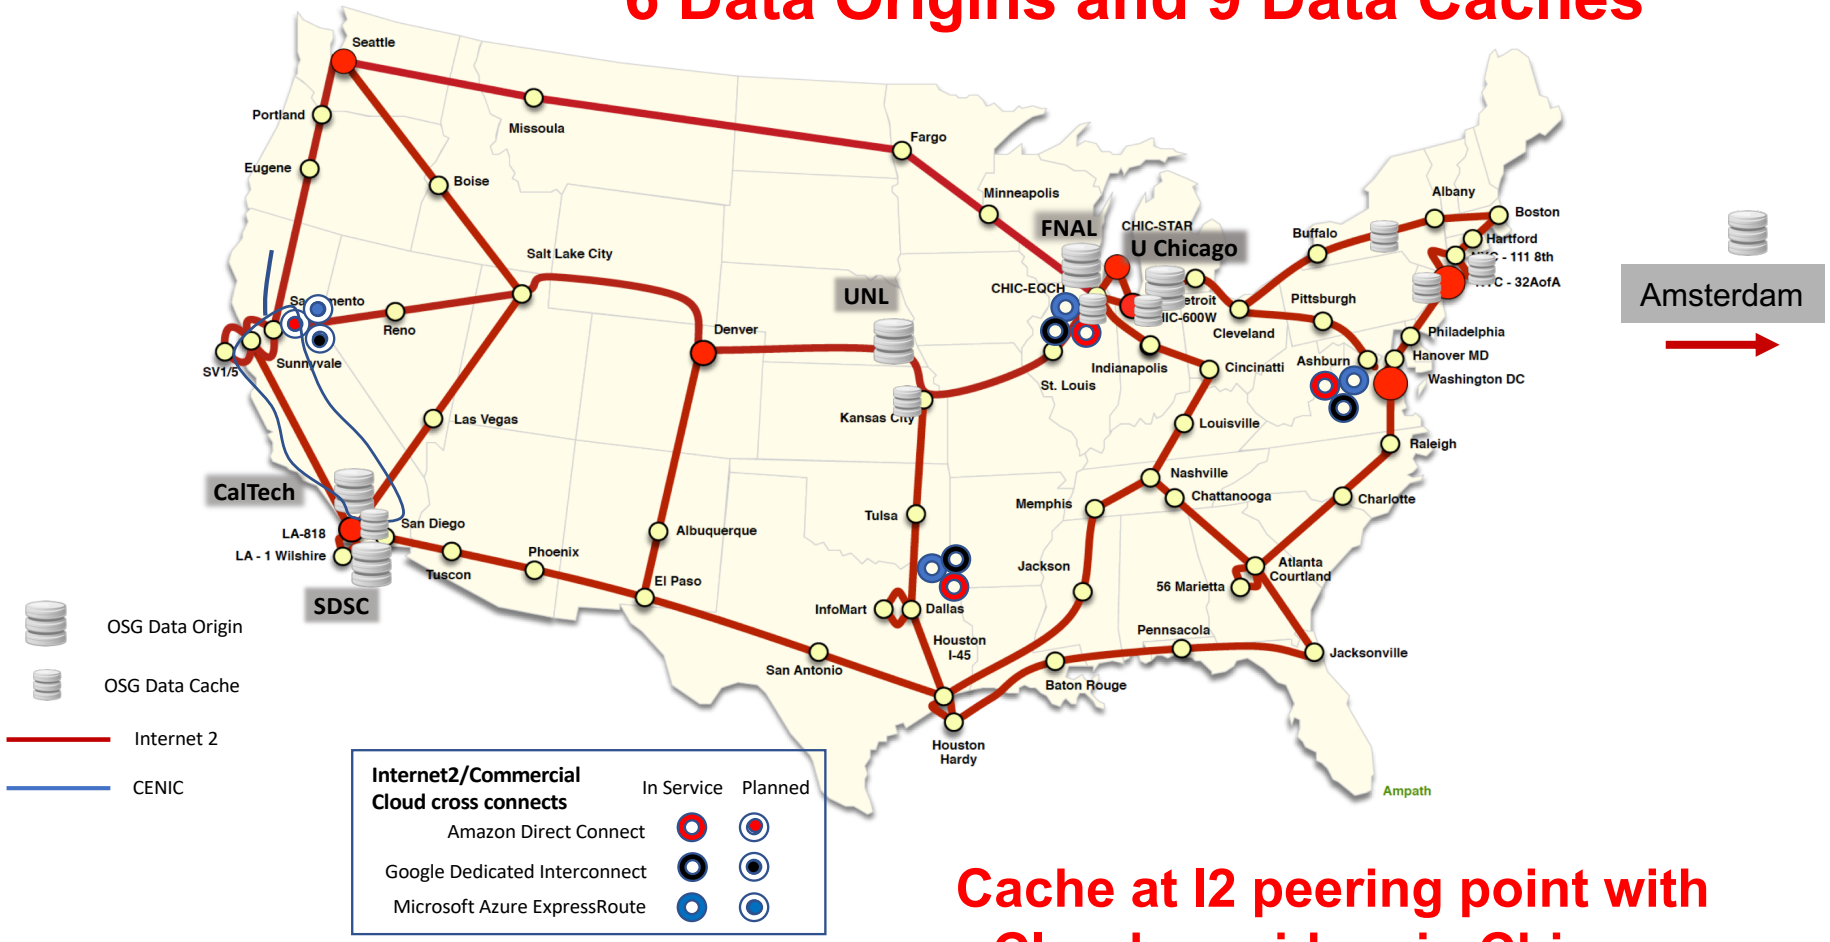

Cache at I2 peering point with Cloud providers in Chicago to support elastic scale out.

Data Federation used by individual researchers and science collaborations

May 2018

XRootD StashCache By Logical Directory

Directory	Bytes
/user/jeanjack	660.125TB
/pnfs/fnal.gov/usr/minerva	197.651TB
/pnfs/fnal.gov/usr/des	159.64TB
/user/ligo	14.078TB
/user/wpoehlm	12.513TB

September 2018

XRootD StashCache By Logical Directory

Directory	Bytes
/gwdata/O1	158.889TB
/pnfs/fnal.gov/usr/minerva	98.705TB
/user/ligo	65.925TB
/pnfs/fnal.gov/usr/des	23.495TB
/user/jeanjack	17.567TB

"Genomics"
Jeanjack
wpoehlm

Public LIGO
gwdata/01

Private LIGO
ligo

Astronomy
des

HEP
minerva

Open Science Grid

UCSD

Last 30 days

Directory	Bytes
/pnfs/fnal.gov/usr/dune	390.114TB
/gwdata/O1	180.208TB
/pnfs/fnal.gov/usr/minerva	99.352TB
/pnfs/fnal.gov/usr/des	63.204TB
	58.211TB

Last 60 days

Directory	Bytes
/pnfs/fnal.gov/usr/dune	396.578TB
/pnfs/fnal.gov/usr/minerva	395.986TB
/gwdata/O1	295.03TB
/pnfs/fnal.gov/usr/des	182.369TB
/user/ligo	104.744TB
/user/dteam	55.197TB
/pnfs/fnal.gov/usr/nova	31.768TB

OSG offers "service bundle" to Campus IT organizations

- **Collaboration**: GaTech PACE, LIGO, OSG @ UChicago, UCSD
- Cluster setup - node provision & config
- Cluster setup - HTCondor batch, login
- OSG HTCondor Compute Element
- OSG Worker Node Client
- OSG Frontier-Squid, CVMFS
- OSO-LIGO Job Submission Service
  - integrated grid & local submissions
- OSG GridFTP Server
- StashCache - authenticated and un-auth
- Integration with existing LIGO LDAP and Shibboleth Services
- LIGO application validations
- gsi-ssh clients & troubleshooting
- Site-specific support documentation, configuration management (salt), staff training

Goal:
Help GaTech deploy a cluster that is part of LIGO production.

OSG helps Science Collaborations develop & debug their Infrastructure

- OSG operates daily satellite data collection and archival for SPT-3G
- In February, poor transfers from USAP to UChicago
- Consulting with OSG, UChicago & UCAR network experts.
- Feb 26, Dave Feldt (USAP) shut down the InterNAP link (one of two commercial uplinks for USAP) and saw 10x improvement in performance.

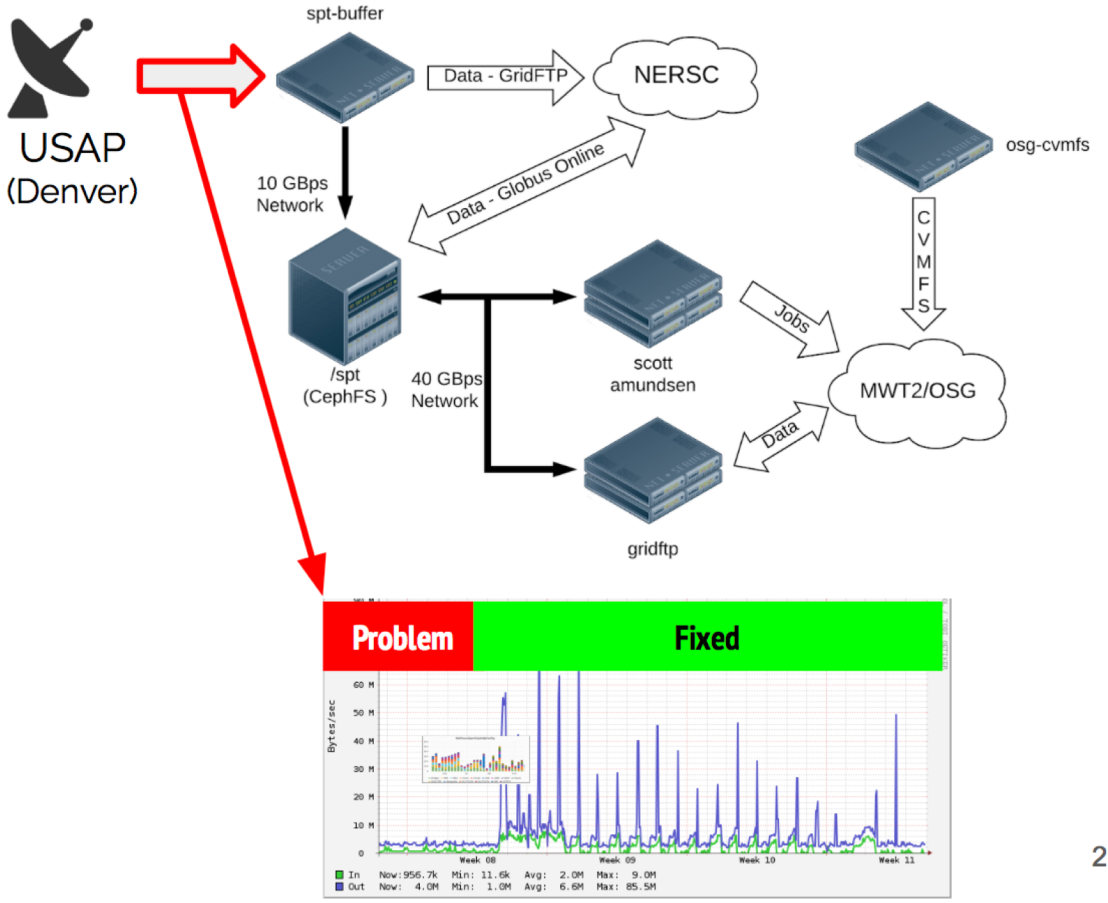

SPT-3G Data Throughput Trouble Shooting

OSG helps Science Collaborations develop & debug their Infrastructure

- OSG operate
  data collecti
  SPT-3G
- In February,
  from USAP t
- Consulting w
  UChicago &
  experts.
- Feb 26, Dave
  shut down th
  (one of two c
  uplinks for U
  improvemen

SPT-3G

OSG enabling international collaboration (last 30 days)

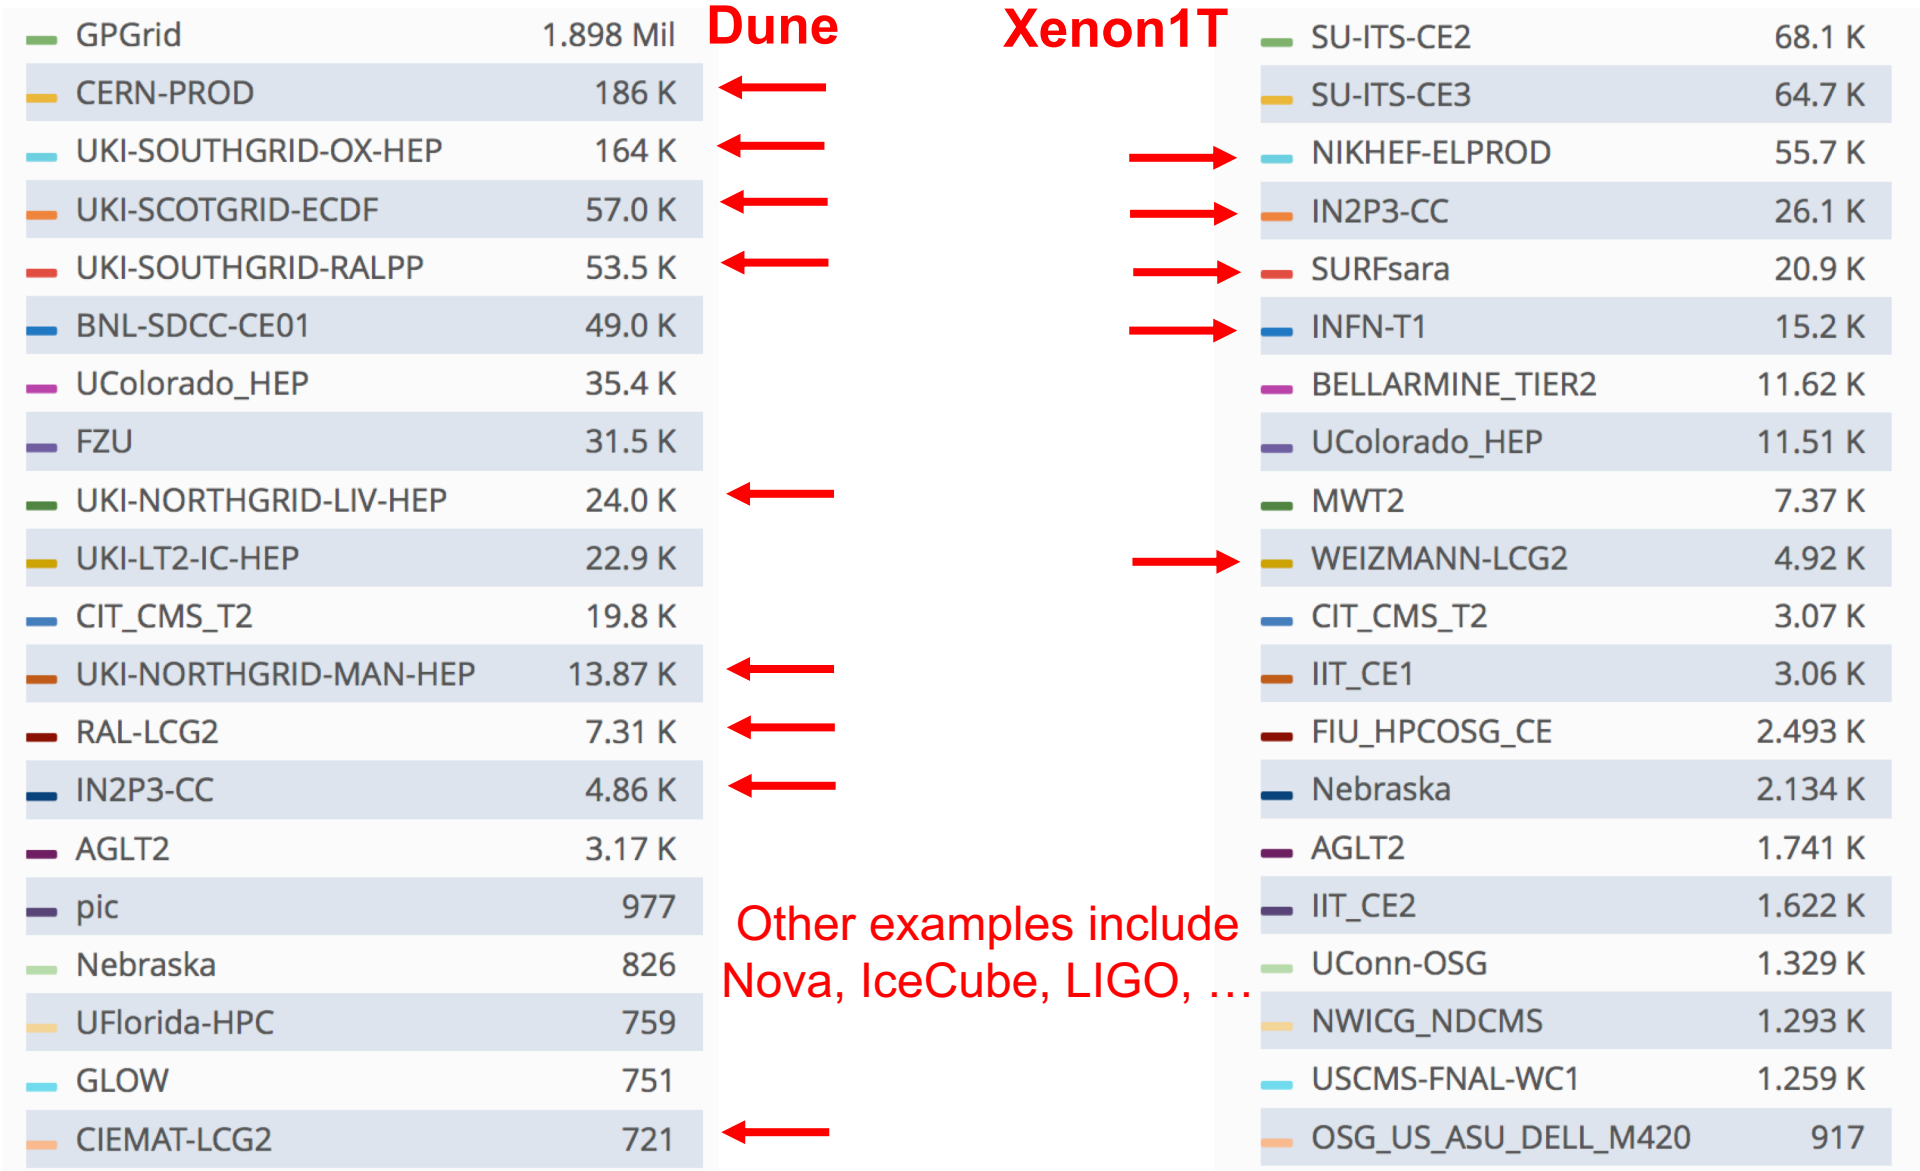

Summary & Conclusion

- OSG's objective is to "Advance Open Science through distributed High Throughput Computing"
- OSG thinks of its science stakeholders in terms of 4 categories:
  - Individual Researchers
  - Campus Research Computing Organizations
  - Multi-campus Science Teams
  - "Big Science" Collaborations
- OSG offers a diversified portfolio of services to support these different science stakeholders.

Some Thanks are in Order

OSG has seen some significant transitions.

And I'd like to use the opportunity to thank the outgoing people and welcome the incoming people.

OSG "User Support"

Rob developed OSG-Connect to support individual researchers on OSG, and has lead OSG's engagement efforts with individual researchers and campuses for the last many years.

Many Thanks Rob !!!

He has handed responsibility for individual researchers over to Lauren Michaels to focus more on HL-LHC. His group maintains responsibility for the Multi-Institutional Science Collaborations.

Welcome Lauren !!!

OSG "Operations"

Rob Quick lead OSG Operations for many years.

Many Thanks Rob !!!

He has handed over responsibility for OSG Operations to Jeff Dost.

Welcome Jeff !!!