LIGO, Shared Computing, and OSG

aka “LIGO and OSG - Take II”

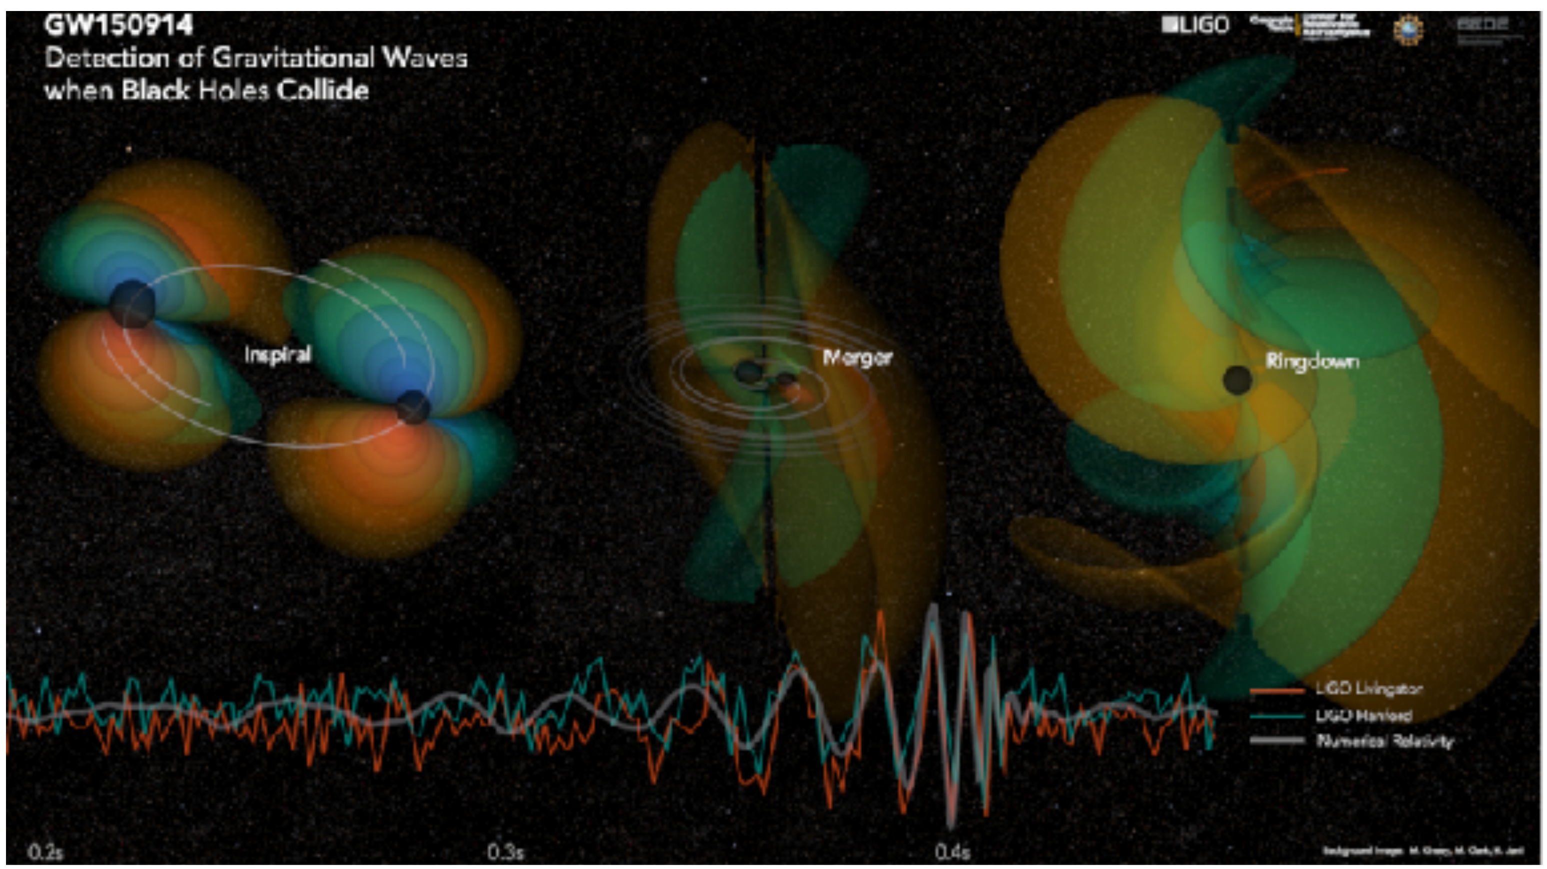

Peter Couvares, LIGO Laboratory — Caltech
LIGO-G1701927-v1

Who am I?

- Former Condor Team staff @ UW-Madison ('99-'08). Involved with early Particle Physics Data Grid (PPDG) and Grid Physics Network (GriPhyN) efforts back in the day.
- At Syracuse University ('09-'15), focused on distributed computing for the LIGO Scientific Collaboration and fostering a research computing community at SU. Helped establish the 10k+ core Orange Grid and Crush campus research clusters and agitated to join them to OSG.
- Now working for the LIGO Laboratory - Caltech ('15-), managing LIGO's data analysis computing with a focus on optimization (broadly defined) and shared computing.

Why am I Here?

- To petition for LIGO membership on the OSG Council.
- To engage more formally and regularly in discussions about the mission, goals, status, and future of the Open Science Grid.
- To have a voice in discussions in which we feel LIGO has a real and growing stake, and a useful perspective to offer.
- LIGO would like to (re)join the OSG Council in order to represent the interests and contribute the perspective of our 1k+ person scientific collaboration, and of the LIGO Laboratory, one of the largest NSF-funded research facilities.

No, Really, Why am I Here?

I'm missing a *really* good party to be with y'all.

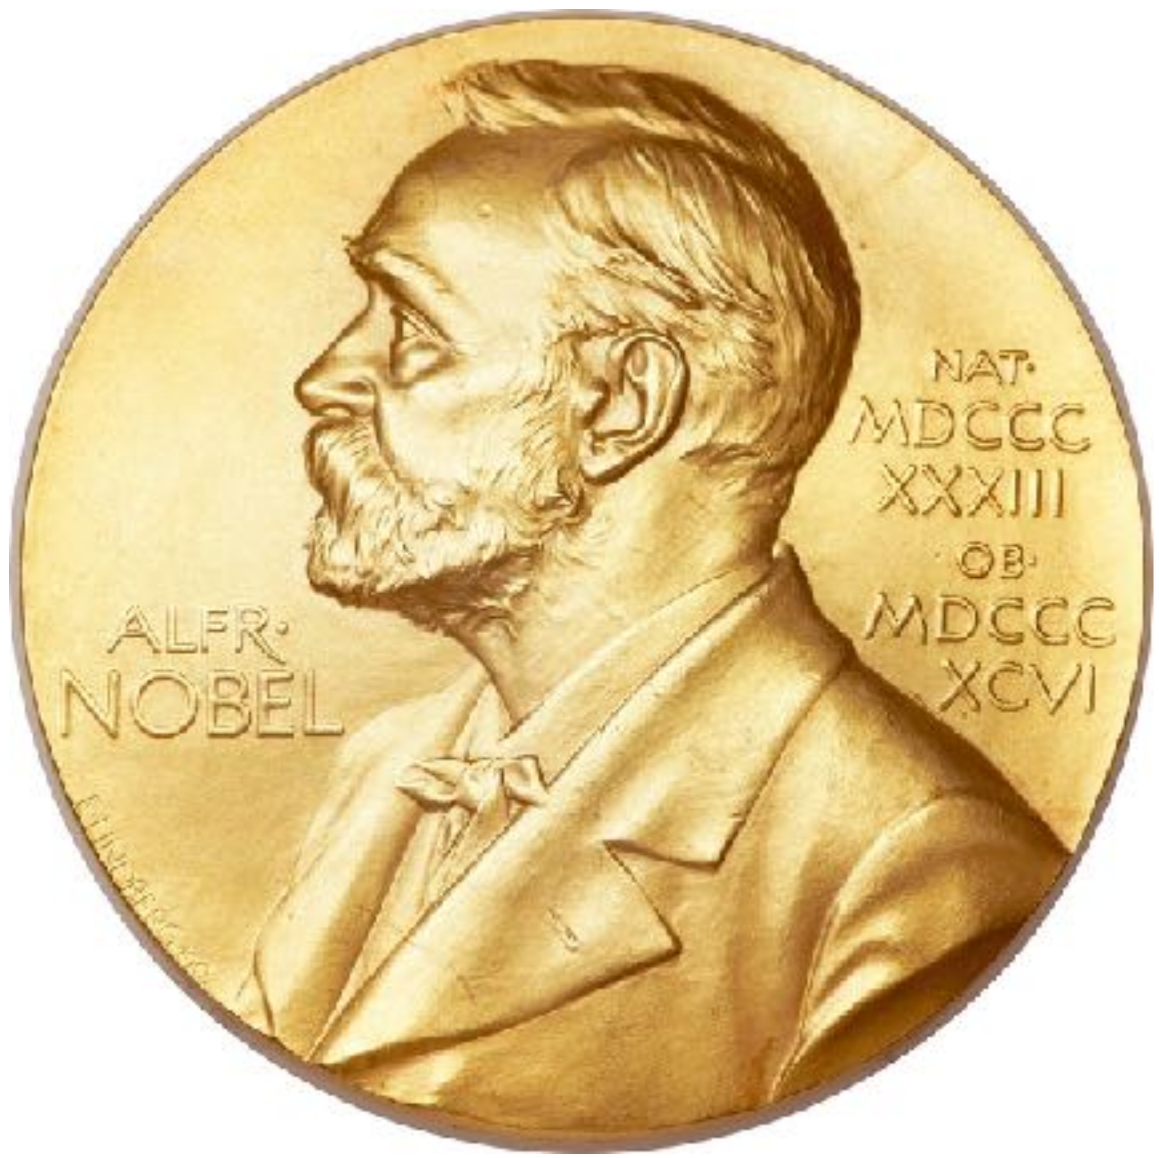

LIGO Science

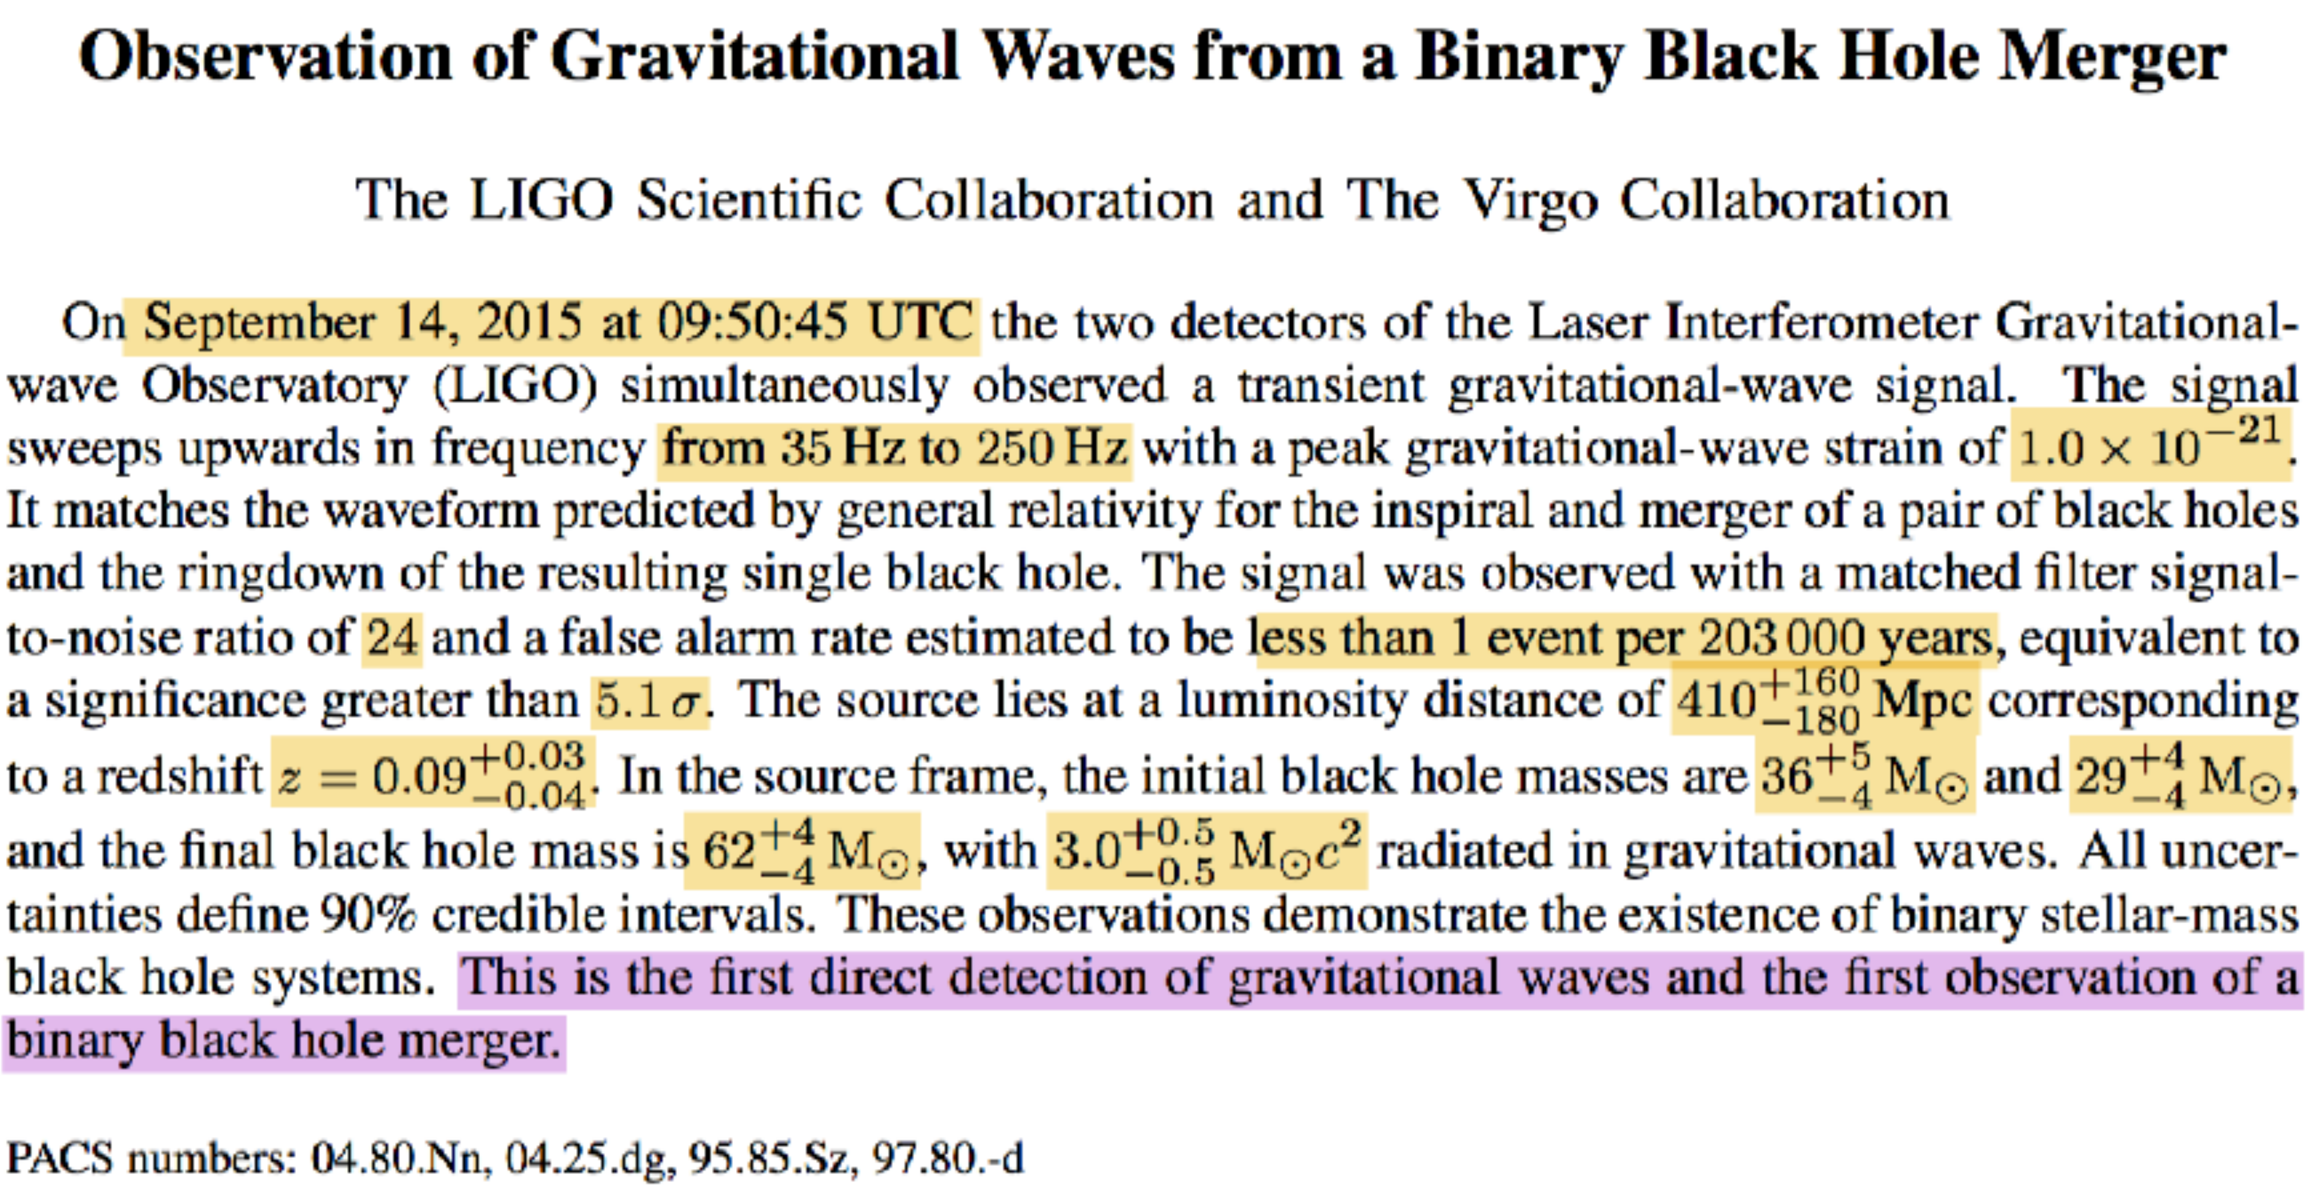
Observation of Gravitational Waves from a Binary Black Hole Merger

The LIGO Scientific Collaboration and The Virgo Collaboration

On September 14, 2015 at 09:50:45 UTC the two detectors of the Laser Interferometer Gravitational-wave Observatory (LIGO) simultaneously observed a transient gravitational-wave signal. The signal sweeps upwards in frequency from 35 Hz to 250 Hz with a peak gravitational-wave strain of 1.0×10^{-21}. It matches the waveform predicted by general relativity for the inspiral and merger of a pair of black holes and the ringdown of the resulting single black hole. The signal was observed with a matched filter signal-to-noise ratio of 24 and a false alarm rate estimated to be less than 1 event per 203 000 years, equivalent to a significance greater than $5.1\,\sigma$. The source lies at a luminosity distance of 410^{+160}_{-180} Mpc corresponding to a redshift $z = 0.09^{+0.03}_{-0.04}$. In the source frame, the initial black hole masses are $36^{+5}_{-4}\,M_{\odot}$ and $29^{+4}_{-4}\,M_{\odot}$, and the final black hole mass is $62^{+4}_{-4}\,M_{\odot}$, with $3.0^{+0.5}_{-0.5}\,M_{\odot}c^2$ radiated in gravitational waves. All uncertainties define 90% credible intervals. These observations demonstrate the existence of binary stellar-mass black hole systems. This is the first direct detection of gravitational waves and the first observation of a binary black hole merger.

PACS numbers: 04.80.Nn, 04.25.dg, 95.85.Sz, 97.80.-d

LIGO Data Analysis Computing

- The bulk of our searches are *embarrassingly parallel.*
- All LIGO analyses and computing resources are managed using HTCondor, which schedules work and handles faults to ensure reliable execution of embarrassingly parallel jobs.
  - Broad use of single tools develops a knowledge base in scientific user community
- LSC computing staff and HTCondor team have a very close 15+ year-old collaboration.
  - regular meetings between senior staff
  - bug fixes / feature development and feedback

Scale

- LIGO Data rates
  - Strain per IFO: ~3GB/yr
    - x N calibrations/reductions
- LIGO-Virgo Computing requirements
  - MSU=1 million E5-2670 core hours
  - O3 (~2019) Projected=~500 MSU
- Users on LIGO-Virgo Computing Network
  - ~600 users, ~300 active past year
  - Top 20 users drive 70% of demand

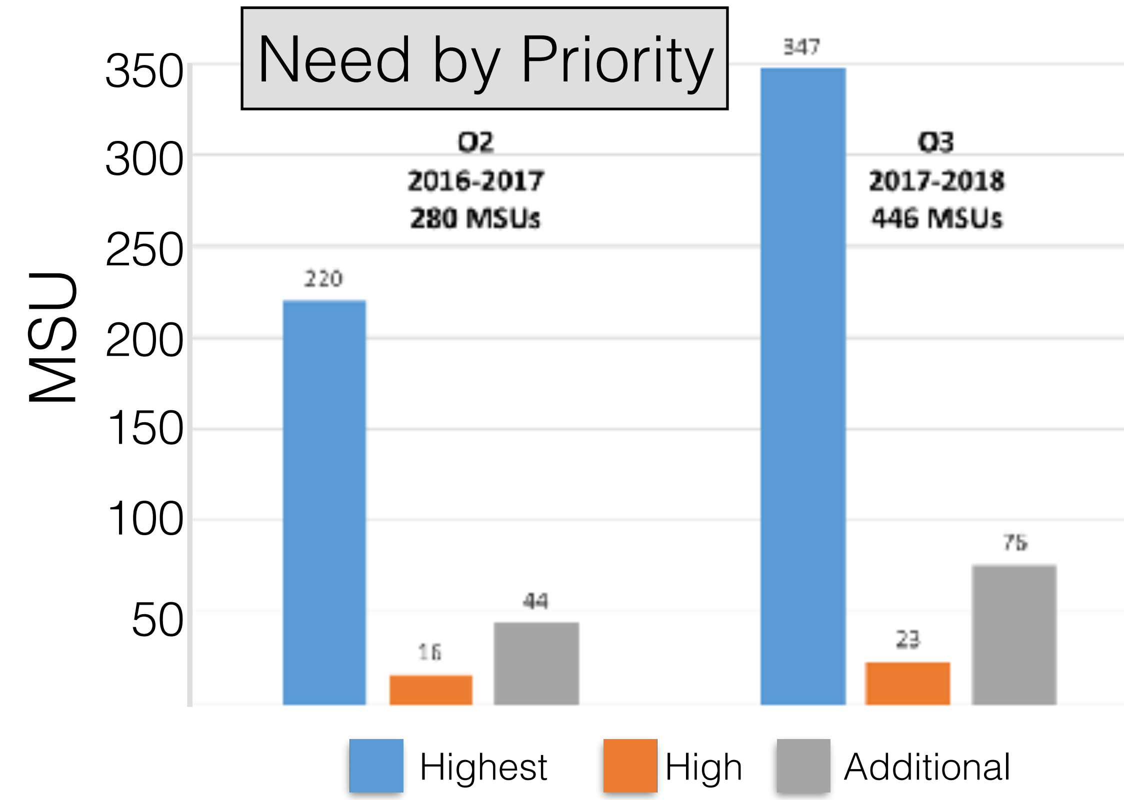

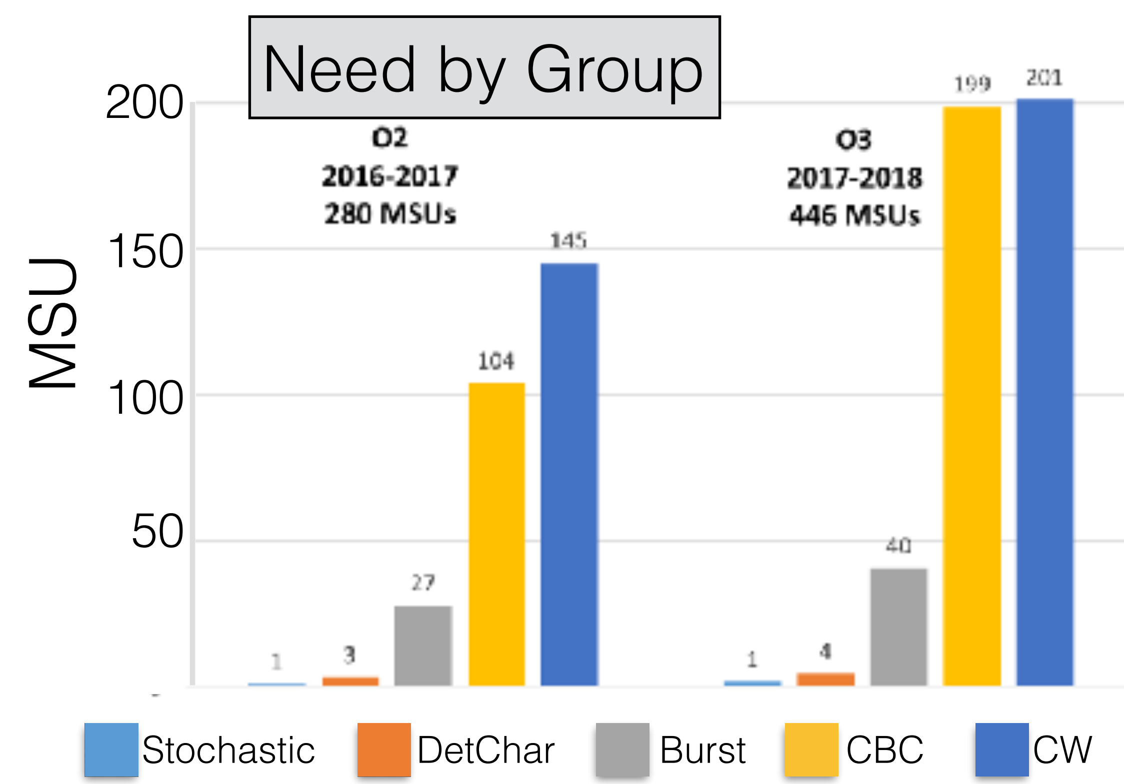

LIGO Long-Term Challenge

- Increasing heterogeneity, complexity of LIGO computing platforms:
  - of **processing hardware** (CPU generations, GPUs, MICs) — due to the opportunities for cost savings, we *must* support multiple generations of CPUs, GPUs, MIC platforms and treat them each as distinct platforms — lowest common denominator code not good enough
  - of **providers** — internal to LIGO, partners & collaborators, institutional, regional/national, commercial, volunteer
  - of provider operating systems and **software environments** — containerization, etc. are tools to mitigate but aren't a silver bullet
  - of provider **batch/queueing** systems
  - of provider **storage and network** interfaces and capabilities
  - of provider **policies** for identity+access management, workflow prioritization
  - of provider **accounting** models and accounting systems
  - of provider **motivations and expectations** — mutual scientific/strategic interest, public or scientific recognition, financial or other compensation, etc. — and not everything is in a MOU, SLA, or contract

LIGO Data Analysis Computing: Supply

- Many types of supply: **dedicated, allocated, opportunistic, volunteer**. Many providers in the US and abroad:
  - Dedicated LIGO Lab clusters (HTC)
  - Dedicated LSC clusters (HTC)

  LIGO Data Grid — ~80% in O1

  - Virgo clusters (mostly allocated on shared resources, HTC)
  - PI clusters (shared, HTC and HPC)
  - Campus/regional shared clusters (allocated, HTC and HPC) e.g., OrangeGrid, PACE, SciNet
  - National shared supercomputers (allocated, HTC and HPC) e.g., XSEDE, Blue Waters
  - Opportunistic cycles (campus clusters, DOE labs, HEP clusters, etc.)
  - future: commercial cloud (EC2, Azure, Google, Rackspace, etc.)?

  via Open Science Grid — ~20% in O1

- Two runtime software environments: LIGO Data Grid, Open Science Grid
- + Volunteer Einstein @ Home computing (~5 PetaFLOPS)

LIGO Use of OSG

- LIGO's OSG computing has contributed directly to results in our detection papers — e.g., false-alarm rates estimated via injection runs performed entirely on OSG.
- The LIGO analyses have run across scores of different OSG resources.
- >20M OSG CPU-hours and counting.
- ~5TB of input data stored at the Holland Computing Center (HCC) at the University of Nebraska-Lincoln.
- The total data volume distributed to jobs from Nebraska >1PB.
- Data rates from Nebraska storage to worker nodes ~10Gbps sustained. (Recently demonstrated >30Gpbs by accident!)

How OSG Helps LIGO

- A "universal adapter" to diverse resource types: dedicated LIGO CPUs, "friendly" campus clusters, "friendly" PI clusters, opportunistic OSG CPUs, XSEDE allocations, and in the future, possibly public and commercial cloud CPUs.

- Outsourced plumbing (factories, CEs, etc.) + expert help = easy to get started without making a huge labor investment.

- Track record of success — HEP forged a path.

- Friendly, enthusiastic, skilled, **results-oriented**, **flexible** OSG staff. Not hung up on boundaries, processes — focused on our goals.

Universal Adapter

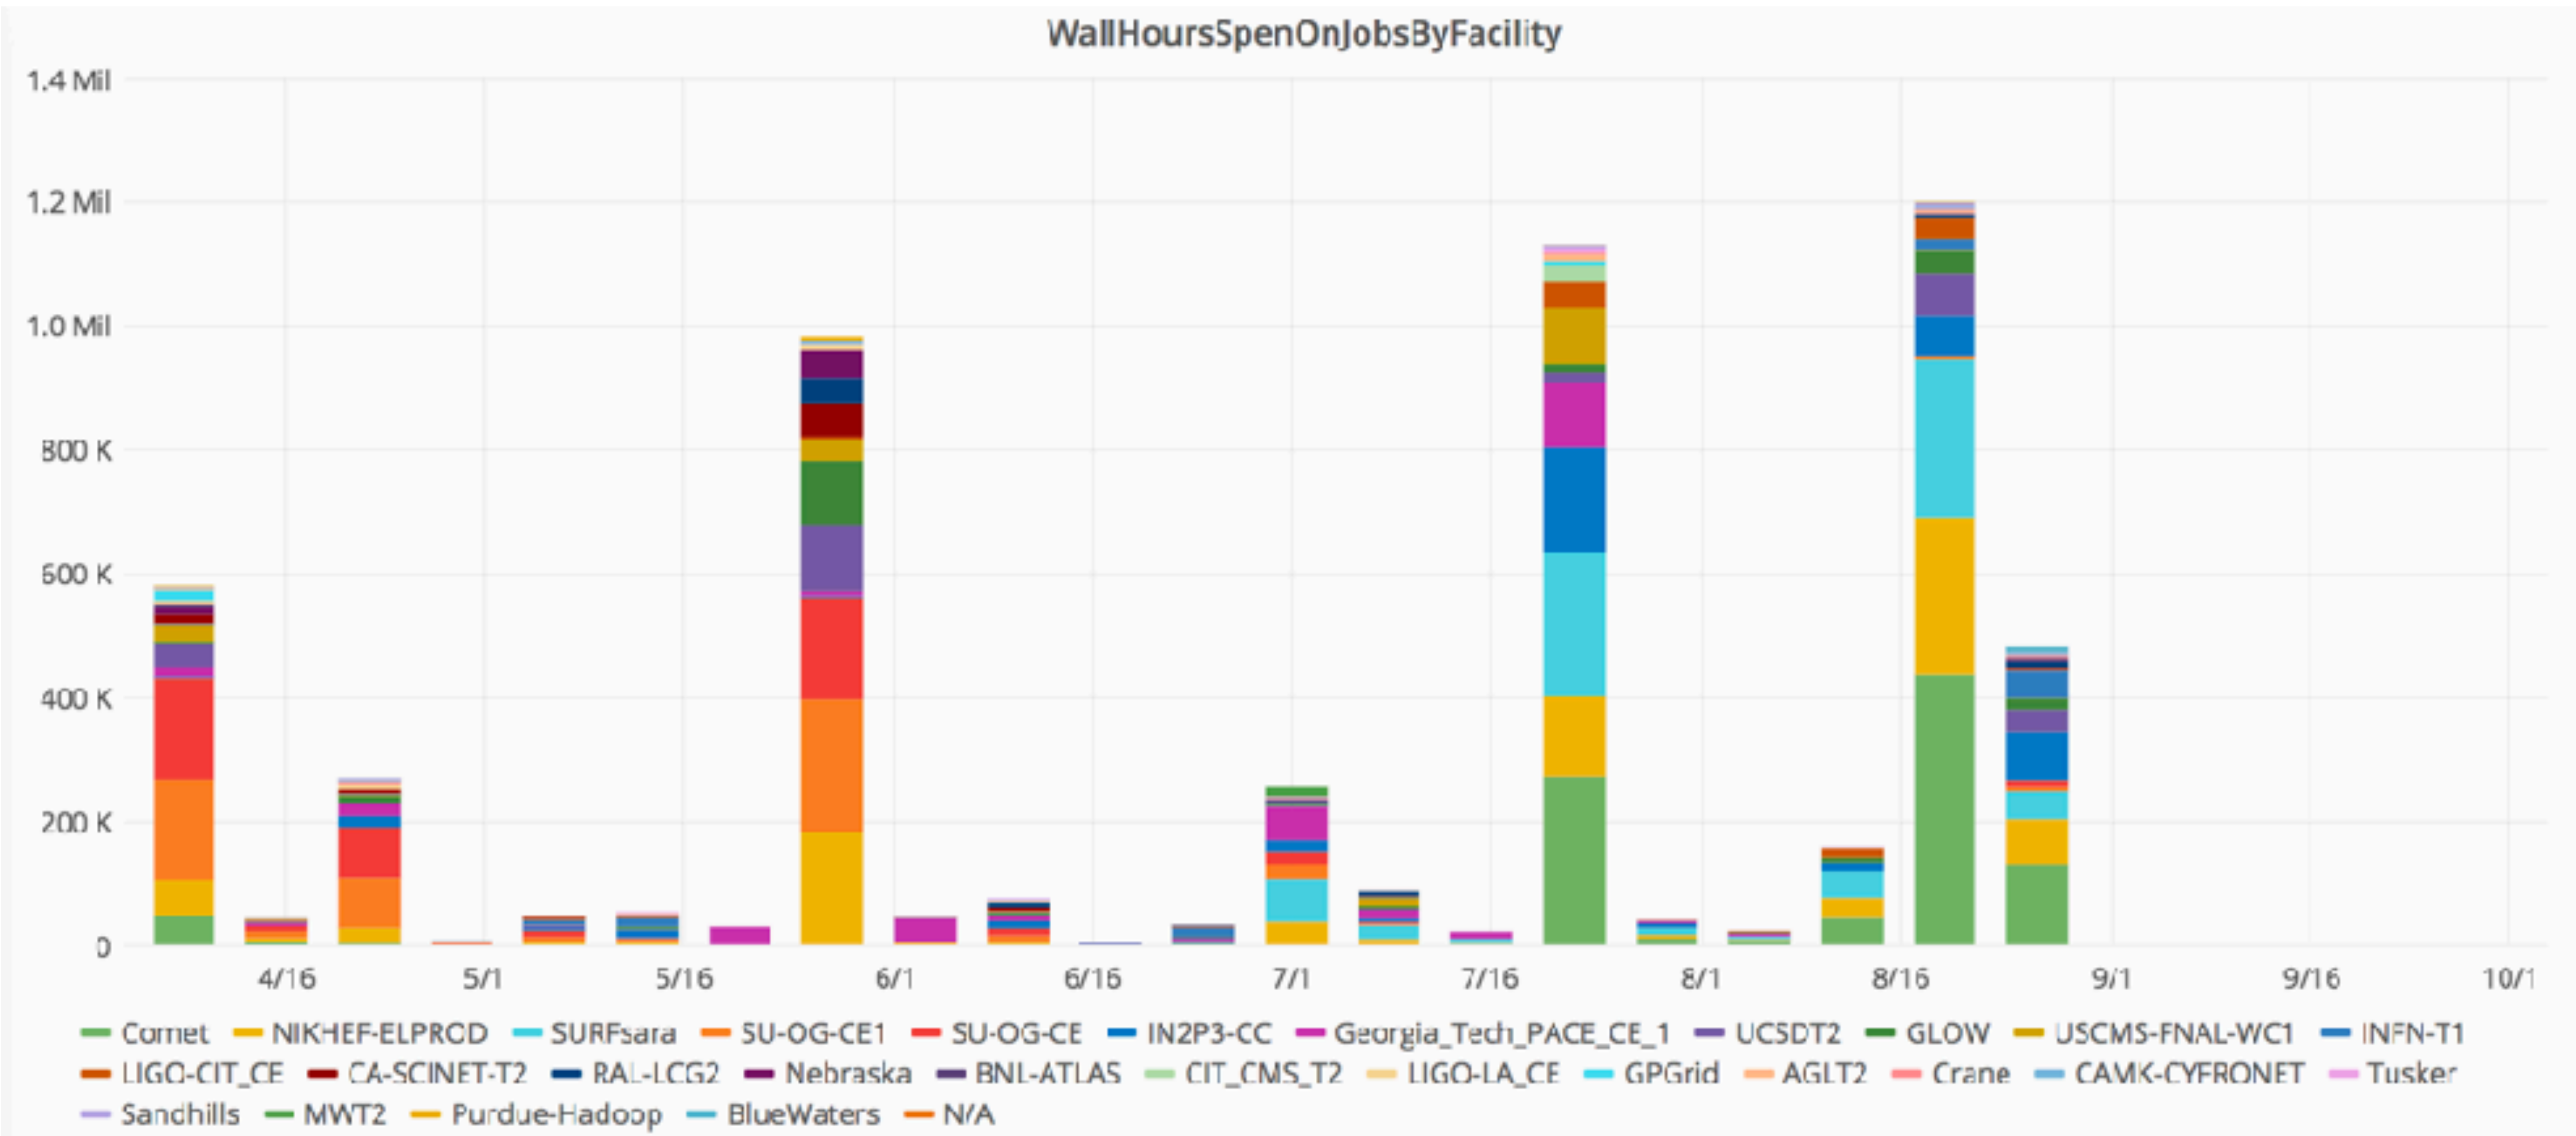

LIGO How OSG Helps LIGO

- LIGO sees OSG as a de-facto, distributed “Center of Excellence” for distributed scientific computing.

- OSG can generalize the experiences of many large and small scientific research projects, not just in terms of architecture but from the trenches.

- OSG staff work alongside many projects’ scientists and technical staff and can identify common problems and devise and disseminate common solutions and best practices.

LIGO Why LIGO is Good for OSG

- Broadening the community: LIGO is one of the largest non-HEP users of OSG.
- LIGO brings new human and computing resources and collaborators to OSG, and evangelizes OSG within the NSF and the larger scientific computing community.
- LIGO brings deep expertise in technologies relevant to OSG's mission: distributed identity and access management, distributed workflow management (e.g., with Pegasus), etc.
- LIGO brings a willingness to experiment and beta-test new OSG solutions (e.g., StashCache) internally and with partners.

Thank you!

LIGO-Virgo Optimization Approach: "The Whole Patient"

Iterate

- Scientific Prioritization and Scoping
- Estimation and Benchmarking of Computational Costs
- Optimization of Data Analysis Methods and Algorithms
- Optimization of Code Implementation and Libraries
- Compiler Optimizations
- Workflow Management Optimizations
- Development, Testing, and Simulation Process Optimizations
- LIGO-Virgo Computing Network Scheduling Optimizations
- Resource Supply Optimizations (make more cycles available)
- Workflow Portability Optimizations (expand usable resources)
- Hardware Procurement
- Pipeline Reviews including Computational Efficiency
- Documentation, Training, Collaboration and External Engagement

OSG Can Help

Neglect nothing, focus on "bang for the buck" and where optimization effort can be most effective. Avoid adding burden where the payback is small.

OSG Challenges for LIGO

- Search pipelines must be “ported” from rich LIGO Data Grid runtime environment to austere OSG environment
  - understand software dependencies and then either trim them down, package them along with payload, deploy them into CVMFS, and/or containerize the runtime environment.
  - new checkpointing challenges (Condor stduniv -> application checkpointing)
- Data access questions: WAN xfer, pre-staged repos, and/or CVMFS/StashCache. Ease of deployment vs. scalability and predictability/robustness.
- Accounting: manual aggregation of two sources of data, with different units and metadata (CPU core hours vs SUs, users, pipelines, etc.), lots of discrepancies to understand.
- May reduce systems administration burden in some ways (we’re using clusters we didn’t have the person-power to integrate into the LDG), but we are now understaffed for “grid admin” — how much more can/should OSG help?
- Complicates our computing model — funding implications — need to be clear that elasticity does not eliminate need for in-house computing to meet steady-state demand, or provide low-latency (quasi-“real-time”) computing.