IGWN Distributed Computing

WLCG Grid Deployment Board
January 2023

Peter Couvares, LIGO Laboratory

LIGO-G2300046-v2

Who am I?

- Peter Couvares — Staff Scientist in the LIGO Laboratory at Caltech
- LIGO Lab's Data Analysis Computing Manager — responsible for computing optimization (broadly defined) and success of distributed high-throughput computing (DHTC) for IGWN.
- Previously IGWN Computing & Software Committee Co-Chair 2015-2022.
- Before LIGO/IGWN, worked in industry (dataviz and telecom) and was part of the early HTCondor team at UW-Madison ('99-'09) and authored of the Distributed Acyclic Graph Manager (DAGMan) tool.

IGWN: International Gravitational-Wave Observatory Network

- Not a formal organization (yet?), but an umbrella designation to represent the common and coordinated data analysis **computing** infrastructure and planning of LIGO, Virgo, KAGRA, and LIGO India.
- Of practical benefit (igwn.org services don't highlight just one collaboration and don't all need to change names if/when we add another detector, unlike ligo.org or LVK)
- Important symbolically to emphasize our increasingly shared computing infrastructure and effort, and a potential precursor to a future more formal umbrella organization.

Why am I here?

- Because you asked!
- To share some of IGWN's operational issues and ask questions — WLCG and IGWN face a variety of common research computing challenges:
  - Distributed computing challenges of large scientific collaborations.
  - Coordination of overlapping computing providers and grids. Scientific workflow monitoring & operations.
  - The evolving intersection of High Performance Computing (HPC) and High-Throughput Computing (HTC).
- To share, learn from you, and find collaborators for future work…

Background: LVK Computing Scope Growth

- **SCALE** 🚀 - With ~1x the computing staff that we had in ~2009, we're supporting a ~5x larger user-base, delivering ~100x more computing power per observing run (10^6→10^9 CPU core hours) from ~10x the number of providers, while doing more science per computing cycle due to systematic optimization.
- **LOW LATENCY** - We support a complex, mission-critical, public-facing low-latency detection and alerting infrastructure on top of our previous almost-entirely batch/offline data analysis capabilities.
- **COLLABORATION** - We support seamless single-sign-on across almost every LVK service (numbering in the 100's) while providing many of the traditional productivity services of a large corporation (email/wiki/chat/document management/user database/etc.).
- **GITLAB** - We support a production git.ligo.org service to power everything from LVK software development, to automated testing, paper-writing, project management, and more.
- **IGWN** - We support the **unified** data analysis computing operations of the member collaborations, having replaced most parallel, duplicative, and incompatible infrastructure services with common IGWN solutions.
- **SHARED COMPUTING** - We transitioned from a labor-intensive "walled garden" of largely internal computing clusters to an open grid of shared resources into which we strive to efficiently incorporate CPUs or GPUs anywhere in the world without requiring data analysts to adapt to the peculiarities of individual providers.

Computing & Software

- We want to **rely on external providers** when it make sense. E.g., for identity and access management (CoManage, InCommon, etc.), distributed computing infrastructure (HTCondor, OSG), etc. — experts who can do things better than we can in-house, and on a larger scale.
- We need to stop making things more difficult than they need to be. **Less customization, fewer proprietary requirements**, etc. There is a large hidden computing cost for anything non-standard. The more we can adapt to match what HEP and others do, the better. (And better yet if we plan together in advance…)
- IGWN has seen enormous benefit in relying on **trusted liaisons** and intermediaries (e.g., OSG) to interface with disparate service providers — especially when share those providers in common with other large projects.
- We'd like to **avoid duplication of effort** and leverage solutions others have developed.

One Success Story: OSG

- LIGO adopted a High-Throughput Computing model over 15 years ago, and the HTCondor ecosystem (HTCondor, DAGMan, etc.) is in the core of IGWN's computing DNA. IGWN's relationship with OSG is anchored on this pre-existing culture, computing model, and software stack.

- OSG provides a "**universal adapter"** to diverse resource types: dedicated internal resources, major shared computing centers (e.g., CNAF, IN2P3, Nikhef), large campus clusters, opportunistic CPUs, supercomputer allocations (e.g., XSEDE/ ACCESS, Compute Canada, etc.) — and in the future commercial cloud CPUs.

- OSG **hosted services** (GlideinWMS, hosted CEs, etc.) + expert support = production infrastructure without a huge IGWN-specific investment.

- **Track record of success** in HEP — High Energy Physics (LHC) computing forged a path, thanks to the work of many of the people in this room.

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

Questions

- IGWN is taking more responsibility for troubleshooting problems for our users.
  - Jobs match and restart for no discernible reason. Jobs go on hold with opaque errors. Jobs appear to be running but make no progress.
- We're increasingly struggling with CE/provider monitoring and operations
  - Knowing when sites are offline or otherwise unable to run jobs successfully, and responding once we know. (Pilot failures are hard to debug — key information doesn't make it back to the Access Point.)
  - Validating data access (e.g., authenticated CVMFS) before a job runs.
- **Can/should we make better use of work that's already been done to monitor WLCG sites, validate resources immediately before jobs run, and/or respond to problems?**

Data and OSDF

- IGWN is making increasingly heavy use of the worldwide Open Science Data Federation (OSDF) network, including ESnet caches at London & Amsterdam POPs — via CVMFS, HTCondor (stash:// file xfer), and the native stash API (stashcp). **How does OSDF relate to WLCG data infrastructure efforts**, and what else should IGWN be using, coordinating, or integrating with?

N-way Communication

- There are so many organizations to coordinate with, and I have to confess I still don't understand all their roles and relationships…
- How should IGWN best interface with multiple worldwide distributed computing organizations and infrastructures?
- What's is the connection (if any) between LIGO being a "CERN Recognized Experiment" (since 2015) and Virgo becoming a WLCG partner (recently)? How do — or *should* — all these things relate?
- **In an ideal world, I'd like to rely on OSG as our first-point-of-contact and liaison** to WLCG, EGI, NeIC, and others so IGWN doesn't have to coordinate everything directly — is this a good idea? How would each of *you* prefer to interact with IGWN?

Usage Accounting

- Centrally reporting — accurately and reliably — how many resources across the world are consumed for different IGWN science goals is a chronic albatross: it's labor-intensive and difficult to get right.
- Is this an area where we should let go of our own requirements and try to adapt to use what others do? Does WLCG have a *labor-efficient* accounting system that works and produces all the information HEP experiments need to report?

Token Auth

- IGWN is “all-in” on the transition from x509 certificates (identity/role-mapping auth) to token (capability-based) auth.
- IGWN is an early adopter with SciTokens and actively working on adapting access to data services accordingly.
- **What is the status of WLCG efforts to integrate SciTokens and WLCG tokens?**

GPUs

- IGWN has a number of data analysis pipelines that run on NVIDIA (CUDA) GPUs.
- We'd like to scale up our use of GPUs but our demand is very bursty — **how can/should we make best use of available European or WLCG GPUs**?
- What's the WLCG vision for provisioning/scheduling/ accounting GPUs?

Questions?

- Questions for me?
- <pcouvare@caltech.edu> or <peter.couvares@ligo.org>
- Let’s talk!

Extra Slides

Long-Term Challenge

- Increasing heterogeneity, complexity of IGWN 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 (LIGO + Virgo + KAGRA + LIGO India), 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

What IGWN Bring to WLCG

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