



NIKHEF SITE UPDATE

NOVEMBER 2019

Mary Hester
Virgo CC Workshop
29 November 2019



OUTLINE

- Overview of GW support at Nikhef and in NL
- Virgo/Ligo Computing at Nikhef
- New dCache storage at Nikhef
- New developments with Condor/OSG
- Conclusions and Questions


INFRASTRUCTURE SUPPORT @ NIKHEF

- PDP operations group: 3 people
- Distributed infrastructure operations with SURFsara
- Resources reserved for local users (including local GW group):
 - Computing: 832 cores
 - Storage: 1.5 PB (dCache)
- Grid resources (IGWN users):

Resource Type	Nikhef	SURFsara
Computing	~7.2k cores	~5-10k cores?
Storage	4 PB	10 PB (Disk)? 41 PB (Tape)?



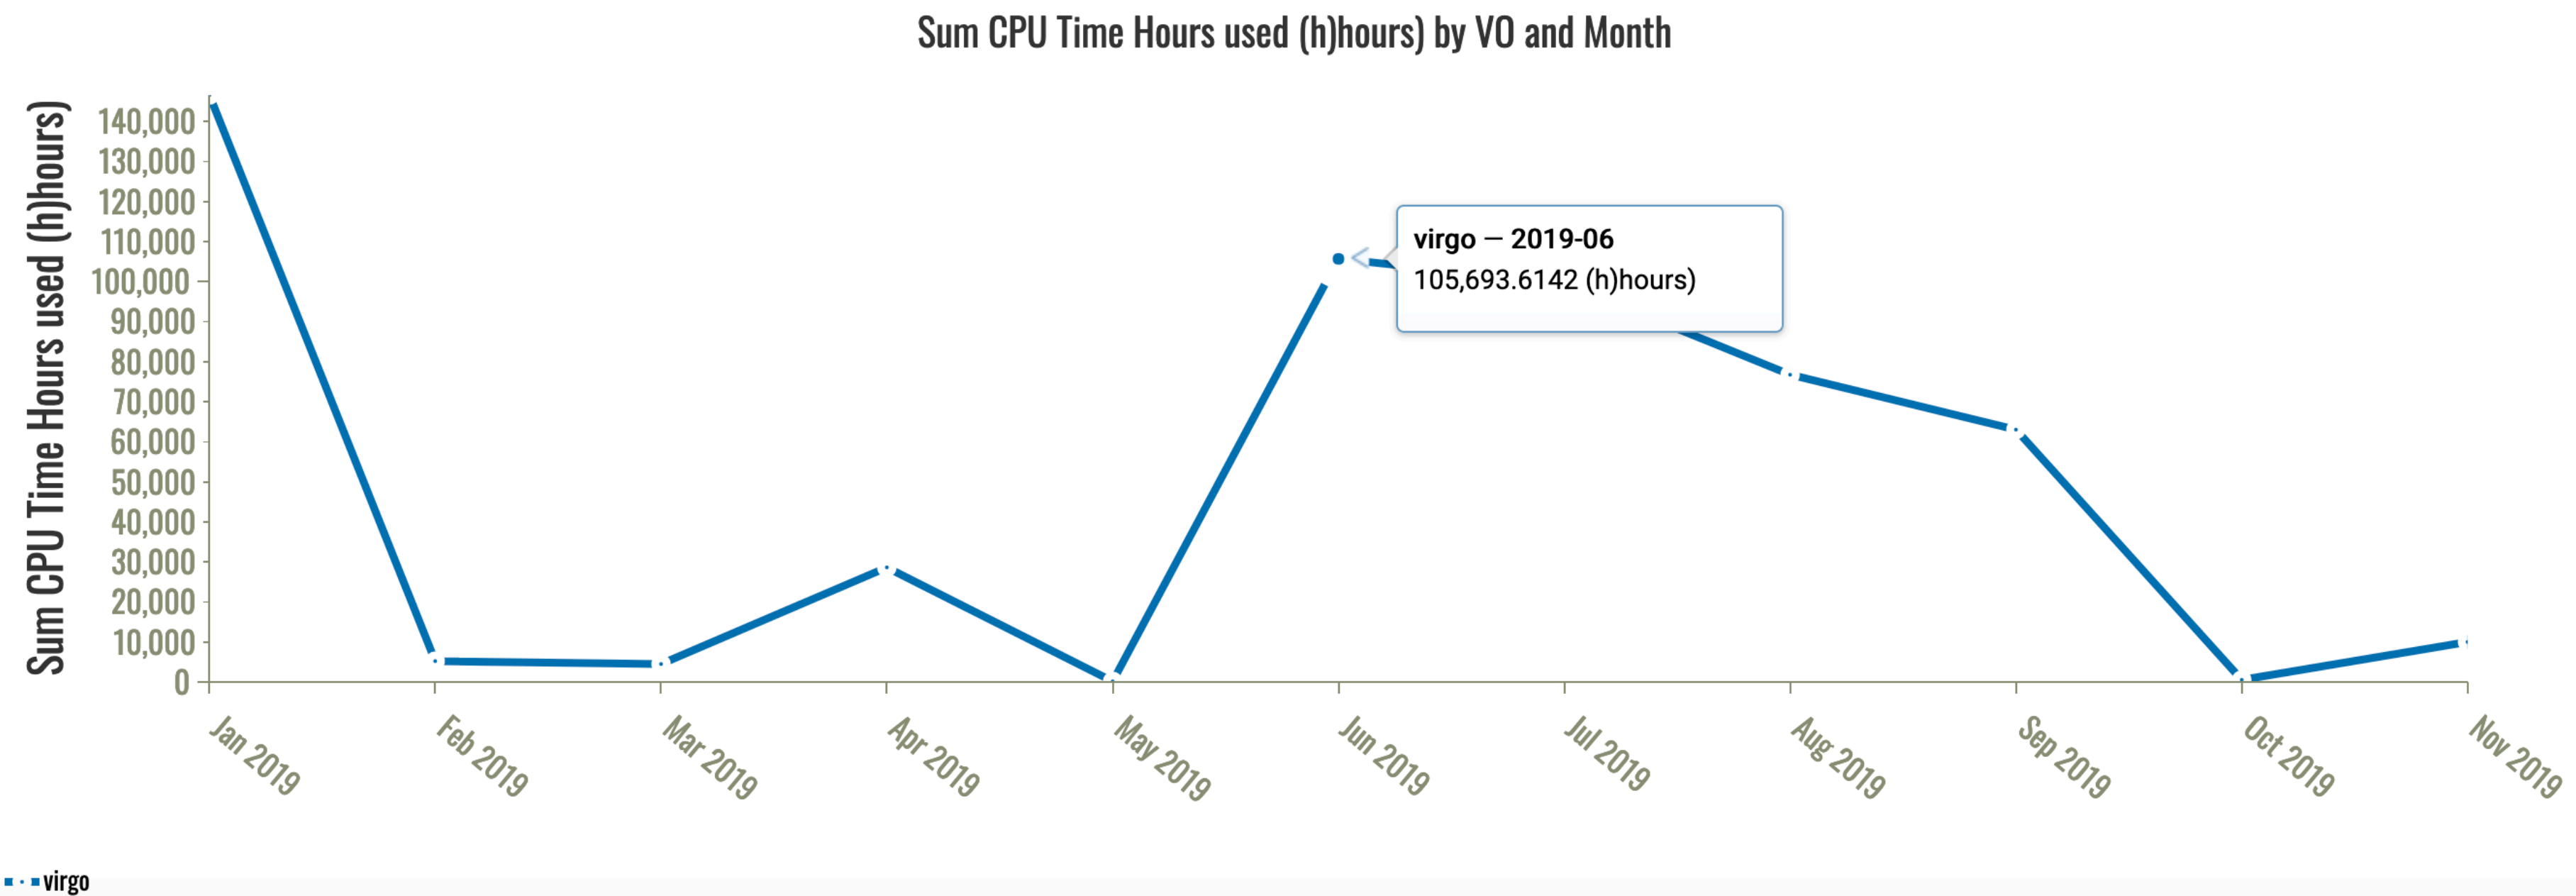
VIRGO & LIGO COMPUTING IN A YEAR



- Virgo computing use (light blue)
- Ligo computing use (dark blue)

COMPUTE RESOURCE ALLOCATION

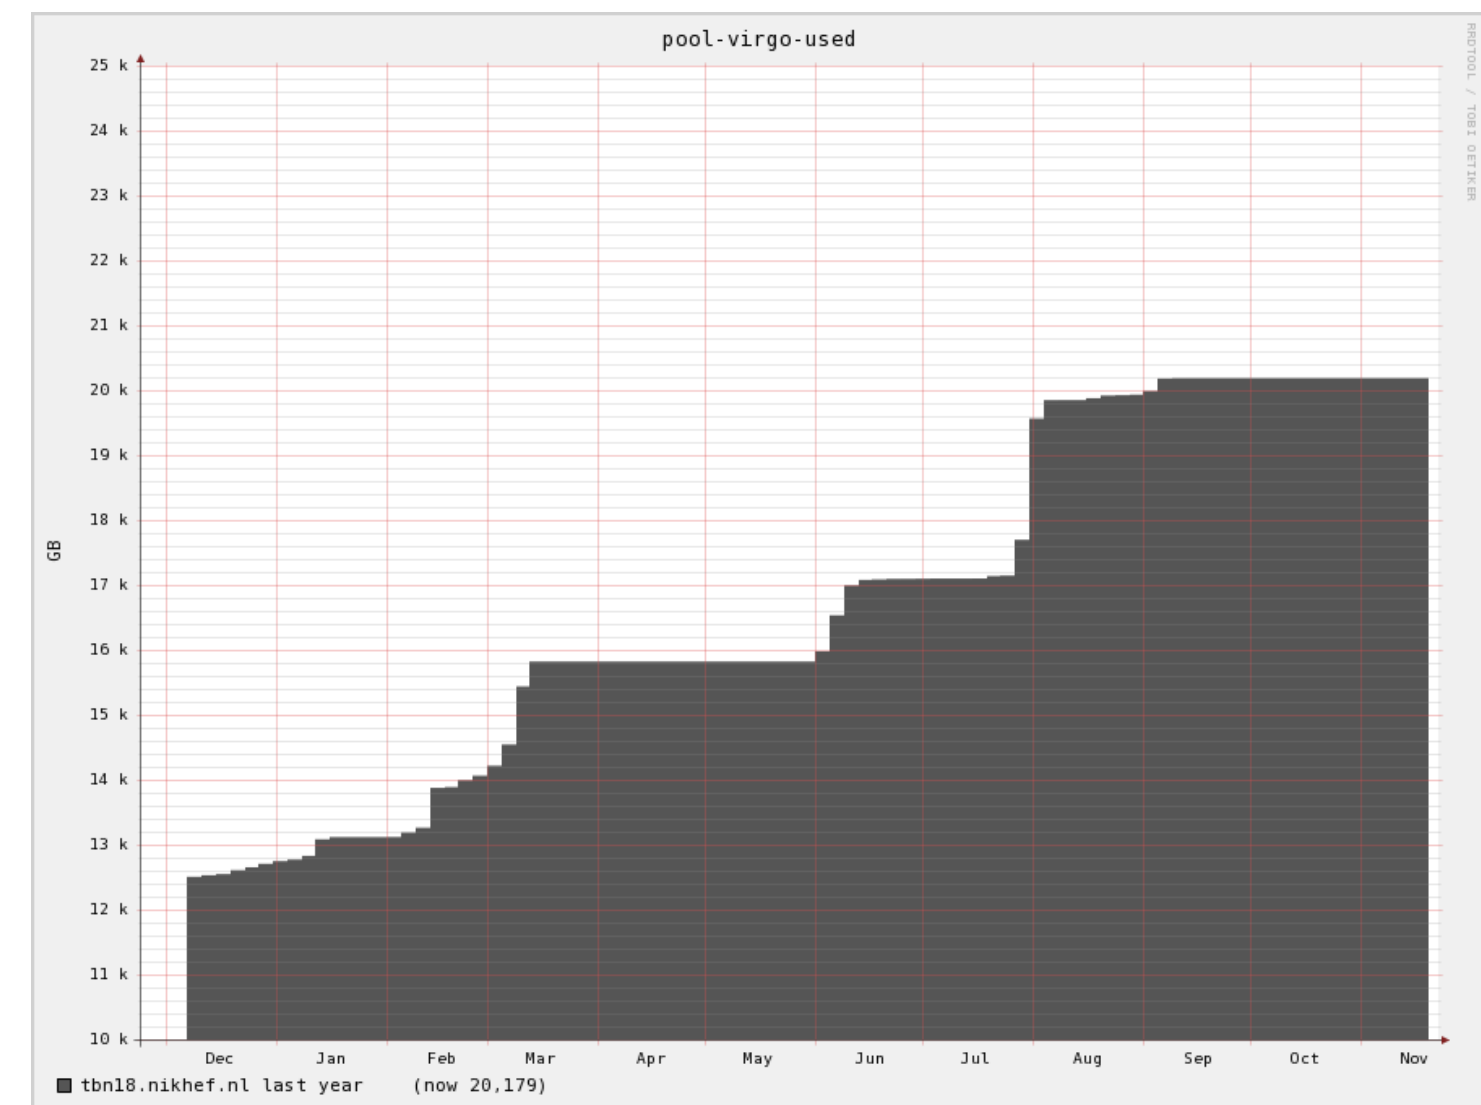

- Virgo allocations @ Nikhef - Plenty of resources for more computing!
 - Max processors available at any point in time @ Nikhef:
 - ~800 cores - available 24/7 (x2 with SURFsara)
 - Peak computing in June: 105,694 CPU hours
 - 146 cores / 800 cores \approx 18% of the possibilities!



VIRGO STORAGE USE


- 29.30 TB allocated on Grid (DPM) storage
 - Using: 18.34 TB
 - Free: 10.96 TB (37%)
- Decommissioning tbn18.nikhef.nl

 - 20 December: read-only access



MEET THE HOOIKANONS

- New Grid storage
- 3 New storage arrays + servers
- Add >1.5 PB to total grid storage resource
 - 556 TiB storage/server
- Server Specs:
 - IBM POWER9 LC922
 - Mellanox ConnectX-4, 100GbE single-port QSFP28 PCIe3.0 x16 LP
 - 2-Socket 2U
 - 2 x 2 TB HDDs
 - Broadcom (LSI) 9305-16E SAS-3 HBA PCIe3.0 x8 LP
 - 16-core 2.9 GHz POWER9 Processor
- Storage Array Specs:
 - NetApp E5700A
 - Controller: 2x 16Gb/10GbE optical onboard ports.
 - HIC, E5700/EF570, 12Gb SAS, 4-ports
 - 4U, 60 disk drives, 12TB, 12Gb



<http://www.hooiberg.info/?page=themas/laden>

MEET THE HOOIKANONS

- New Grid storage
- 3 New storage arrays + servers
- Add >1.5 PB to total grid storage resource
 - 556 TiB storage/server
- Server Specs:
 - IBM POWER9 LC922
 - Mellanox ConnectX-4, 100GbE single-port QSFP28 PCIe3.0 x16 LP
 - 2-Socket 2U
 - 2 x 2 TB HDDs
 - Broadcom (LSI) 9305-16E SAS-3 HBA PCIe3.0 x8 LP
 - 16-core 2.9 GHz POWER9 Processor
- Storage Array Specs:
 - NetApp E5700A
 - Controller: 2x 16Gb/10GbE optical onboard ports.
 - HIC, E5700/EF570, 12Gb SAS, 4-ports
 - 4U, 60 disk drives, 12TB, 12Gb



DCACHE



- New dCache allocations on hooikanon storage:
 - 28 TB pool partitions
- Access the storage resources through “doors”:
 - i.e., virgo.dcache.nikhef.nl
- These will be accessible by xrootd, webDAV, and gridFTP (no SRM support)
- New directory path for the storage is:
 - [/pnfs/nikhef.nl/data/virgo/pegasusdisk/](https://pnfs.nikhef.nl/data/virgo/pegasusdisk/)
- Admin user currently: virgo000 (from VOMs mapping)
 - read/write access to anyone in the VO

MIGRATION: DPM 2 DCACHE

- Decommissioning DPM node - tbn18.nikhef.nl
 - 20 December 2019: read-only access
 - Asking VOs to start moving data now
- How does Virgo want to do this migration?
- Who would like to work on this?
- For additional questions and help, please email:
grid.sysadmin@nikhef.nl

NIKHEF IGWN SUBMIT NODE

- This is NOT a traditional:
 - CE
 - UI
- So what is it?
 - An HTCondor submit host
 - Forwards / “flocks” jobs to an OSG-managed submit host
 - osg-ligo-1.t2.ucsd.edu
 - Rescheduled on OSG resources (at the moment) via a glidein-wms
 - Provides an easy way to access more Grid resources via OSG
- Nikhef flock node: stro.nikhef.nl
- More info at <https://opensciencegrid.org/docs/submit/osg-flock/>



Open Science Grid

IGWN Grid Architecture

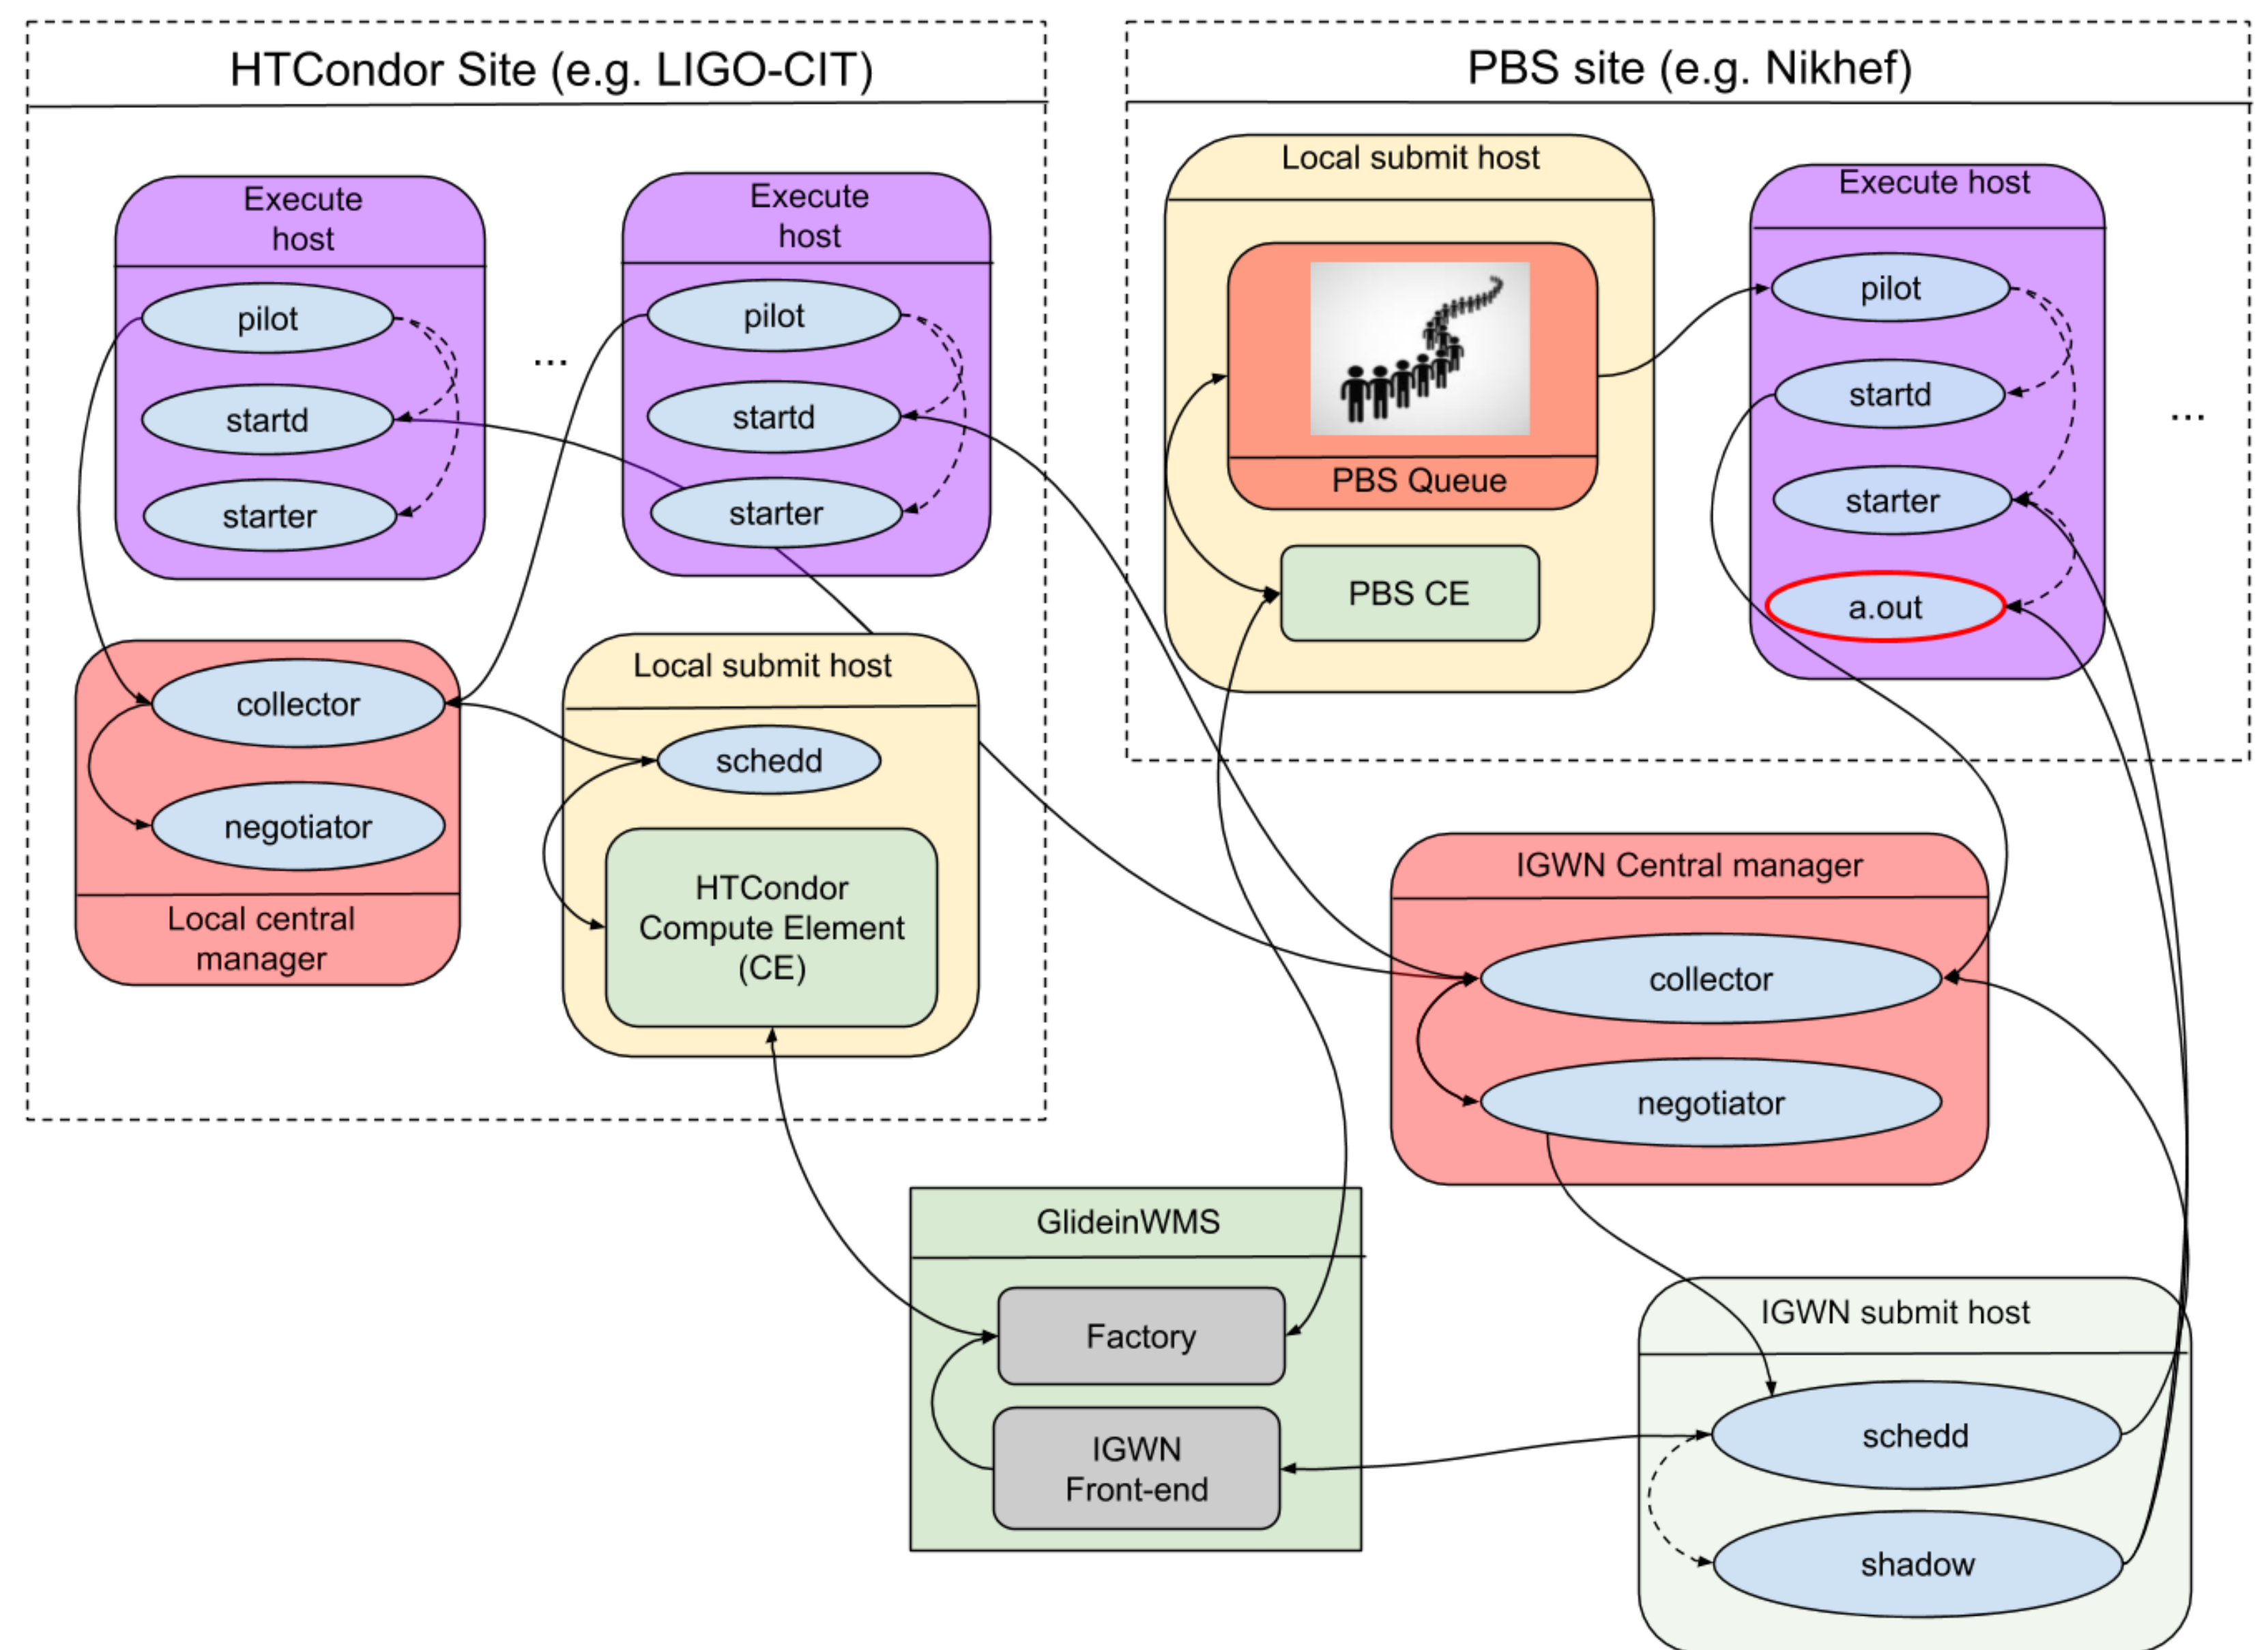
IGWN-grid comprising 1 HTCondor site & 1 PBS

1 job **a.out** running @ Nikhef

Compute element: “gateway” at remote sites, installed on local submit hosts

GlideinWMS provides *factory* and *frontend*, operated by OSG consortium (under k8s on Nautilus)

Factory submits pilot jobs to local batch system CE, spawns HTCondor daemons → slots in the user pool!



From James Clark's slides: <https://docs.google.com/presentation/d/11ipja4CeQValf6BCc1g8vS8udrsRb8kIQHWDcInQB5U/edit#slide=id.p>

SOME QUESTIONS ON THE TABLE...

- How does the Virgo collaboration want to make use of these flock nodes like the Nikhef IGWN submit node?
- What is the model/workflow that Virgo wants to develop in Europe?
- How do we make this easy/accessible to users?
- Is there a centralised way to make this accessible to both Virgo and Ligo users alike?

- Thoughts:

- Normal model for access to Nikhef Grid resources:
 - Users have certificates -> VO can sign proxies -> User belongs to VO
 - Site accepts VO -> users with valid vo-signed proxies are allowed
 - Is there/could there be a corresponding IGWN model?



QUESTIONS?
COMMENTS?
FEEDBACK?

THANK YOU!

GRID.SYSADMIN@NIKHEF.NL

