



IGWN SUBMIT NODE: DESIGNING EU-BASED
GATEWAY NODES FOR DISTRIBUTED COMPUTING
FOR VIRGO, LIGO, KAGRA

OCTOBER 2020

Mary Hester
HEPIX 2020 Online Workshop
12 October 2020

OUTLINE

- Virgo/LIGO Computing at Nikhef
- What is IGWN?
- Developing an IGWN submit node
- Conclusions & Future work
- Questions

INFRASTRUCTURE SUPPORT @ NIKHEF

- Part of the Dutch National e-Infrastructure (with SURFsara and others)
 - Nikhef PDP operations group: 3 people
- Resources reserved for local users (including local GW group):
 - Computing: 832 cores
 - Storage: 1.5 PB (dCache)
- Grid resources (LHC, IGWN, WeNMR, DUNE, Xenon, ...):

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



IGWN

The International Gravitational-Wave Observatory Network (IGWN) community includes members of

- KAGRA
- LIGO Laboratory
- LIGO Scientific Collaboration (LSC)
- Virgo Collaboration

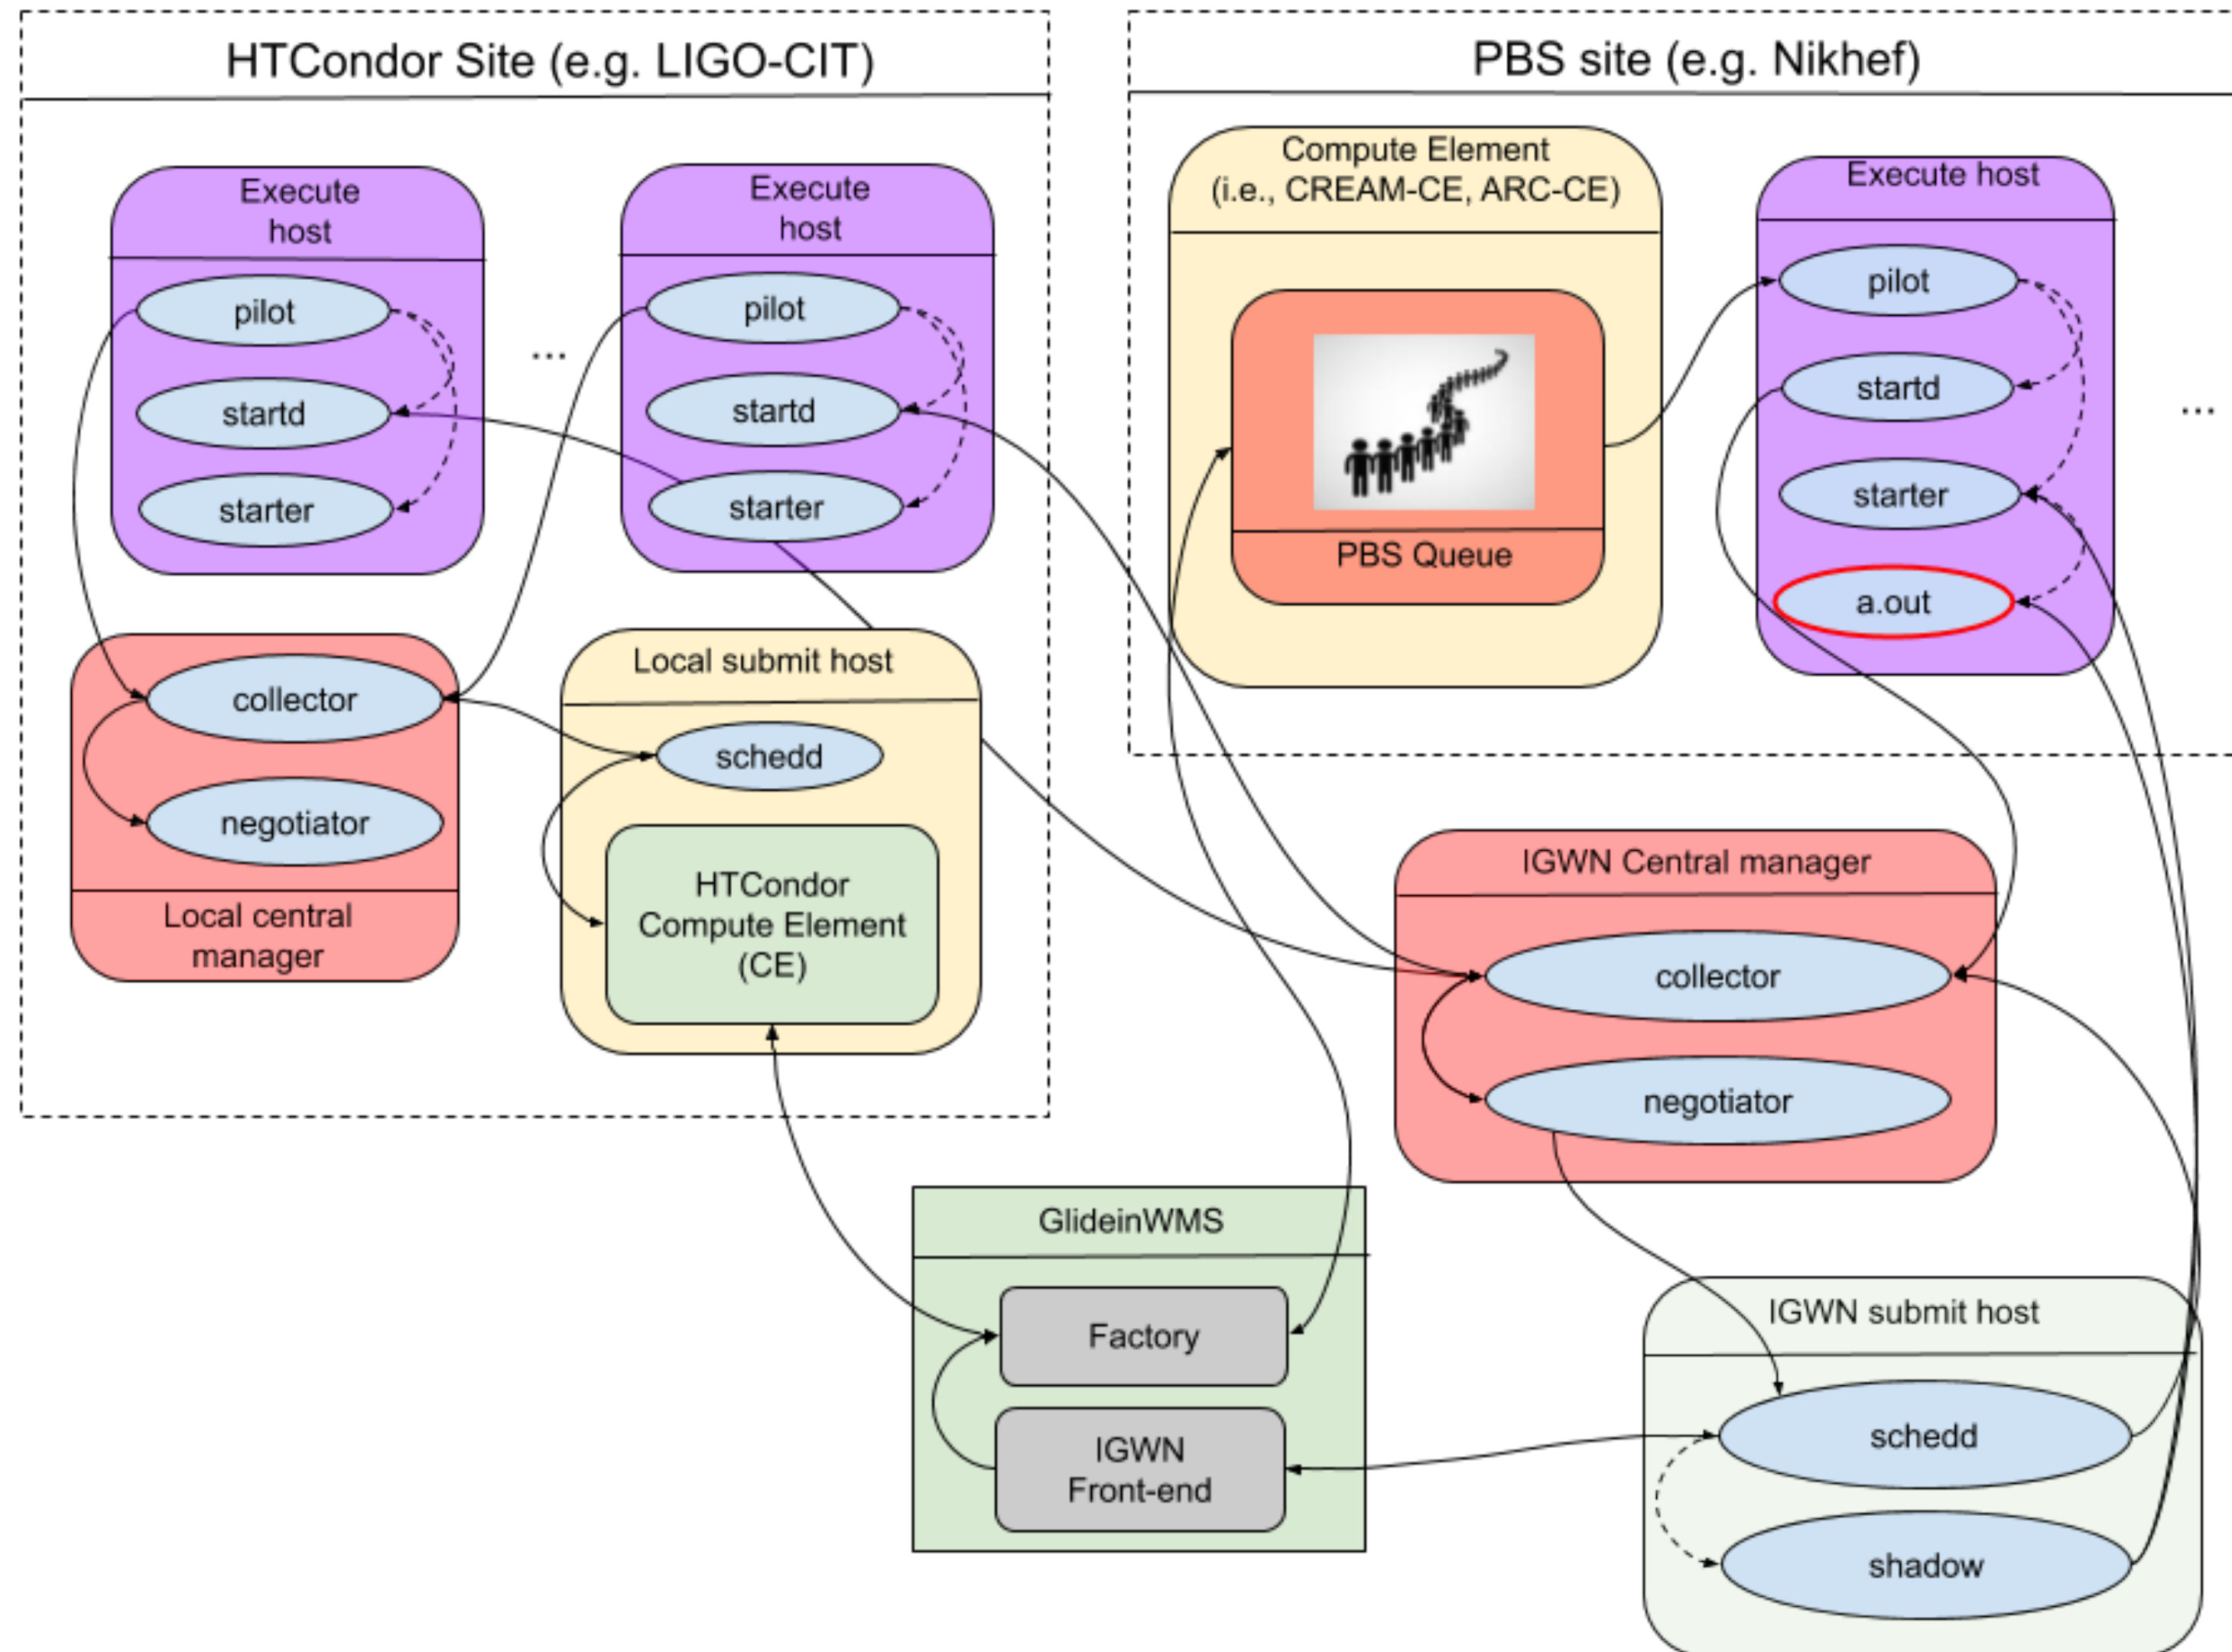
PURPOSE OF AN IGWN SUBMIT NODE

Enable GW pipelines to submit jobs to distributed computing resources globally — easily

Production submit nodes (or these gateway nodes) only in US

- 3-4 available for users
- Not ideal for user support (i.e., timezone offsets, etc)

DISTRIBUTED COMPUTING W/ IGWN



IGWN SUBMIT NODE

So what is it?

- Built on HTCondor flock node Requirements
- Forwards / “flocks” jobs to an OSG-managed submit host
 - In this case: osg-ligo-1.t2.ucsd.edu
 - Rescheduled on OSG resources via glideinwms

Easy to configure:

1. RHEL/centos7 machine
2. At least 1.2TB provisioned for the submit machine
3. Port 9618 open for condor connections
4. IPv4/v6 address - dual stack desirable
5. A host certificate from an IGTF-accredited CA
6. Repos — mixture of umd, osg, epel

More at <https://opensciencegrid.org/docs/submit/osg-flock/>



Open Science Grid

AAI - IDENTITY VETTING LIGO V. VIRGO

LIGO LDAP manages most user identities

Most Virgo users also in the LIGO LDAP

- Less trusted in the LIGO LDAP
- Need different way to get a certificate since its not sufficiently trusted via LIGO

NEXT IS TESTING — DOES IT WORK?

Friendly pipeline to test Nikhef IGWN submit host

- Blew up the home directory (now 200 GB of that 1.2 TB)
- Aggressive checkpointing is also an issue (1 TB now allocated) which is saved to `/var/lib/condor/spool/`.

Added extra partitioning and some more space...

- Great for the time being but not scalable
- We would like to implement user quotas but this requires a change in how the pipelines normally run

SCALABLE SPACE SOLUTIONS

- Nikhef operates with the WLCG model of writing output data to a storage element
- Data management for LIGO is usually archiving users' submit-node home directories for long periods of time
- Potential solution:
 - Can the users write to our dCache storage element?
 - This requires changes to user jobs
 - Need to test if writing out over https will work (i.e., via webdav)
 - Maybe we can mount the Virgo dCache pool on the node?
 - Proxies are not always equal
 - LIGO has CILogon Silver now (ligo-proxy-init will work)
 - Virgo's LIGO proxies are CILogon Basic (voms-proxy-init)

ACCOUNTING

User data available to all collaborators in LIGO accounting db

- Cannot send this data along from the submit host
 - Usually this data should only be available to VO managers

Not really GDPR/privacy compliant

- “Privacy by design; privacy by default”
- Need to redact user names before sending data
- Group- or pipeline-level granularity is ok

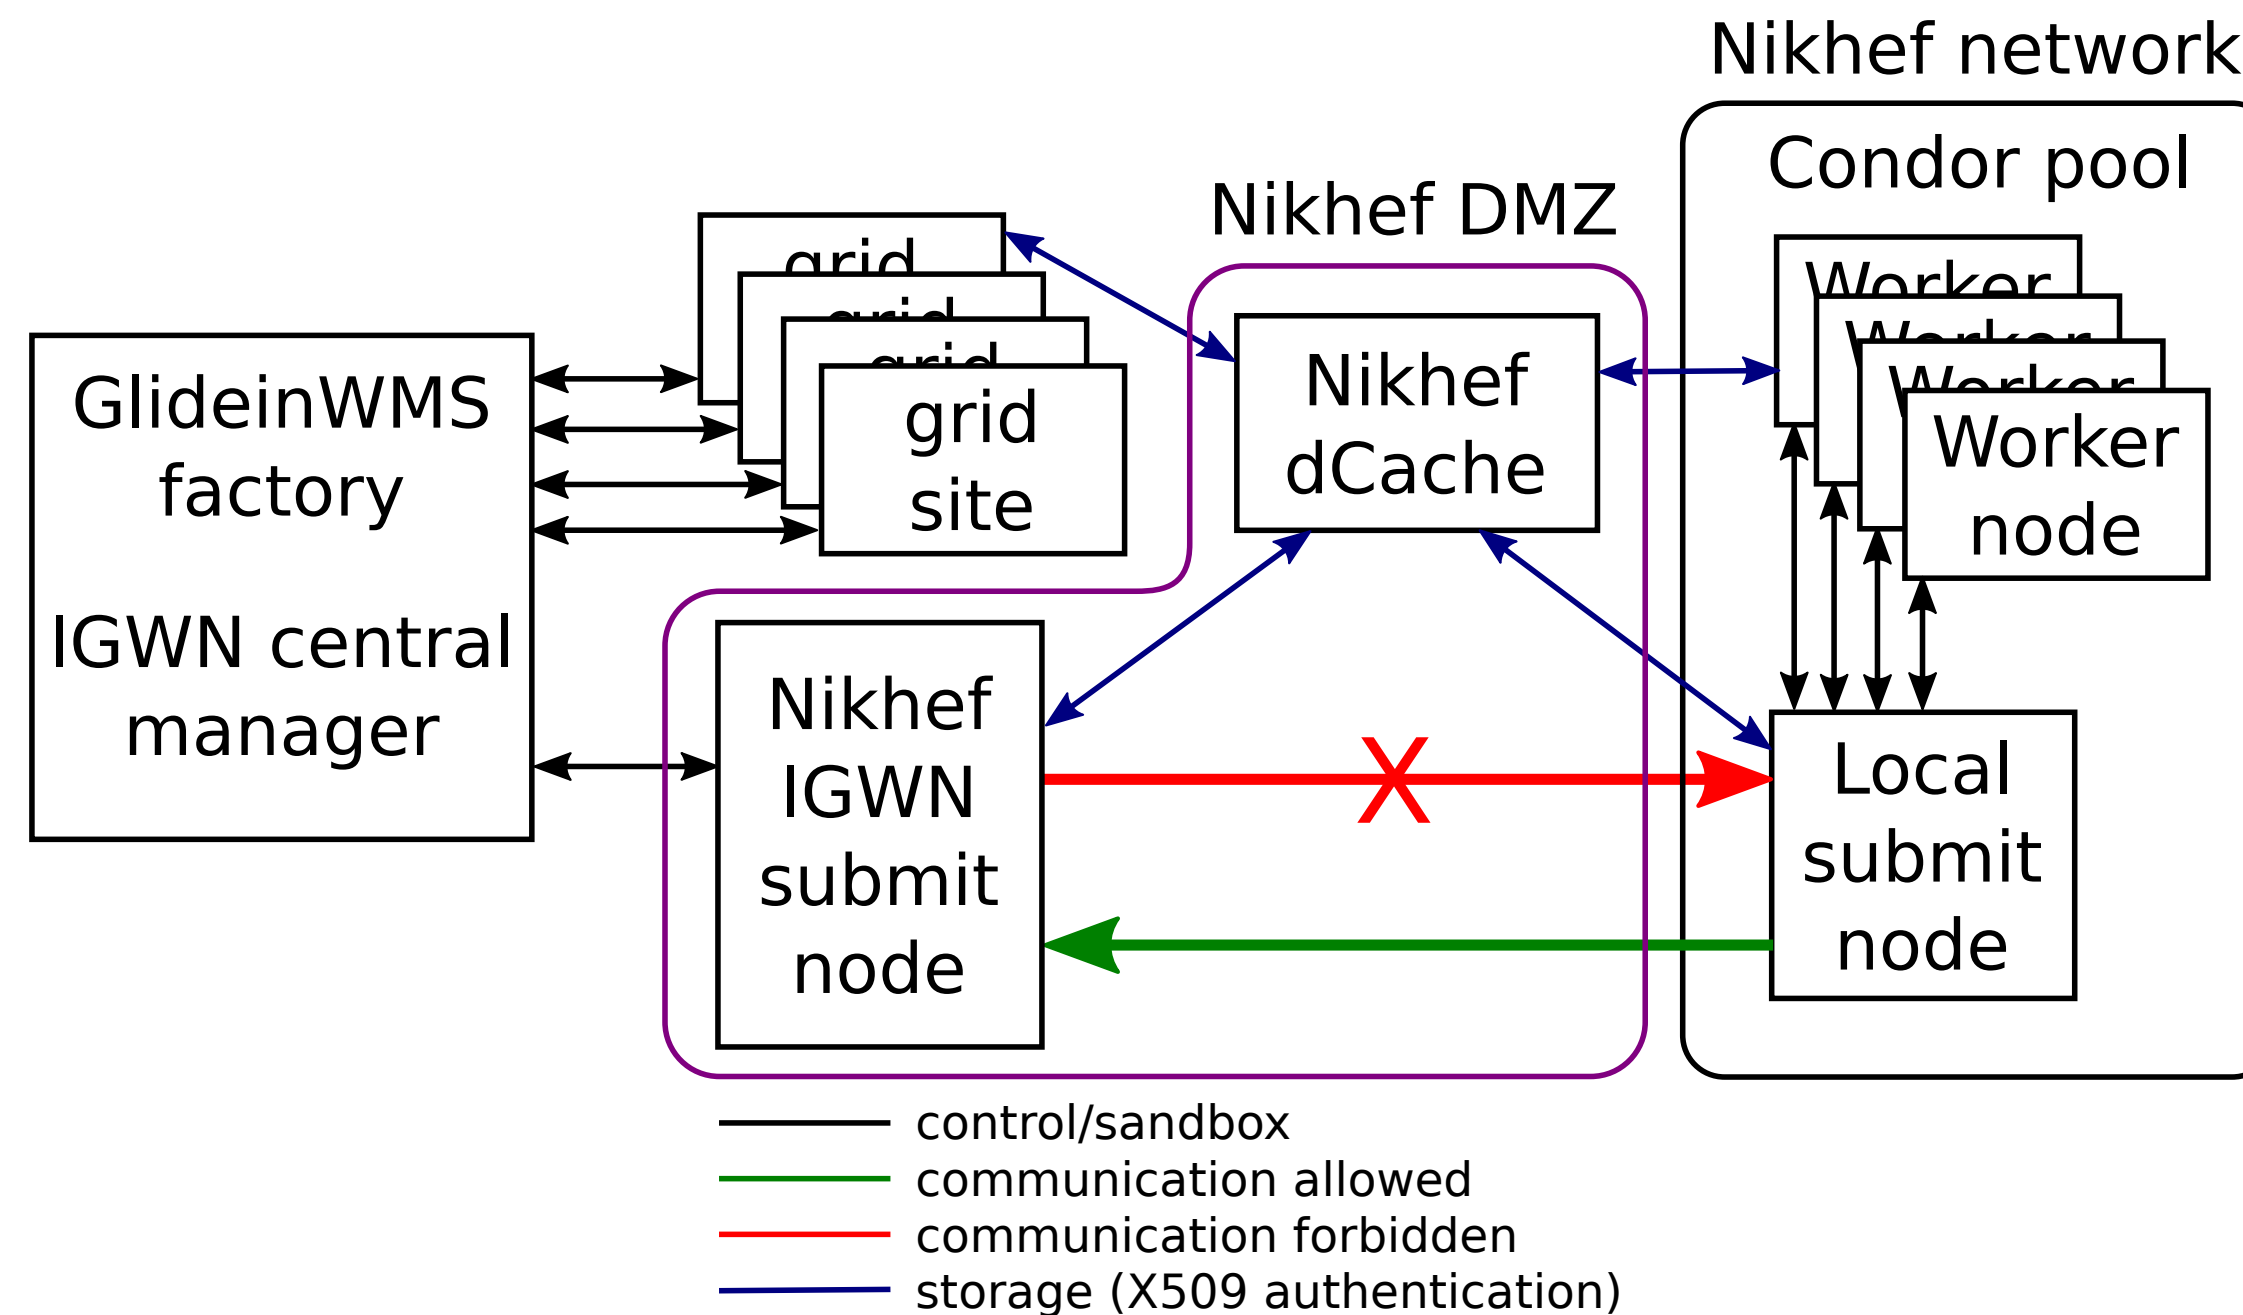
CONCLUSIONS

- Key differences that require(d) some extra work and working with pipelines/users:
 - Redirecting data products to scalable storage (dCache)
 - Accounting details send all information to US server; viewable by all users (not just VO managers)
- Solutions:
 - NFS mount for dCache pool — however could have strange x509 issues for some users
 - Redact usernames from accounting data

FUTURE WORK — BIG PLANS

Plans to integrate local gravitational wave cluster with the world through IGWN submit host

- see Jeff Templon's talk from European HTCondor week: <https://indico.cern.ch/event/936993/contributions/4020898/>
- Using condor_c
- First draft:
 - **Local v. external submit node tbd.**





QUESTIONS?
COMMENTS?
FEEDBACK?

THANK YOU!

GRID.SYSADMIN@NIKHEF.NL

