

GlueX in the integrated global infrastructure for distributed computing

Richard Jones, University of Connecticut



- need for elastic capacity
- myths and realities
- adaptation to OSG+OSDF



Background: *the Gluex experiment*



GlueX at Jefferson Lab

- 9 GeV photons, fixed target
- *Scientific program*
 - search for exotic mesons
 - threshold charmonium states
 - dark matter searches, rare η decays
 - Primakov, pion polarizability

commissioned: 2014-2016

Phase 1 run: 2017-2019

Phase 2 run: 2020-2025

Phase 3 run: *approval pending*




Evolution: *beyond simulation?*

GlueX offline computing resource needs (*GlueX-doc-3813*)

- 35 Mcore-hr/yr - **detector simulation**, primarily on OSG
- 130 Mcore-hr/yr - experimental **event reconstruction**
 - *Jefferson Lab compute facility (total 70 Mcore-hr/yr, calibration and analysis)*
 - *NERSC (proven option, but allocations are limited)*

Data Production at NERSC in 2024



GlueX need: *elastic capacity*



Open Science Grid

from GlueX Computing Report, A. Austregesilo slide 3, May 14, 2024:

2024 ERCAP Request: 350k node hours

2024 ERCAP Award: 150k node hours

2024 NERSC hours used so far: 46k node hours



GlueX-II 2023-01 data set: 300k, will require additional resources

Options: Request more at NERSC, use JLab batch farm, OSG?

Looking down the road to GlueX Phase 3, starting in ca. 2026?

- ***beam intensity x2 relative to Phase 2***

GlueX need: *elastic capacity*



Open Science Grid

Two possibilities:

1. wait 1-2 years for 2023 results from **HPC**
2. exploit untapped capacity with **HTC**

GlueX tradition: *our mythology*



Open Science Grid

myth: a shared story that guides and unites


GlueX tradition: *our mythology*



Open Science Grid

myth: a shared story that guides and unites

- Monte Carlo simulation is appropriate for OSG.
- Event reconstruction needs HPC resources.



GlueX tradition: *our mythology*



Open Science Grid

myth: a shared story that guides and unites

- Monte Carlo simulation is appropriate for OSG.
- Event reconstruction needs HPC resources.

BUT:

- GlueX event reconstruction is inherently HTC!

GlueX tradition: *our realities*



realities: shared experience behind our myths


- our recon framework was developed for HPC
- our recon workflows assume HPC (entire nodes)

new effort required:

- adaptation of applications, workflows to HTC
- better means for access to raw data

GlueX adaption: part 1

1 HPC job (CORI)



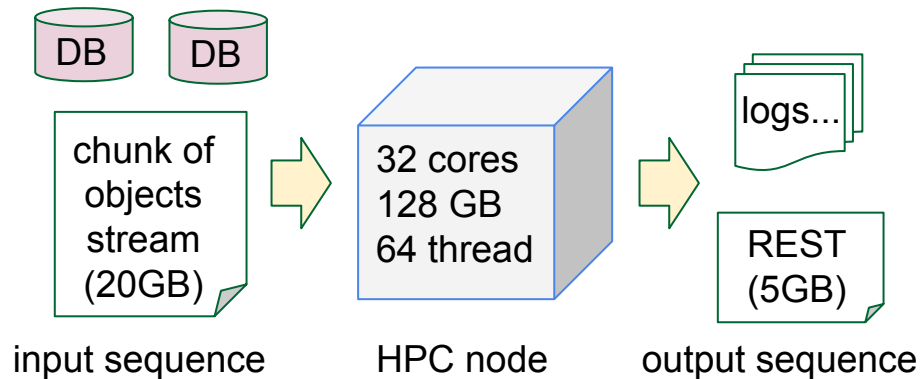
1 HTC job (OSG)

GlueX adaption: part 1




Open Science Grid

1 HPC job (CORI)




1 HTC job (OSG)




GlueX adaption: part 1

1 HPC job (CORI)




- 8 hour jobs
- 10 MB/s i/o rate
- whole-node alloc.

1 HTC job (OSG)



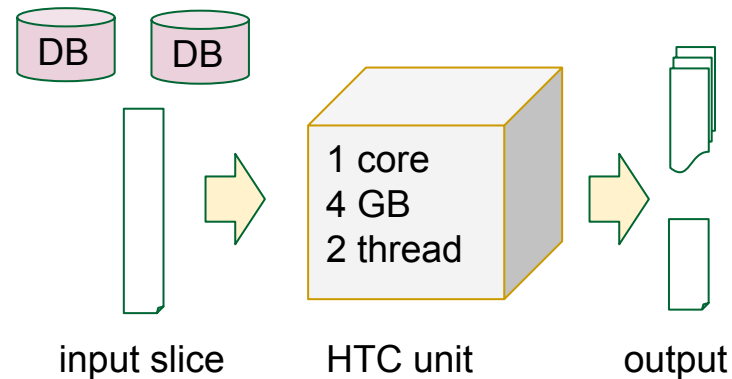
GlueX adaption: part 1

1 HPC job (CORI)



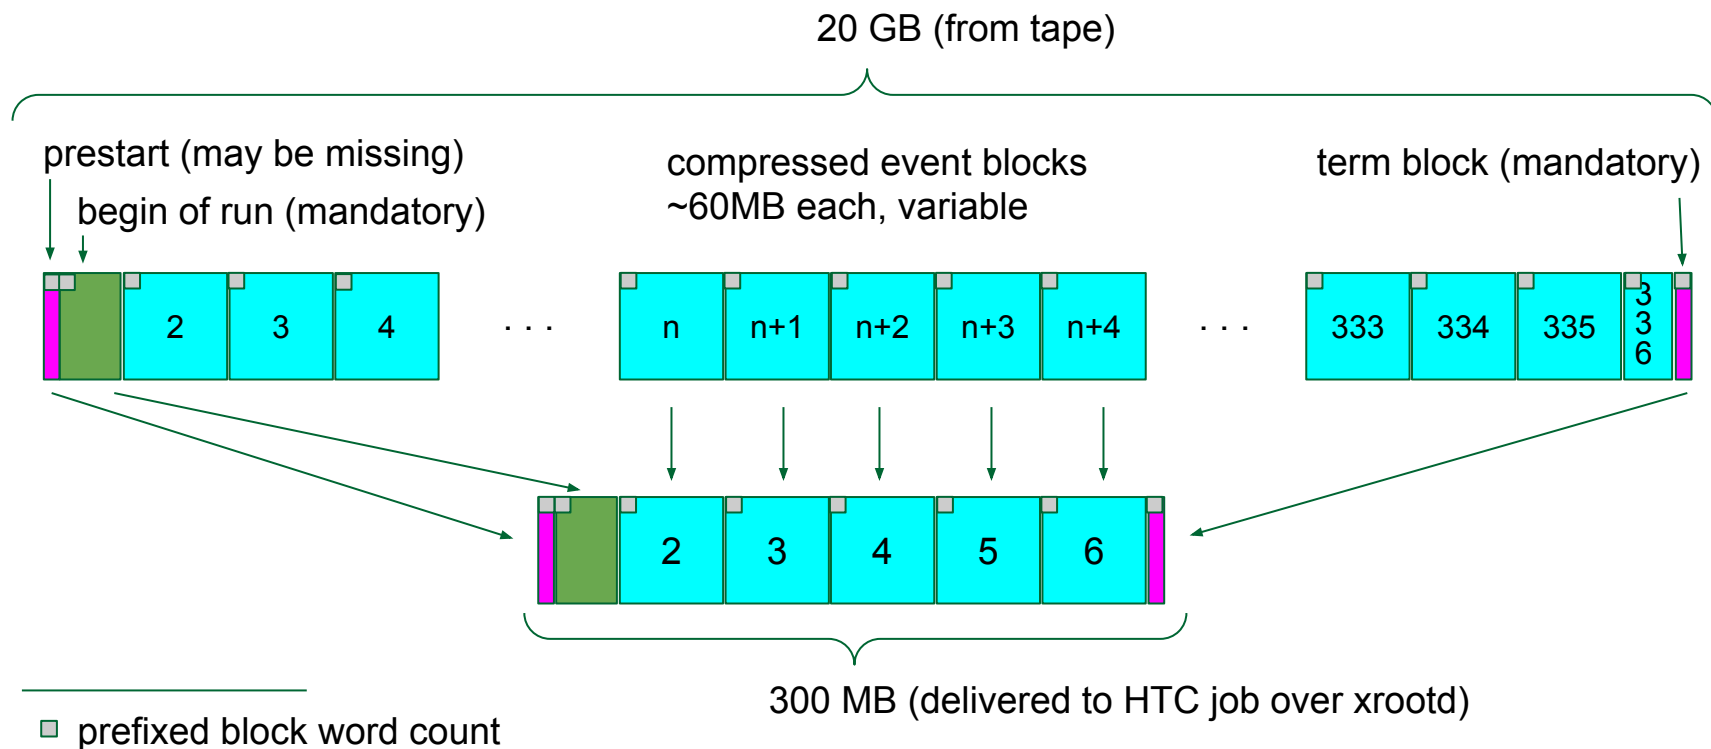
- 8 hour jobs
- 10 MB/s i/o rate
- whole-node alloc.

1 HTC job (OSG)



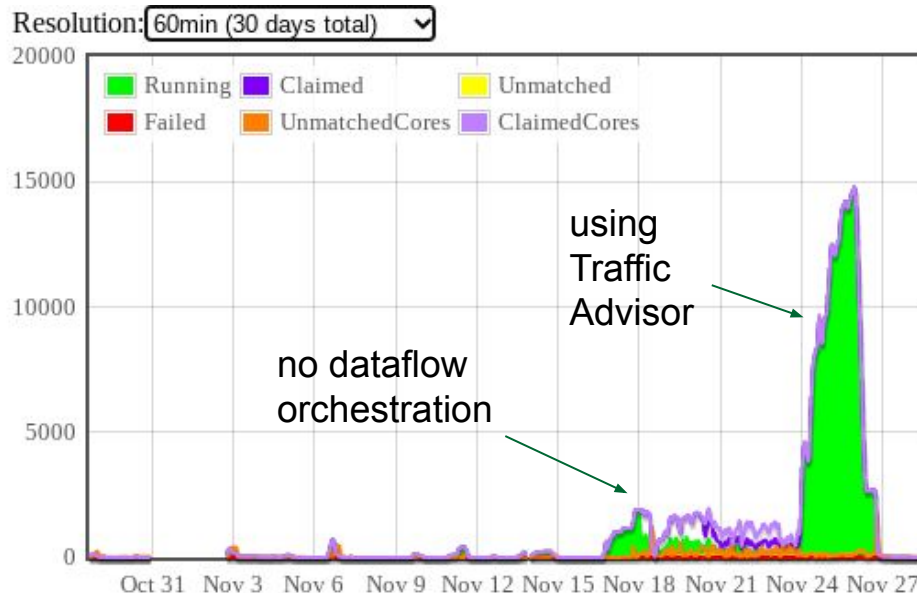
- 2-4 hour jobs
- 0.3 MB/s i/o rate
- single core alloc.

GlueX adaption: raw data stream over xrootd

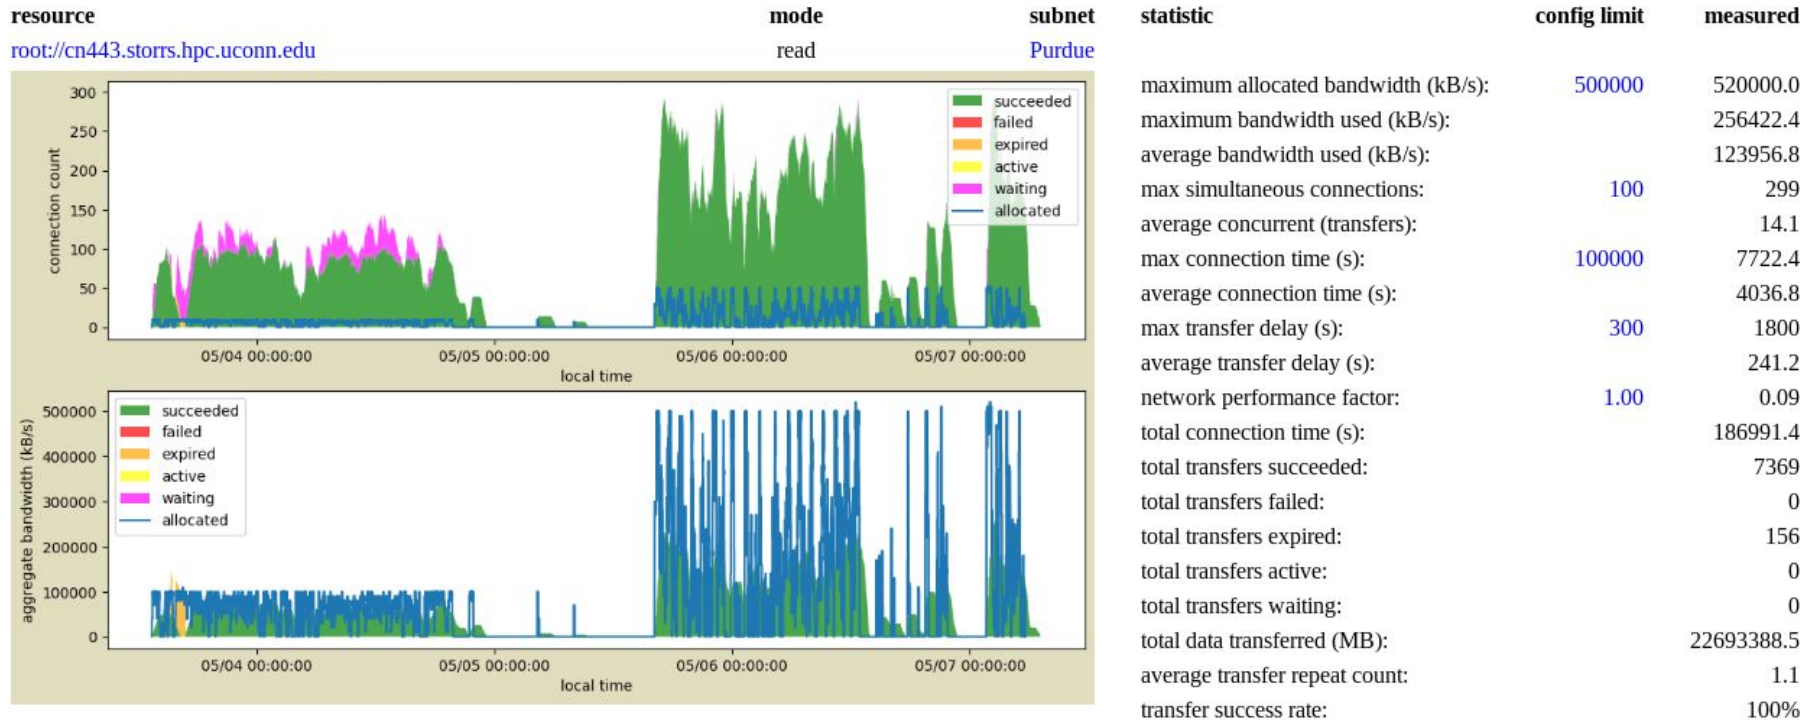


GlueX adaptation: a feasibility study


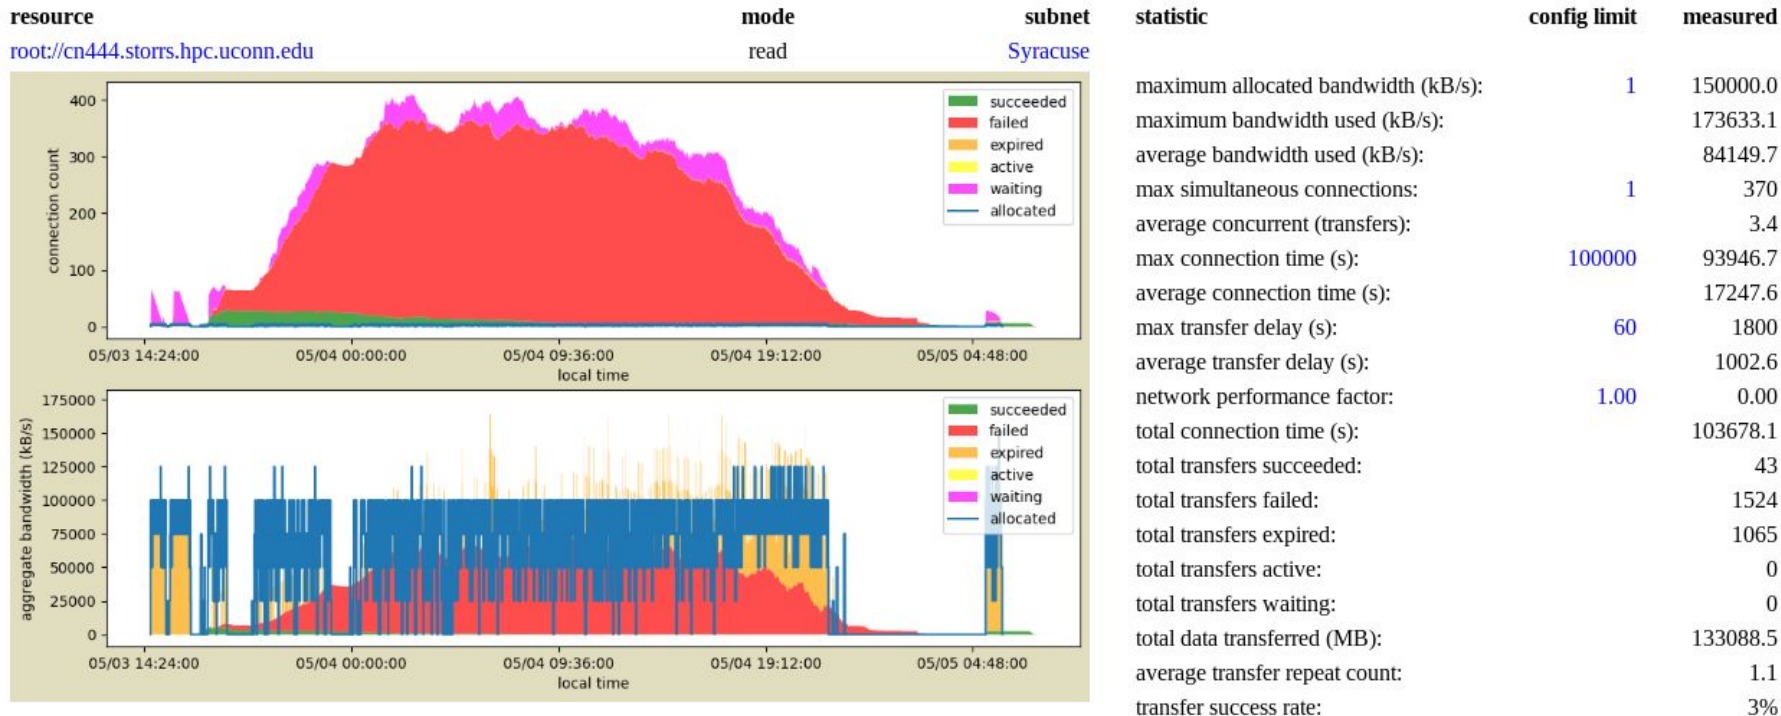
- 10 runs reconstructed on OSG
 - 1880 raw 20GB files
 - 35 TB, 600M events
 - 500k core-hr
- mapping onto osg user jobs
 - 4 cpu-hr / 1-core job
 - 68 jobs / evio file
 - 126,000 jobs
- highlights
 - completed in 3 days, see plot
 - 15,000 running cores, peak
 - job completion rate >90%




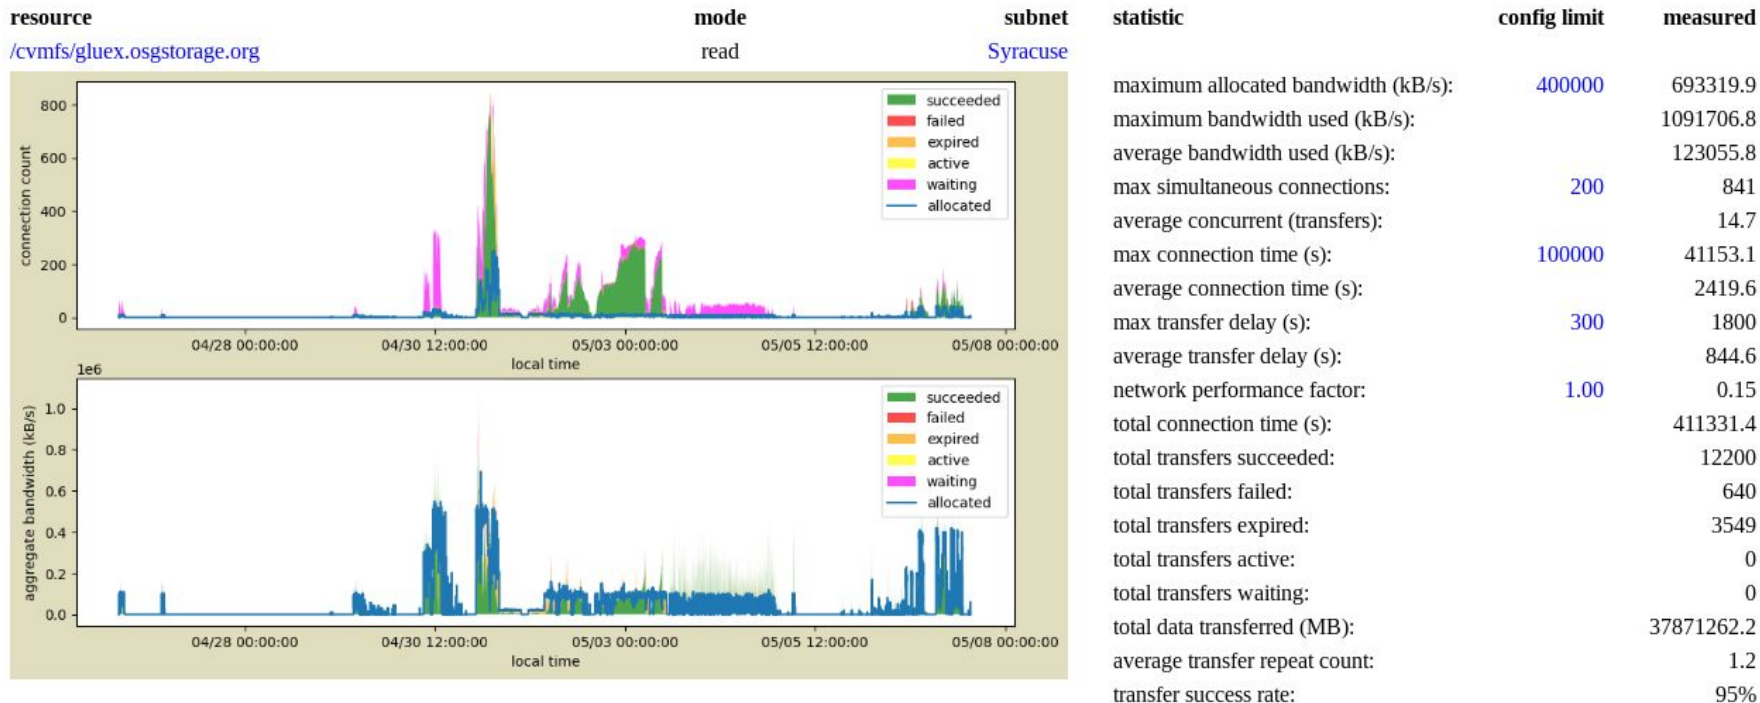
GlueX adaptation: *Traffic Advisor web service*



GlueX adaptation: *Traffic Advisor web service*



GlueX adaptation: *Traffic Advisor web service*



GlueX adaptation: outcomes of feasibility study

myth dispelled?

- yes and no...
- effective deployment awaits **enhanced data access**



The Open Science Data Federation

Backed by the Pelican Platform, the Open Science Data Federation (OSDF) is an OSG service hosting data origins and caches across the globe.

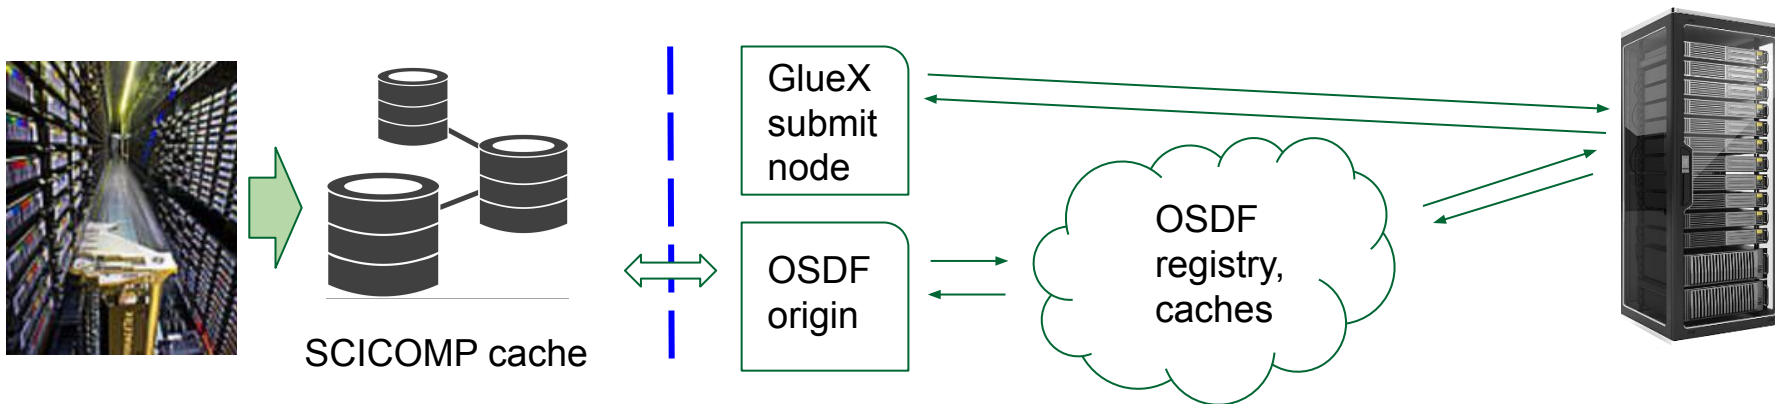
Integrated with other OSG services such as the OSPool, the OSDF facilitates the distributed nature of a national compute pool. In addition to the performance improvements, purpose built plugins for HTCondor make it easy for researchers to harness these caches.



[Learn More →](#)

GlueX data access: *reducing friction with OSDF*

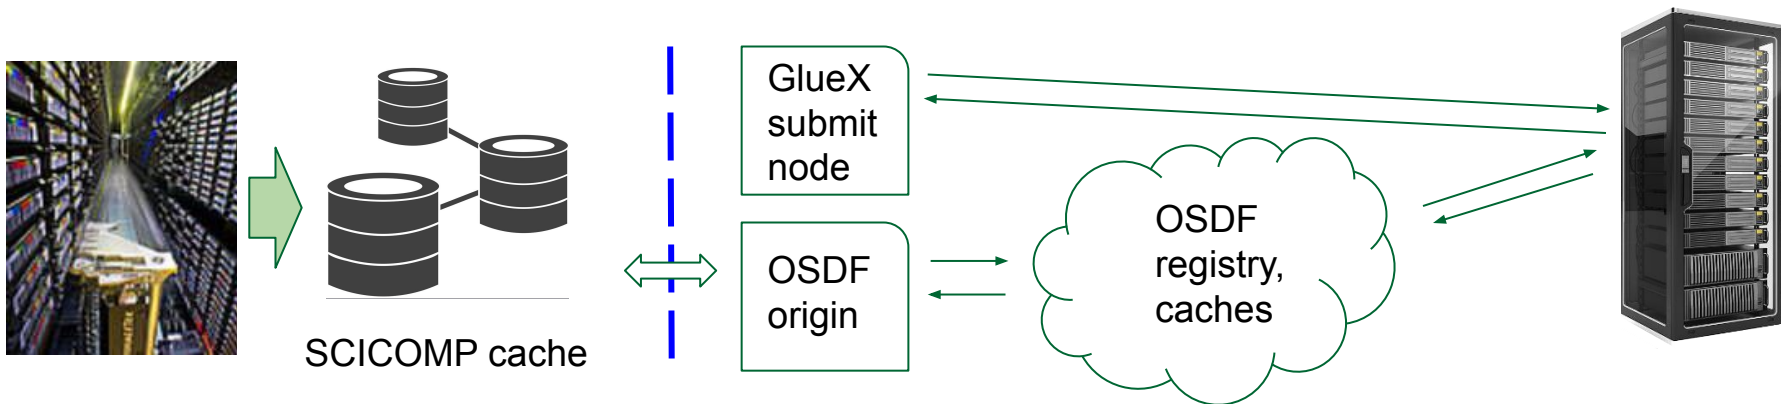
Coming soon for GlueX@jlab:



GlueX data access: *reducing friction with OSDF*

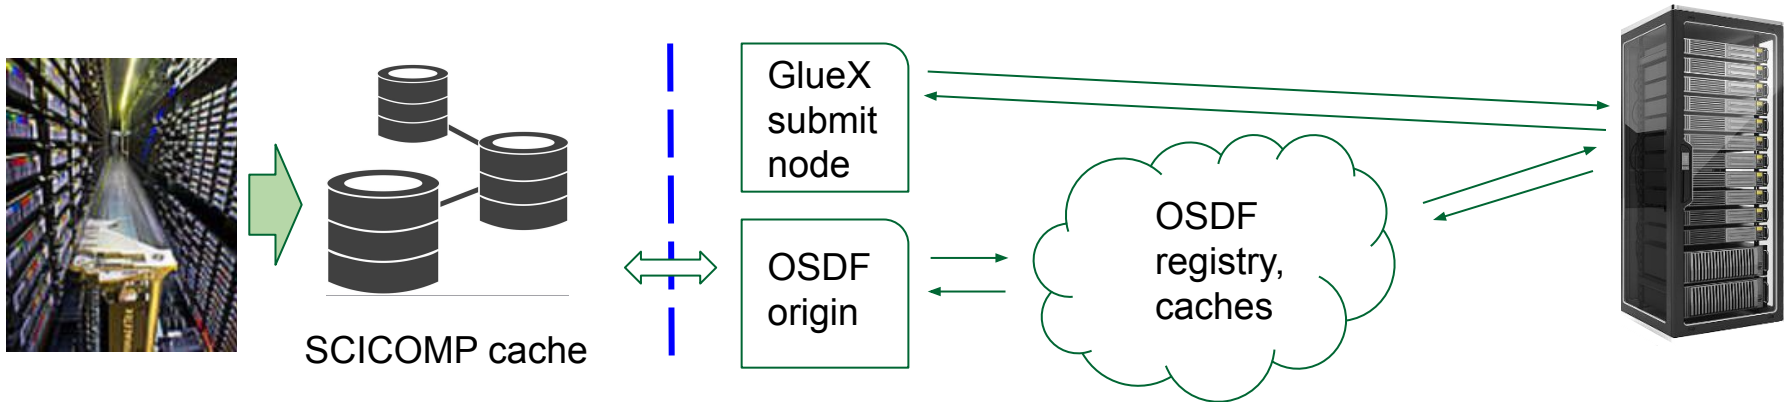
Coming soon for *GlueX@jlab*:

- lab-wide data release policies that regulate access
- cilogon-based token infrastructure to support above
- a new JLab OSDF origin for raw data access



Conclusions

1. GlueX is underserved by current HPC resource allocations
2. HTC on OSG is a demonstrated but untapped resource
3. data access barriers are the present bottleneck
4. JLab storage integration into OSDF will have near-term impact



Backup

GlueX challenge: dataflow choreography


Goal: *port an application optimized for HPC to run on HTC*

- output must be byte-for-byte identical to HPC for same input data
- source code must be identical to HPC, binary compatibility not required

Data distribution strategy:

1. **Split** up the input data into schedulable chunks (2 core-hr)
 - 20GB evio file = 1.4M events = 130 core-hr
 - 340 evio blocks / file
 - 5 evio blocks / job
 - 70 jobs / evio file
2. **Merge** results after processing
 - much smaller data volume
 - can be performed on the SE as part of post-job validation
 - utilities already exist for merging (eg. *hadd*, *eviocat*, etc.)

raw data reconstruction on shared HTC



Open Science Grid

Submit locally, run globally.


OSDF
x
r
o
o
t
d



the challenge of data distribution:

- 1.0GB/s staged on tape at Jlab
- 150MB / core-hr reconstruction rate
- 65k cores in steady-state to keep up
- *Is this scale feasible on shared HTC?*

HPC



PSC



NERSC

an *effective, light-weight* solution

what changed in the xrootd client interface?

- all existing preload library functionality retained
- added: recognition of a specially formatted xrootd url at open
- overlap of processing, fetching next block

To access the full input file, use the standard xrootd url

root://cn440.storrs.hpc.uconn.edu/gluex/rawdata/Run071728/hd_rawdata_071728_000.evio

To get a subset with just 2 evio blocks + BOR + term

root://cn440.storrs.hpc.uconn.edu/gluex/rawdata/Run071728/hd_rawdata_071728_000.evio :14,16

how NOT to subset a 20GB raw data file

- ✗ Split data to 300MB files on the SE prior to job submission
 - ✗ introduces extra step in processing
 - ✗ doubles load on the SE
 - ✗ not necessary
- ✗ Tell event reader to skip forward N events before start of processing
 - ✗ increases total data transfer load by factor ~ 35
 - ✗ large cpu load increase from block unpacking
 - ✗ assumes we know in advance how many events per file
- ✗ Tell event reader to skip forward N blocks before start of processing
 - ✗ still need to know when to stop
 - ✗ requires modification of the reconstruction software
 - ✗ violates goal to run ***the same software release on all production sites***

an effective, light-weight solution

customize the XRootD Posix preload client library

- no modifications to the Gluex software stack
- no modifications to the xrootd protocol, OSDF infrastructure
- changes are local to the OSG singularity container
- correctness can be verified prior to running production
- maintainable as a fork of the XRootD client github repo

what happened in this time?

- data had already been staged from JLab to UConn (*staging time not included*)
- data distribution governed by a postgres database
 - a. jobs “checked out” chunks of data via database transaction at job start
 - b. successful completion of chunks registered to database at job exit
 - c. reservations expired periodically, uncompleted chunks recycled
- *all jobs were identical, no predefined job-data association*

what happened in this time?

1. jobs were submitted from UConn submit node `gluex.phys.uconn.edu`
 - a. shared same gluex factory with regular GlueX simulation jobs
 - b. standard gluex queue, no special boosted privileges
 - c. input data pulled during processing by osg worker
 - d. output copied back to UConn SE at job exit
2. job outputs automatically merged
 - a. performed by cron job running on the UConn-HTC cluster (colocated with SE)
 - b. one-step merging, verification, and push of merged results to Jlab
3. not everything worked perfectly...

what worked well?

1. postgres database as a management hub
 - a. jobs requested their work from the database when they started up
 - b. jobs returned their exit code and completion time when they finished
 - c. no problem with scaling
 - i. ~5 jobs starting/finishing per second
 - ii. merger cron job multiplicity scaled to keep up with job output
 - iii. management was hands-off -- **watch but do not touch!**
 - d. only failed slices were reprocessed (1% of a full raw data evio file)
 - e. recycling unprocessed slices was immediate -- within seconds!

2. output merging + verification + push to Jlab done in one step

where were the difficulties?

- ★ reconstruction app. memory footprint: RSS = 5.5GB @ 1 worker thread !!
 - reconstruction library alone = 1.4GB -- **good!**
 - there are a lot of plugins -- 29 of them
 - most of the bloat comes from root histograms -- 17,000 mostly TH2I
 - remove the 5 worst offenders = 2.5GB -- **kinda ok, not ideal**
- ★ thread explosion
 - reconstruction code had a strange behavior, required intervention
- ★ hanging transfers pushing output back to JLab
 - engaged with UConn NetOps to resolve this
 - very difficult to diagnose problems at the JLab firewall from outside
 - connections were being unpredictably dropped -- **hanging sockets**
 - investigated further with JLab NetOps -- **hanging communication**

Results and assessment

- robust performance by redundant dcache xrootd server “doors”:
 - 6 instances: cn440...445, but only needed one
 - *door* functions only to forward connections to “movers” on “pool” hosts
 - input + output storage spread equally over 38 hosts
 - internal network 100Gb/s infiniband
 - external network 10Gb/s ethernet
- at peak, load remained reasonable (50 simultaneous transfers, 2 min. max)
- no long-lived sockets, all connections fetch just one block and disconnect
- no problems observed with database transaction rates, all good!
- ***BUT*** results were *not* byte-by-byte identical to HPC production -- *unrelated to OSG*

concluding comment: enabling role of OSDF

- UConn dcache SE is a OSDF origin (UConn-HPC_StashCache_origin)
 - 750TB of shared network-visible storage
 - used for staging input + output to/from tape @ Jlab
- data are read-once during reconstruction
 - raw data not suitable for caching
 - *nevertheless, OSDF data access was a major advantage!*
- **xrootd access via OSDF worked well**
 - many sites have limited direct network access to offsite
 - accessing through the OSDF url often factor 10 x faster than direct
 - sites have configured special routes for OSDF access -- *enables this capability*

Backup slides

an effective, light-weight solution

what XRootD components needed to be touched?

- github fork rjones30 / xrootd from master on 11/12/2020
- updates enclosed in `#ifdef EVIO_BLOCK_SUBSET_EXTENSION`
 - `src/XrdPosix/XrdPosixFile.hh` (40 added lines)
 - `src/XrdPosix/XrdPosixFile.cc` (37 added lines)
 - `src/XrdPosix/XrdPosixXrootd.cc` (237 added lines)