
TACC Frontera Integration Efforts for ATLAS

Lincoln Bryant, Doug Benjamin
9/4/2020

ATLAS
EXPERIMENT

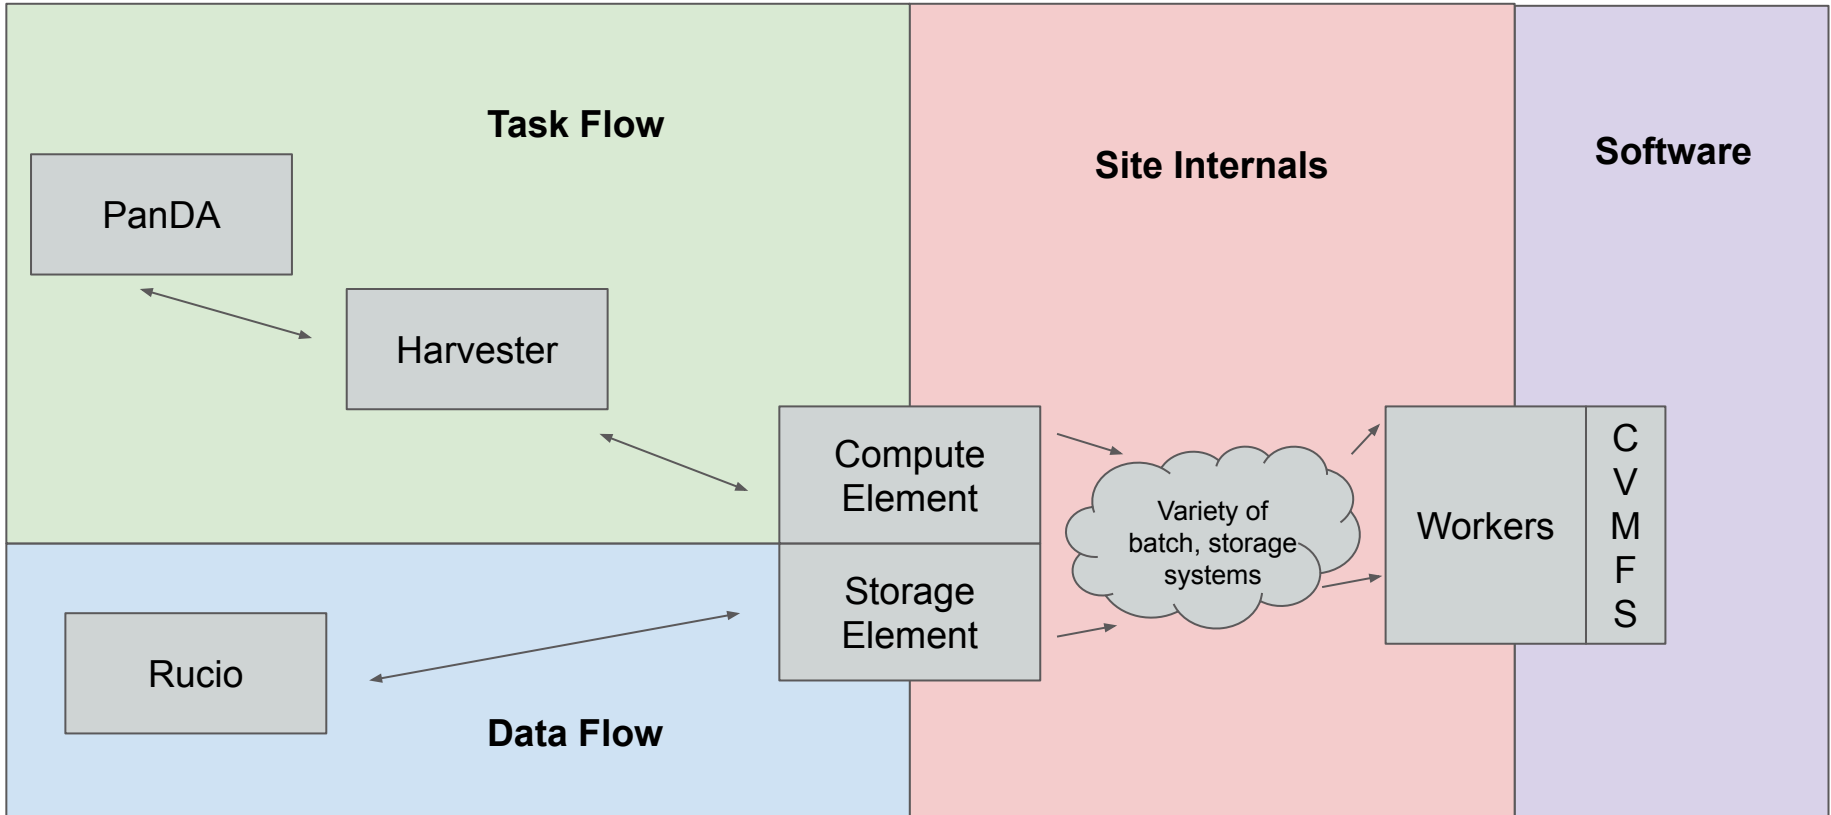
Frontera

- \$60M machine deployed in 2019 at the Texas Advanced Computing Center (TACC)
- About 450,000 Cascade Lake Xeon Platinum CPU cores and all the trim

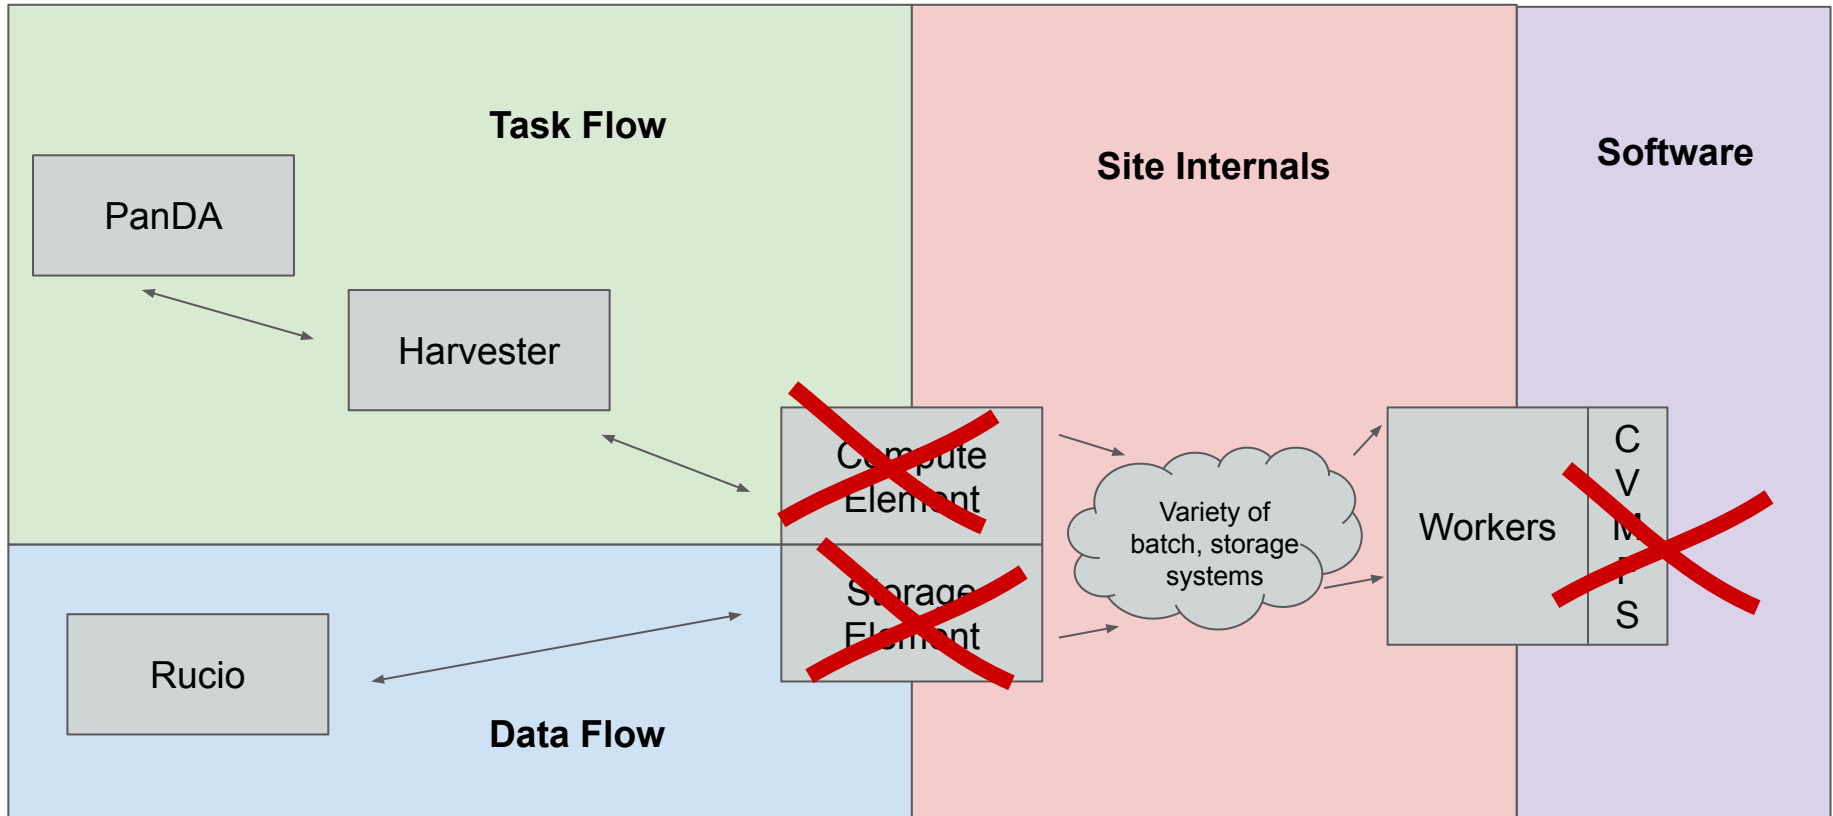


**ATLAS and CMS both
awarded Large Scale
Community Partnership
awards for a sizeable chunk
of service units**

High-level overview of ATLAS workflow system



Much of this does not exist at TACC and other HPCs



Challenges and Solutions

- To provide software
 - **We bring our own Singularity container with the ATLAS environment**
- To stage data in and out
 - **We proxy data (via SLATE) between Globus and Rucio**
- To submit and manage tasks
 - **We run Harvester locally to communicate over the shared file system**

SLATE @ TACC

- Part of this work took advantage of the SLATE platform
- In a nutshell: **SLATE provides a federated Kubernetes platform on which to launch various long-lived services at computing sites like WLCG sites and HPC centers**
- Offered to send SLATE hardware to TACC to support ATLAS and others in the community
 - now also being explored by CMS (more in Dirk's talk)
- TACC was enthusiastic and the SLATE security proposal passed muster
- Hardware was sent and set up in March just before COVID19 hit
- 192GB RAM, ~35TB of NVMe, Mellanox 50Gbps card

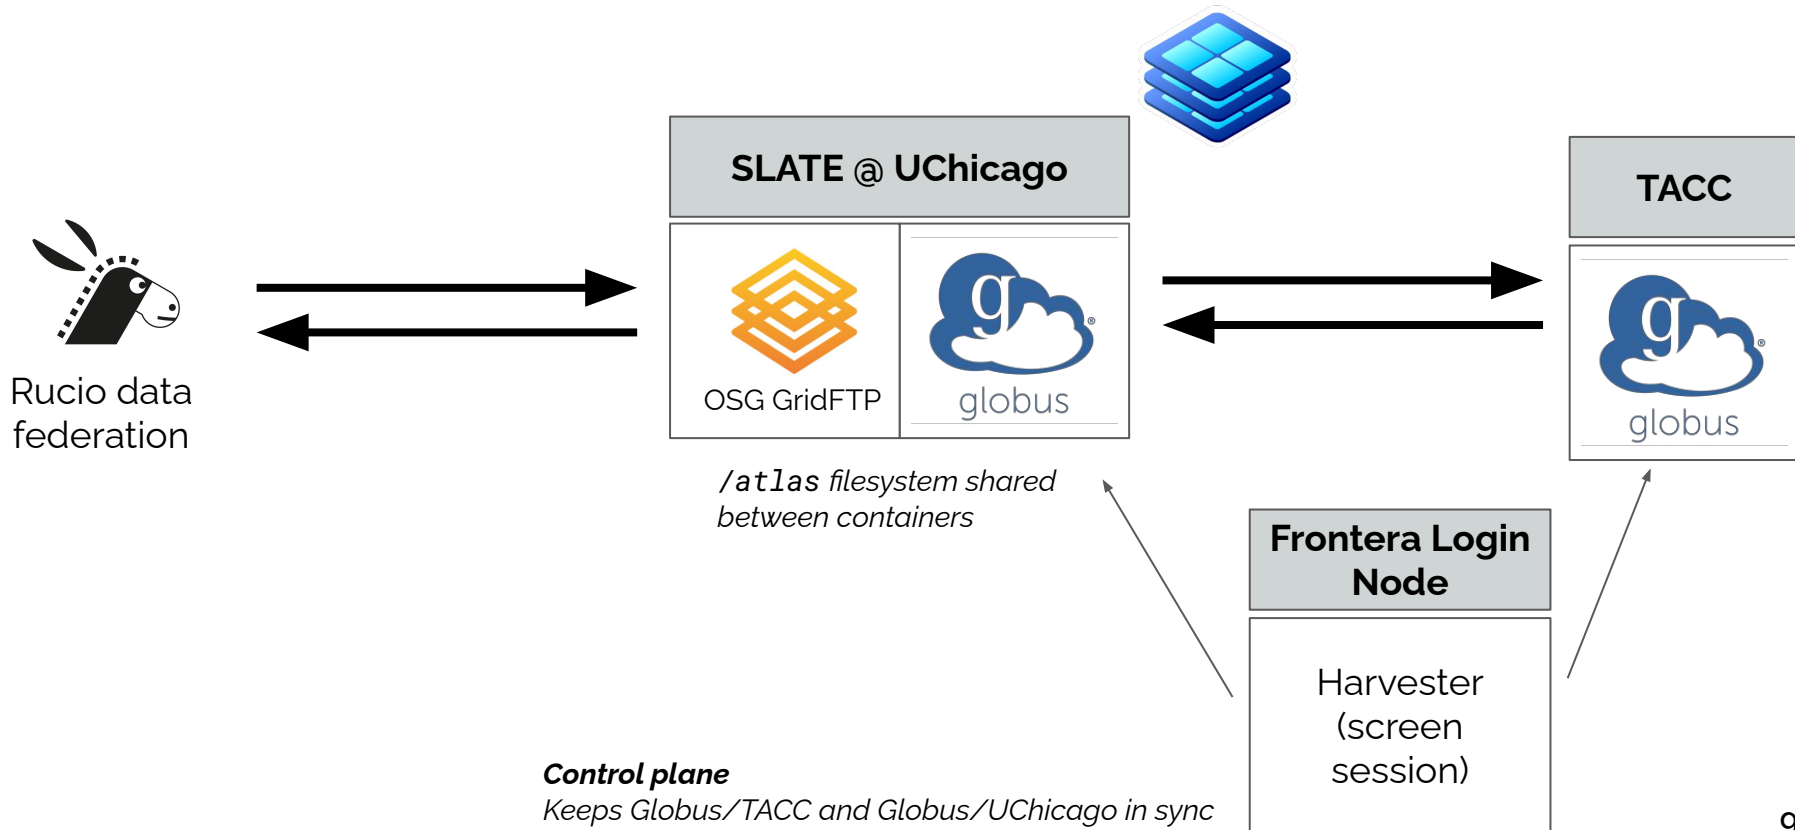
Software Environment

- ATLAS produces Simulation software releases
 - Published in two ways: Docker images and via CVMFS.
- These Docker images are converted (by us) into Singularity images and staged to local storage at TACC
- Prodsys2 is use to create tasks that are submitted via PanDA to TACC via Harvester.

Data Environment

- Frontera doesn't have a traditional GridFTP endpoint in the WLCG/OSG sense but does have a Globus Online endpoint
 - Common to many HPCs
- Harvester developers implemented a plugin to synchronize Globus Online with GridFTP endpoints
- **Globus Online and OSG GridFTP containers launched on SLATE node at UChicago, sharing a common filesystem for keeping TACC in sync with the rest of the ATLAS data federation**
- R&D effort ongoing to have Harvester/Rucio directly use Globus Online

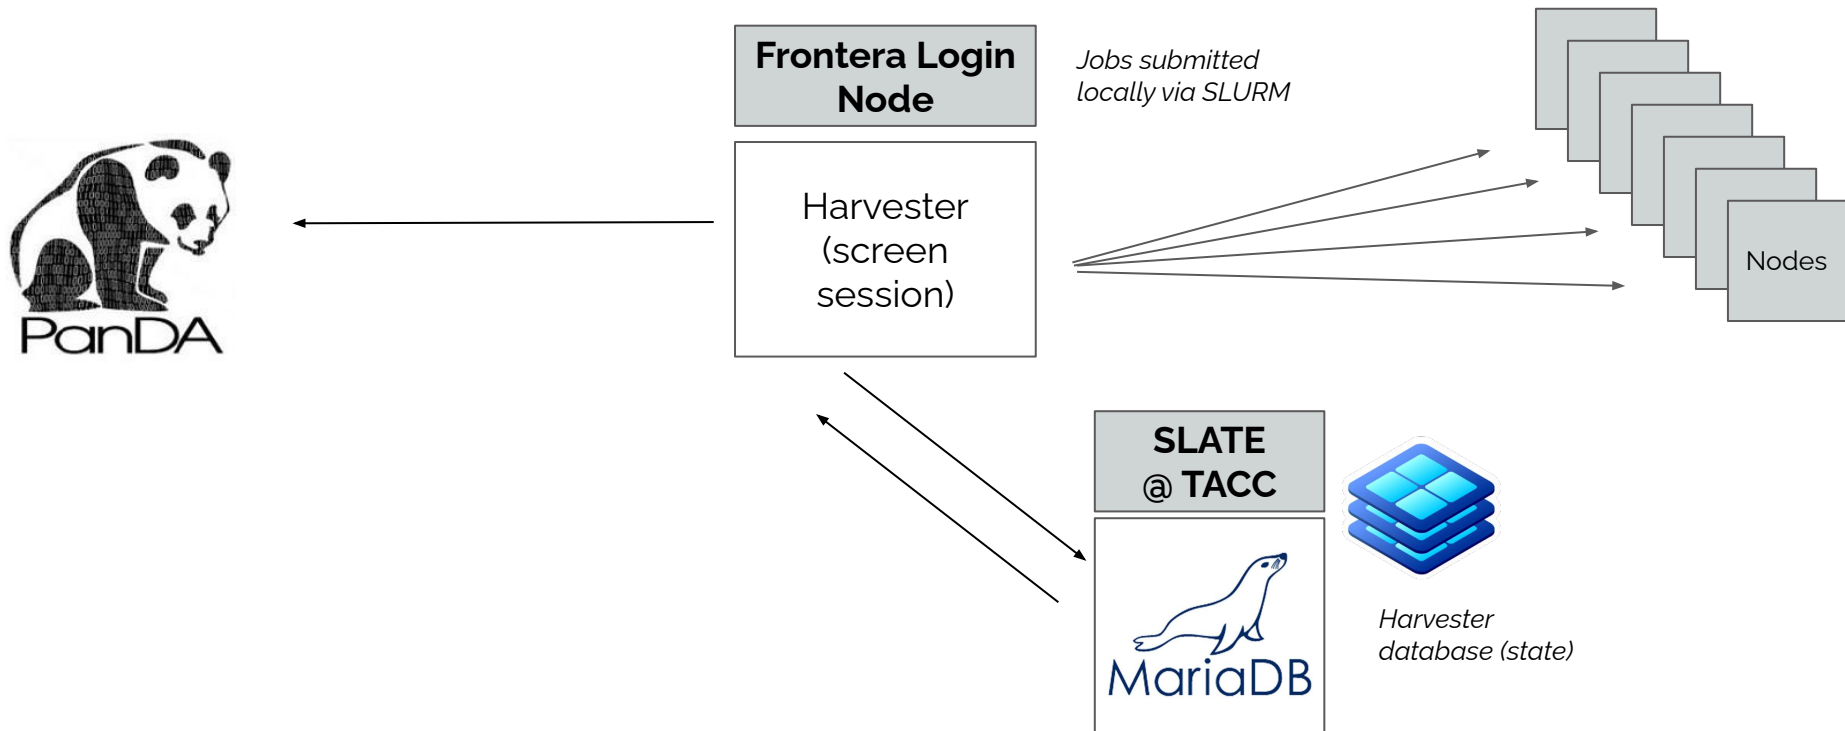
Data Flow



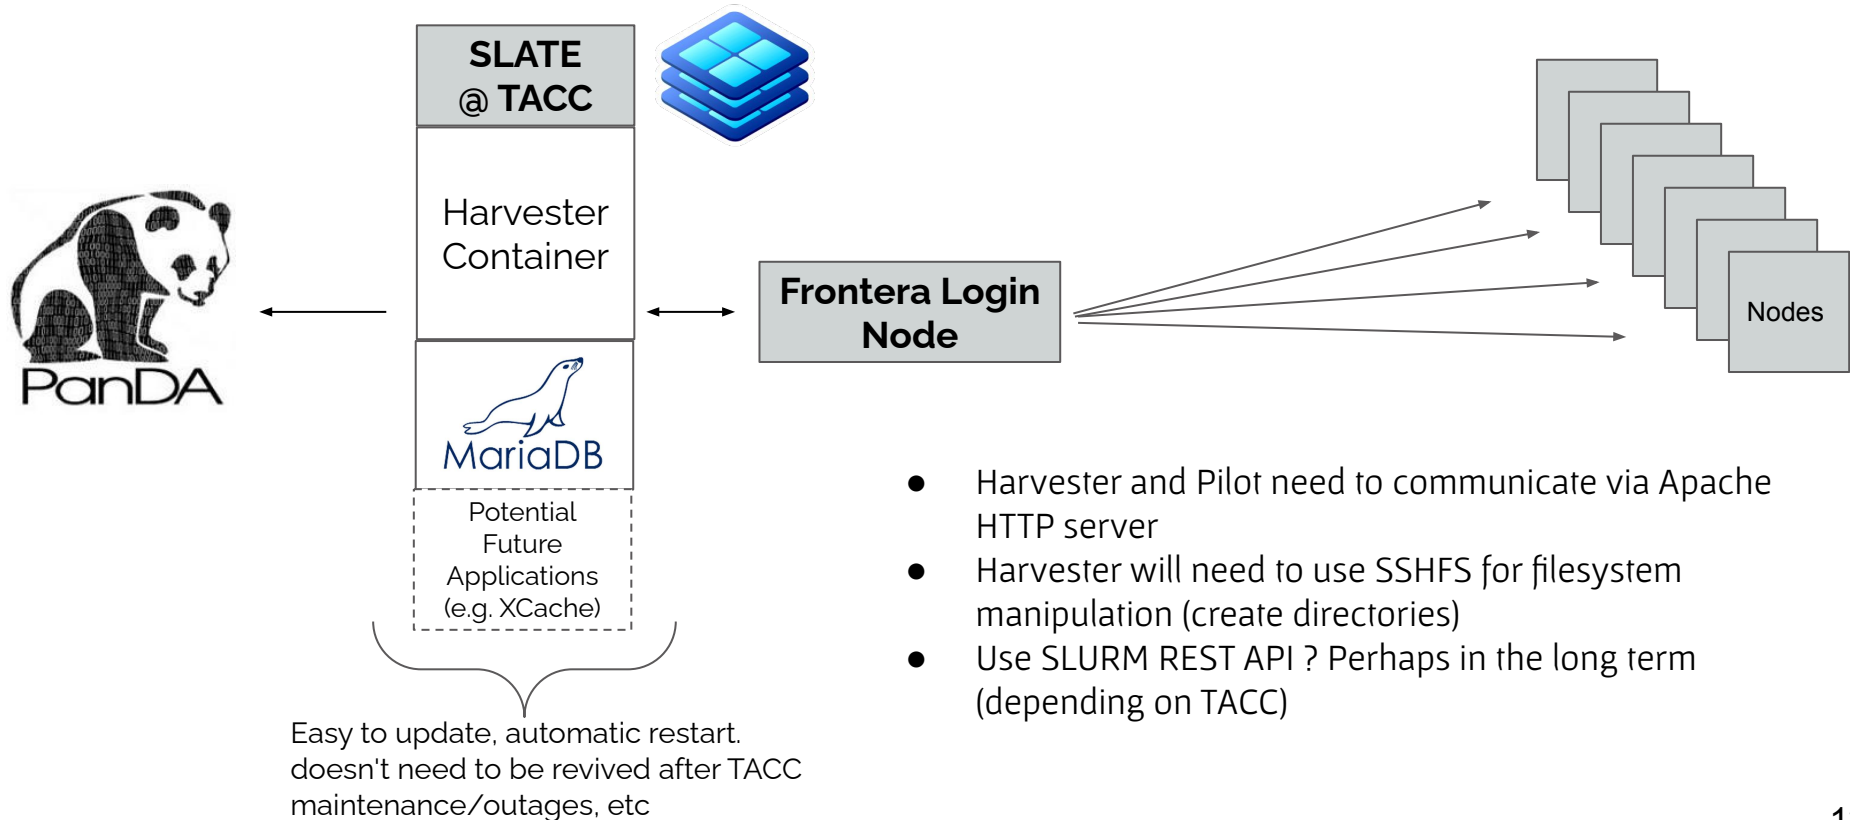
Job Environment

- Frontera doesn't have a Compute Element as we might expect of an OSG site
- Moved Harvester to the Edge at TACC
 - History of managing HPCs in ATLAS in US and Europe, common solution
 - Compliant with site policies
- Harvester mediates and monitors:
 - Task submission to SLURM
 - Triggering Globus transfers to/from our Rucio endpoint at UChicago/SLATE
 - Shared filesystem-based communication to/from Frontera worker nodes

Job Flow (today)



Job Flow (planned)



Developing Kubernetes-based edge services


- Kubernetes edge platforms becoming very important & popular at sites
 - SLATE at TACC
 - Spin at NERSC
 - PetrelKube at ALCF
 - Slate at OLCF (unrelated to NSF SLATE)
- Use programmatic interfaces to leadership computing facilities
- Helm etc -> Code reuse

Status and new challenges

- Foundational pieces are in place, can successfully run jobs
- We are slowly scaling simultaneous nodes per job on Frontera, to get access to thousands of cores while being trying to not crash the Lustre filesystem
- Serial phases of the execution knock down the efficiency
- Will be investigating other modes of running, including "Raythena" (Ray distributed task framework + Athena framework), Jumbo jobs, and cycle scavenging queue

Results / thoughts

- Have scaled up to 250 nodes simultaneously before Frontera went down for "TEXASCALE" mode this week.
- TACC support have noted that (perhaps) a concern area is 'inode_permissions', from traversing deeply nested directory structures.
- Ready to go to production for simulation
- Will almost certainly continue to use containers. We are working on getting additional containers produced so that we can do more of the MC chain at TACC. file merging, reconstruction etc



Job list. Sort by **PandaID**, time since last state change, ascending mod t

PanDA ID Attempt# of maxAttempts#	Owner Group	Request Task ID	Transformation	Status
4793142253 Attempt 1 of 1	bdouglas AP_VALI	15259 21950602	Sim_tf.py	finished
	Job name: valid1.424108.Pythia8B_A14_CTEQ6L1_Jpsimu6			
	Datasets: In: mc15_13TeV:mc15_13TeV.424108.Pythia8B_A14_CTEQ6L1_Jpsimu6 Out: valid1.424108.Pythia8B_A14_CTEQ6L1_Jpsimu6			
4793142251 Attempt 1 of 1	bdouglas AP_VALI	15259 21950602	Sim_tf.py	finished
	Job name: valid1.424108.Pythia8B_A14_CTEQ6L1_Jpsimu6			
	Datasets: In: mc15_13TeV:mc15_13TeV.424108.Pythia8B_A14_CTEQ6L1_Jpsimu6 Out: valid1.424108.Pythia8B_A14_CTEQ6L1_Jpsimu6			
4793142250	bdouglas AP_VALI	15259 21950602	Sim_tf.py	finished

Thanks!

(and special thanks to Todd Evans at TACC for all of the support)

Any questions?