



# WBS 2.3.5 Continuous Integration and Operations (CIOPS)

Ofer Rind/BNL

US ATLAS Software and Computing Operations

June 2022



## 2.3.5: Overview of Activities

### **CIOPS is Continuous Integration (of new services, infrastructure & updates) and Operations (of existing services)**

- ❖ **Provide crucial operational support to the US Cloud**
  - Includes management of ADCoS Shifts, troubleshooting site problems, improving monitoring and documentation for admins
- ❖ **Support and coordinate the deployment of ATLAS and OSG service updates across the US Cloud**
  - For much of the past year, the focus has been on preparing the US Cloud for Run 3 data taking
  - Includes new deployments/capabilities such as those involving the transition to token-based AAI infrastructure and deprecation of the GSIFTP protocol in favor of HTTP-TPC, Containerized Edge Services (SLATE), OSG Site Topologies, and more
  - **Cross-cutting efforts with 2.3.2**
- ❖ **Develop new ATLAS workflow components**
  - Includes Virtual Placement Queues/XCache, Analytics Infrastructure and ServiceX development
  - **Cross-cutting efforts with 2.4 and 2.3.4**



## 2.3.5.1 ADC Ops FTE Summary

	FY22	FY23	Comment
Ofer Rind	0.30	0.50	Coordination & US Cloud Support Shift of effort from 2.3.4
Mark Sosebee	0.50	0.00	ADC Liaison & US Cloud Support Mark retired in 1Q22
Qiulan Huang	0.25	0.00	US Cloud Support Shift of effort to Tier-1
Xin Zhao	0.20	0.00	US Cloud Support Left SDCC
Armen Vartapetian	0.25	0.25	US Cloud Support
Armen Vartapetian	0.50	0.50	Central ADC Ops Coordination
Total	2.00	1.25	



## 2.3.5.2 Service Development and Deployment FTE Summary

	FY22	FY23	Comment
Ilija Vukotic	0.30	0.30	ADC Analytics Platform Development
Ilija Vukotic	0.30	0.30	Caching DevOps, Virtual Placement R&D
Wei Yang	0.25	0.25	XRootd/XCache Development
Andrew Hanushevsky	0.20	0.20	XRootd/XCache Development
Armen Vartapetian	0.25	0.25	Site level Kubernetes
Ed Dambik	0.50	0.50	FTS Development and Operations
Rob Gardner	0.2	0.2	Facility R&D Design & Coordination, Federated operations
Total	2.0	2.0	



## 2.3.5: Selected Highlights

- ❖ Completed transition of all US sites to HTTP-TPC as primary transfer protocol with no GridFTP fallback
  - Provided extensive development support for XRootd sites
- ❖ Supported the update of most site gateways to supported OSG software versions with deprecated X509 proxy support
  - Coordinated with OSG and ADC to prevent the US cloud from operating in a mode with insufficient token infrastructure support
- ❖ Maintained, with 2.3.2 management, a site status [spreadsheet](#) for centralized oversight of facility Run 3 readiness
- ❖ Removed obsolete Tier 3 storage elements from CRIC and assisted remaining sites in transition from GridFTP to HTTP-TPC
- ❖ Implemented GitOps for management of SLATE services and analytics infrastructure
  - SLATE Squid Ops [manual](#) under development
- ❖ US Cloud participation in successful WLCG Network Data and Tape Challenges
- ❖ Development, deployment and testing of ServiceX at UC



## 2.3.5: FY22 Milestone Status

### ❖ MS55. Multi-site Kubernetes platform (A. Vartapetian, R. Gardner)

Delayed - Sep 2021 (orig) - Sep 2022 (new, *see below*)

- Use the Kubernetes container orchestration platform, with a special instance of Harvester, for the purpose of scheduling and running ATLAS Panda jobs on a large cluster at UTA
- Currently accepting test jobs on a 600 core cluster at UTA, but technical issues preventing job starts (details in [backup slide](#))
- Plan:
  - Run a variety of ATLAS production workloads and workflows with Kubernetes to understand inefficiencies and bottlenecks
  - Tune system to optimize cluster performance and evaluate benefits of adopting the Kubernetes model
- Risks: Kubernetes has not been tested in this mode at scale
  - **Impact of intractable scaling problem would be reversion to standard batch operation**



## 2.3.5: Ongoing M&O Activities

- ❖ Coordination of 2.3.5 effort and facility-wide tasks and updates: Rind (0.30)
- ❖ US Cloud Support and ADC Liaison: Sosebee (0.50)
- ❖ US Cloud Support: Huang (0.25), Zhao (0.25), Vartapetian (0.25)
- ❖ Central ADC Ops Coordination: Vartapetian (0.50)
- ❖ ADC Analytics Development: Vukotic (0.30)
- ❖ Caching DevOps, Virtual Placement R&D: Vukotic (0.30)
- ❖ Xcache and Xrootd software support: Hanushevsky (0.20), Yang (0.25)
- ❖ Site-level Kubernetes development: Vartapetian (0.25)



## 2.3.5: New Activities

- ❖ Continue to evolve operational structure and integration with ADC
  - This would greatly benefit from a 0.5 FTE focused on this role
- ❖ Continue progress toward implementing token-based auth [FY22-24]
  - Deployment of HTCondor 10 (all X.509 dependencies removed) - **22/23 Winter shutdown**
  - Monitor/support continuing development of storage token infrastructure
- ❖ Support upcoming Data and Throughput Challenges
  - Testing/support of migration to Tape REST API (**23/24 winter shutdown?**)
- ❖ Technical investigations into modern tools for http based caching (eg. Varnish)
- ❖ Migrate Site Admin knowledgebase and DevOps support documentation off the SDCC website to the new US ATLAS website





## 2.3.5: Risk Analysis

- ❖ Risk #1: Site performance at risk due to continued reliance on EOL software
  - Technical or manpower issues have caused delays in updating some critical software infrastructure
  - Mitigation: Continued support for facilitating upgrade path
  - Costs: Potential for downtimes or service degradation
  - Probability: Low
  - Owners: Rind, Luehring



## 2.3.5: Summary and Conclusion

- ❖ Continued smooth facility operations as we begin Run 3
- ❖ Several cross-cutting development activities ongoing with 2.3.2 and 2.4
- ❖ Steadily improving engagement and cooperation with the central ADC team...
  - ...but *additional effort is needed here* to improve further
  - Shift of some effort for Rind back to 2.3.5, but not enough to account for loss of personnel



# Backup Slides



## 2.3.5: FY22 Milestone Status

- ❖ **MS55. Multi-site Kubernetes platform** (A. Vartapetian, R. Gardner)
  - During 2021 and beginning of 2022 there was no hardware available to setup Kubernetes (K8s) cluster, and configure it for ATLAS production.
  - Small amount of servers were made available mid-February 2022, to start an initial Kubernetes cluster, before the main hardware arrival from UTA\_SWT2 site decommissioning. K8s setup on startup cluster completed successfully, with all K8s components showing healthy.
  - An ATLAS queue in CRIC, SWT2\_CPB\_K8S, was created, and later used to submit ATLAS grid test jobs. A Harvester service account in the K8s cluster was created, setting up configuration to communicate with the ATLAS Harvester. Cluster instance was setup on the ATLAS Harvester side.
  - There were a couple of weeks of downtime during April during which all the K8s servers were updated with additional memory. Also part of the hardware from decommissioned UTA\_SWT2 cluster was racked in CPB and added to the K8s cluster.
  - There were networking issues during May-June, affecting all the worker nodes in the cluster, eventually resolved by replacement of the network switch.
  - There was a configuration issue with the K8s cluster internal network, provided by the Calico networking service, affecting one calico pod on the master node. The configuration was updated, and all the networking components are now showing healthy.
  - Cluster currently has ~600 cores (to be increased once production starts). ATLAS grid test jobs can reach the cluster, and are distributed to the workers, but remain stuck in a waiting state, trying to bring a container. Seems to be a network configuration issue for worker nodes on the private network.



# Proposed ADC Liaison Role

US ATLAS is looking for someone to support distributed computing operations and service development across the five data centers in the US Cloud. The essential duties and responsibilities of this position include:

- ❖ Serving as an active liaison between US site admins and ATLAS Distributed Computing (ADC) operations at CERN, i.e. facilitating communication in both directions on topics such as site or ADC problems and concerns, software changes, and facility or ADC maintenance windows
- ❖ Serving as a US ATLAS point of contact for the Open Science Grid (OSG) and LHC Grid teams, as well as other projects providing site software
- ❖ Assist sites in maintaining up-to-date information in various configuration registries (e.g. the ATLAS/WLCG Computing Resource Information Catalog (CRIC) and OSG Topology)
- ❖ Proactive monitoring, triaging and troubleshooting of grid operations in the US Cloud
- ❖ Contributing to documentation of site monitoring and troubleshooting as a resource for local site admins
- ❖ Providing occasional assistance to site admins with software upgrades and site-specific issues
- ❖ Participating actively in support and development of US Federated Operations activities such as the SLATE Frontier-Squid operations project
- ❖ Attending and reporting at regular operational and software development meetings; providing occasional admin training sessions
- ❖ Assisting management in planning the evolution of US computing operations