



WBS 2.3.2 Tier 2 Operations

Fred Luehring



2.3.2 Introduction

- ❖ No information was received in over a month from the BU team at NET2.
 - In some slides, I leave blank space where the NET2 information would be presented.
 - Fruitful discussions were held with new members of the NET2 team from the University of Massachusetts during SW&C workshop at CERN two weeks ago.
- ❖ I did these slides in a non-standard format because of the decision to let each site present individually.



2.3.2 FTE Summary

- ❖ The individual sites will present their FTEs.
 - The FTEs are largely unchanged with these exceptions:
 - MWT2 has added 0.25 FTE IU (Dambik) and moved 0.67 from UIUC (Haney) to UC (Golnaraghi).
 - SWT2 had 0.5 FTE on CPB retire and then return (Sosebee).

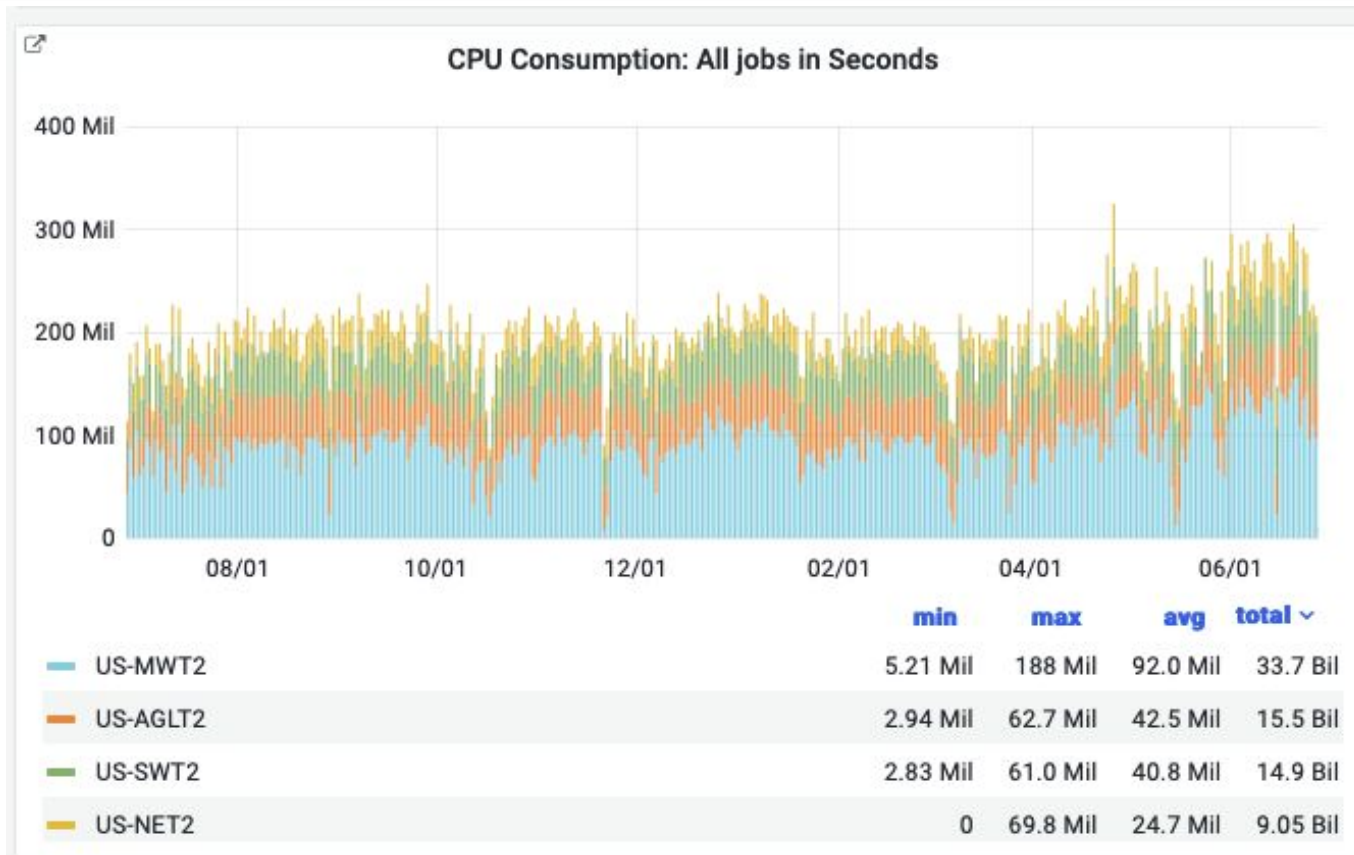


2.3.2: Overview of Activities

- ❖ The tier 2 complex produced A LOT of data processing at a very reasonable cost.
- ❖ The sites worked hard and most are ready for run 3 data taking:
 - All sites have put online all servers that have been received.
 - Dell is struggling and compute servers are still on backorder.
 - Most sites are generally are ready or close to ready with software and service updates.
- ❖ Several sites were instrumental in debugging/validating new service versions:
 - OSG 3.6/Gratia
 - WLCG Token Usage
 - dCache 6.2.x and 7.2.x
 - XRootD 5.3.x
 - ALRB Versions



2.3.2 A LOT of Data



- ❖ The plot is the CPU seconds/day over the last year.
- ❖ The sites also transferred ~7 PB of data out and ~9 PB of data in each week with 95%-98% efficiency for the year. They also transferred ~250 PB internally each week.



2.3.2 Hardware Purchases

Site	FY22 \$ Spent	Slots Added	Slots Total	HS06 Added	HS06 Total	TB Added	TB Total
AGT2	\$374,000	2,784	18,064	47,000	223,948	0	12,000
MWT2	\$1,137,000	4,320	41,776	74,952	589,592	5,441	15,466
NET2			16,968		212,773		8,776
SWT2		4,608	22,136	90,811	339,615	3,200	13,982
Total	\$1,511,000	11,712	98,944	212,763	1,365,928	8,641	50,224

❖ Notes:

- AGLT2 totals incl. 29 R6525 (dual EPYC 7413) shipped/not received.
- NET2 has received CEPH storage and will retire their old GPFS soon.
- SWT2 is receiving their FY21 purchases now but they are included.

❖ Full details can be found in these documents:

- [The v60 Jun 22 & WLCG-v60 tabs of the Facility Capacity Sheet](#)
- [The Procurement Plan for FY22 & FY23 adds descriptive text to the capacity sheet information](#)
- [US ATLAS Equipment Evolution Sheet](#)



2.3.2 Run 3 Readiness

- ❖ AGLT2 is completely ready.
- ❖ MWT2 is ready.
 - Partly updated to HTCondor 9.0.13 (will finish updating this week.)
- ❖ NET2 has not been communication recently but I believe:
 - They still need to update from OSG 3.4 to OSG 3.6.
 - They are using an unsupported version of HTCondor-CE
 - They still are not dual stacked for IPV6.
 - They are using an unsupported version queueing system (SGE 2011)
 - They need to remove their old GPFS storage (aging hardware)
- ❖ SWT2 OU has not put into service servers that they bought for a new GK (received January) and squid (receive March).
 - Horst is has been traveling for most of past 3 weeks. This week he left for a couple more weeks of travel.
 - I only learned that OU would not be updated last Wednesday.
 - OU is still at OSG 3.4 & an unsupported version of HTCondor-CE
 - The old GK does not support IPV6 (the rest of site does).



2.3.2 Run 3 Readiness (continued)

- ❖ SWT2 UTA has progressed well getting ready for run 3 but they are down to the wire for some updates:
 - Still using LSM instead of Rucio mover.
 - Need to finish setting up IPV6.
 - Still waiting to receive compute servers ordered in last September.
 - Need to install networking gear received last year...

- ❖ Details about Run 3 readiness and about what SW versions are in use for each site is using can be found at:
 - [The run 3 readiness sheet](#)
 - [The services sheet \(shows SW versions for each site\)](#)



2.3.2 Worries

- ❖ Some sites resist updating to current software versions.
 - These sites can be difficult to communicate with.
- ❖ We are scaling up and up and up which leads to:
 - Staff being under significant stress.
 - Gatekeepers and other control servers being overloaded.
 - Possible network capacity issues
 - Though this is mitigated by the recent network upgrades.
- ❖ The supply chain issues make it difficult to predict delivery.
 - In September, all 4 sites ordered the compute servers with a CPU that Dell said was available.
 - The 1U configuration server were delivered starting in January with the last being delivered this week (~9 months of waiting).
 - A chassis configuration was delivered in less than 2 months.
- ❖ Staff leaving is a really worry because industry is offering high paying work from home jobs.



2.3.2: Summary and Conclusion

- ❖ The Tier 2 sites remain an economical way to process data on a massive scale.
- ❖ Great strides have been made getting ready for Run 3
- ❖ A 10% cut can be handled by buying less equipment.
 - Reducing personnel levels would hurt the facility much more.
 - We can't afford to lose the historical knowledge and experience of the current team.
 - As the facility grows we are really stretched thin.
- ❖ However the progress is not uniform across all sites.
 - It is a real challenge to solve this discrepancy.

- ❖ All following slides are template ones and are not altered.



2.x.y: Ongoing M&O Activities

- ❖ List ongoing M&O activities e.g.
 - Monitor performance
 - Respond to support requests
 - Provide analytics platform
 - Coordinate with OSG/CCE/whatever

Highlight changes in
personnel/effort level

For each activity indicate who is working on it at which
fraction



2.x.y: FY21 Milestone Status

❖ MS1. Complete FY22 scrubbing

On time - Aug 2022 - Aug 2022

- Describe milestone goals and work status
- Highlight risks
- Explain delays and extra costs if any
 - Describe impact

- Describe corrective actions (e.g. increase FTE, buy more disk, etc)

❖ MS2. Meet FY21 Disk Pledge

➤

❖ Use as many slides as you need

Refer to [milestones spreadsheet](#)

Include both open and completed milestones. If completed add original milestone date, when it was actually completed, why it was delayed and the impact of completing it late.

Put technical details in backup. Highlight who is working on what



2.x.y: New Activities

- ❖ Describe new developments/ideas/collaborations not yet ready to be captured as milestones
 - e.g. US ATLAS/Google collaboration
 - e.g. fabulous new idea to use GPU

One slide only,
details in backup



2.x.y: Risk Analysis

- ❖ Risk #1: HEP-CCE fails to deliver a PPS
 - Describe potential problems
 - e.g. can't run ATLAS code across different accelerators
 - Describe mitigation strategy and its effects
 - e.g. support only CPU, can not use exascale machines
 - Quantify extra-costs and delays as relevant
 - Guesstimate the probability that the risk event will happen
 - Identify risk owner (or owners if it's multiple people)
- ❖ Risk # 2: etc

Refer to the [US ATLAS Risk Register](#)

This may help you organize your thinking about priorities, and the benefits of R&D projects.



2.x.y: Priorities

- ❖ One slide to describe priorities for FY22 and beyond
 - Each FY22 priority should be captured by a new milestone (next slide)
- ❖ Connect your priorities to
 - Risk analysis
 - CDR priorities
 - DOE/NSF priorities
 - Facility requirements
 - ...
- ❖ For each priority estimate FTEs/budget needed and # of years.
- ❖ For each priority identify person in charge.

One slide only,
details in backup

For example: a priority could be “save RAM by running our workflows using athenaMT”



2.x.y: Upcoming FY22 Milestones

- ❖ Same format as slide 4 (FY 20 Milestones)

For example: a FY22 milestone could be “validate athenaMT pileup/digitization for production”