WBS 2.3.2.4

Southwest Tier-2 (SWT2)

University of Texas at Arlington
University of Oklahoma

Mark Sosebee (on behalf of the SWT2 Collaboration)

2.3.2.4 FTE Summary

	FY21 US ATLAS Funded	FY21 Other Funding	FY22 US ATLAS Funded	FY22 Other Funding	Job description
Patrick McGuigan	1.0		1.0		UTA sys admin
Mark Sosebee	0.5		0.5		UTA sys admin
Horst Severini		0.5		0.5	OU sys admin
Chris Walker	0.5		0.5		OU sys admin
Kaushik De		0.2		0.1	Management
Total	2.0	0.7	2.0	0.6	

- Other activities:
  - Patrick, Mark, and Chris - none
  - Horst - Oklahoma supercomputer system administrator
  - Kaushik - faculty/US ATLAS management/NSF

SWT2: Summary of Current Capacities

- SWT2 pledged 5.11% of ATLAS requirement for 2022
- CPU deployed:
  - Job slots (threads): 22,136
  - Total HEP-SPEC06: 260,662 (pledged 81,164)
  - Non-dedicated - opportunistic (OU): typically 500-1500 / 10.9 k
- Disk storage deployed:
  - 13.1 PB (pledged 7.3 PB)

- Note: On 4/1/22, UTA SWT2 consolidated to a single campus site, after 18+ years of operating two sites (on Arlington campus & at Ft. Worth).
- Over the past year we retired 1.8 PB of storage (Dell MD3*xxx*-based) in parallel with the commissioning of ~5.4 PB of new storage.

SWT2: Plans & Concerns

- Additional Worker Nodes arriving currently:
  - UTA: 45 Nodes, 2880 Job slots, ~47K HS06
  - OU: 15 Nodes, 1920 Job slots, ~28K HS06
- Current network capacities:
  - UTA: 2 x 40 Gb to the WAN
  - OU: 100 Gb to the WAN
- Potential needs:
  - UPS refresh at UTA in the next year
  - Dedicated chilled water system for the CRAC's at UTA
    - Two sudden campus chilled water outages in past 12 months
  - Networking bandwidth
    - New VPR at UTA - discussions underway soon

SWT2: Plans & Concerns (II)

- ATLAS Run 3 readiness:
  - Fully ready for Run 3 MC (ongoing) and data processing (soon)
    - All systems/services operational, all workflows are supported
- Other activities:
  - Decommissioning of LSM (local site mover) underway
    - LSM maintained for 15 years to optimize storage access
    - Testing rucio mover as a replacement for the LSM
  - IPv6 has been been deployed
    - Fully at OU
    - perfSONAR at UTA ready, SE gateway hosts next
- Power capacity at the rack level is becoming a concern - making plans to efficiently distribute power loads in the machine room
- The LAN networking infrastructure at SWT2_CPB will be upgraded later in summer 2022 for improved throughput and redundancy

SWT2: Unique Characteristics

- Both SWT2 sites (UTA and OU) use XRootD (+ xfs) for the back-end storage (SWT2 is only US ATLAS T1/T2 to do so). Hence we work closely with the developers of XRootD in all phases of deployment & operational aspects of the software.

- Over the past few purchasing cycles we have prioritized storage over CPU, in preparation for Run 3. Going forward we will revisit this question and adjust our mix as needed.
- A Kubernetes (K8) cluster for testing has been deployed. Armen Vartapetian is leading this effort. Currently working with PanDA developers to enable ATLAS workflows on the system.
- Dedicated test queues for pilot developers are provided.

SWT2 Performance

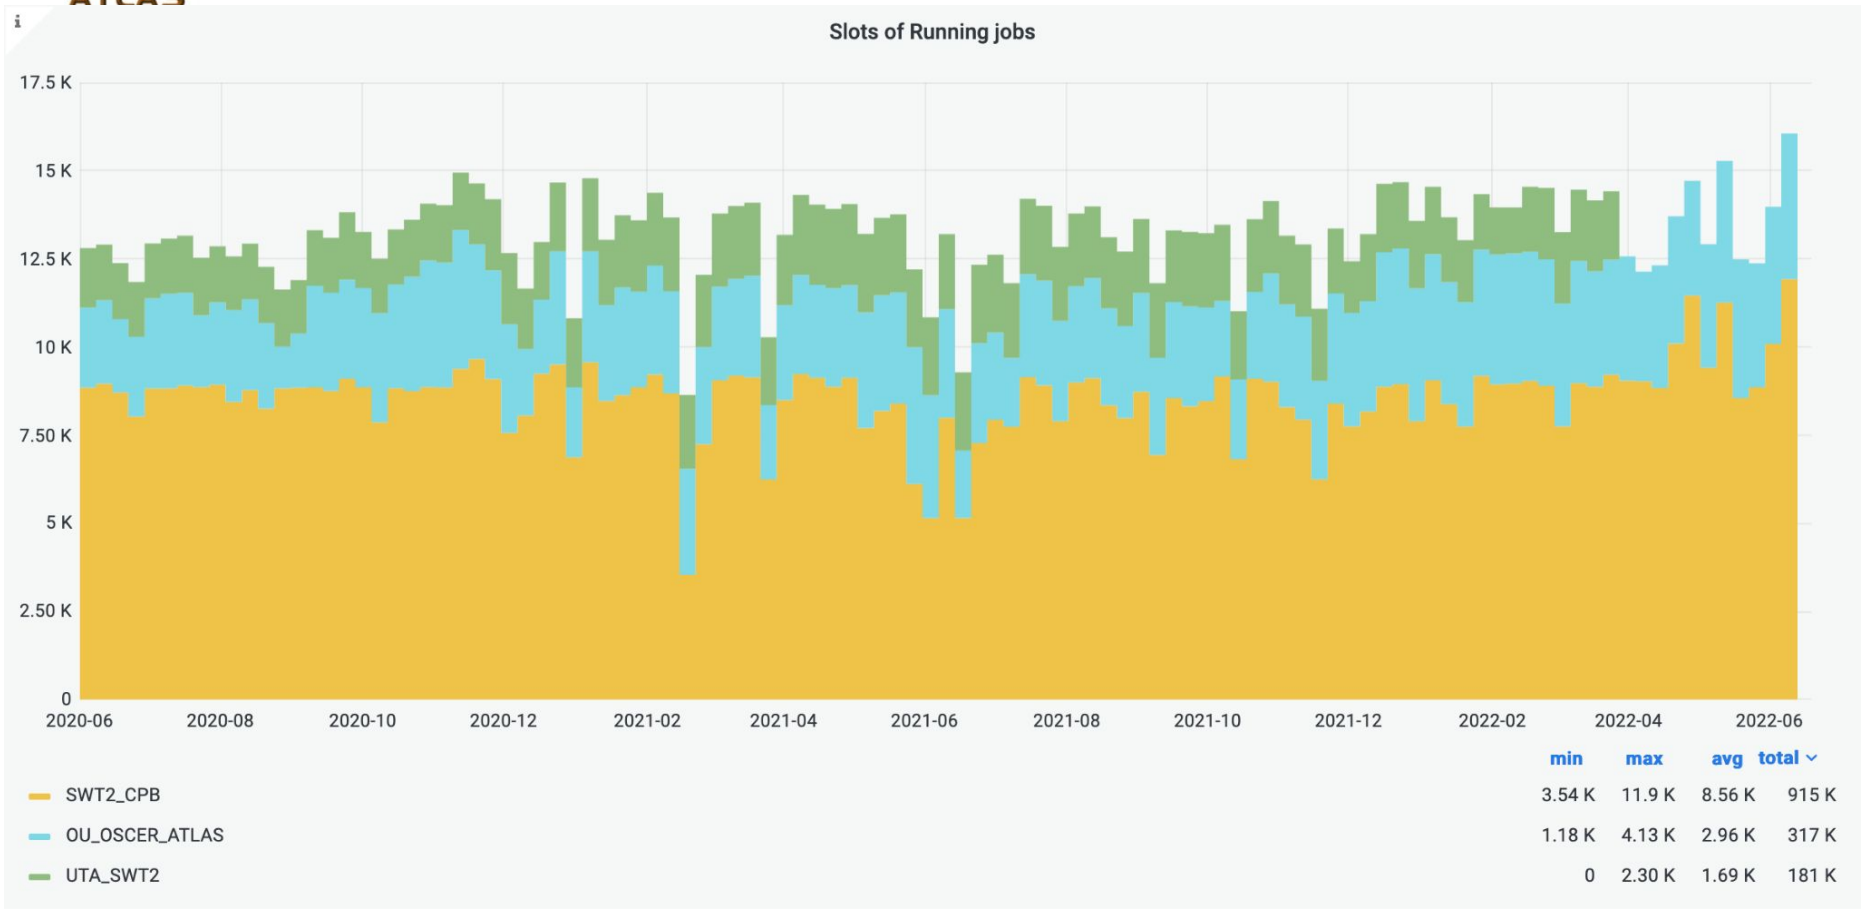

- Steady performance (averaged weekly) for past 24 months
- Old Ft. Worth site (in green) fully migrated
- CPU capacity rising with recent (delayed by supply-chain problems) deliveries

SWT2 Performance (II)

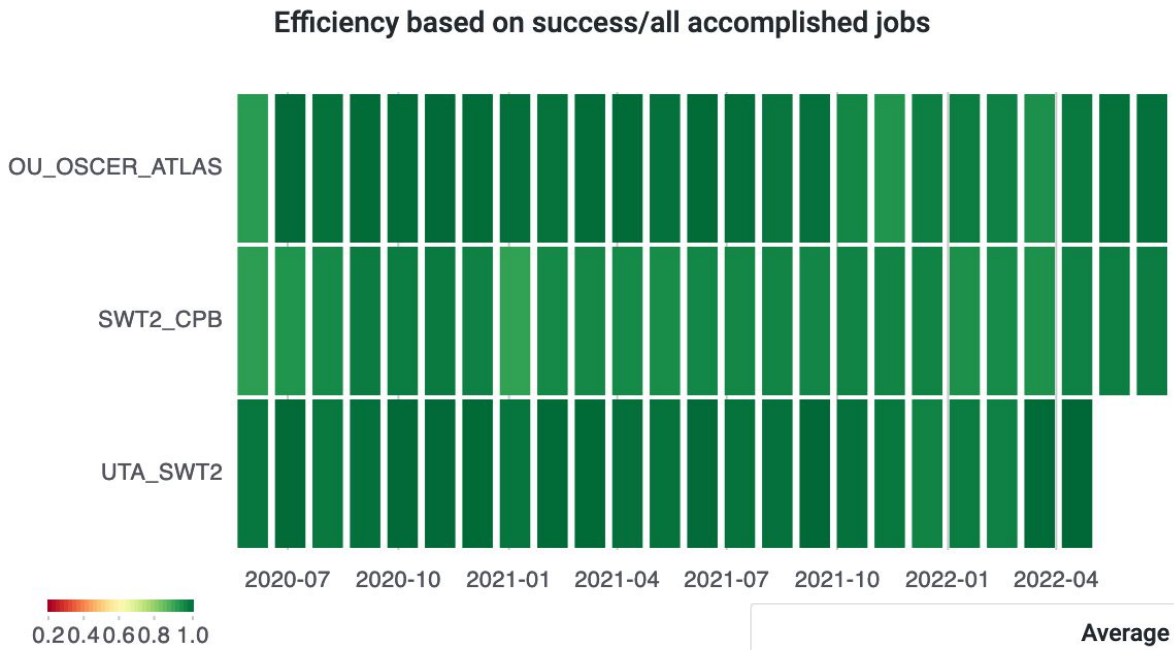

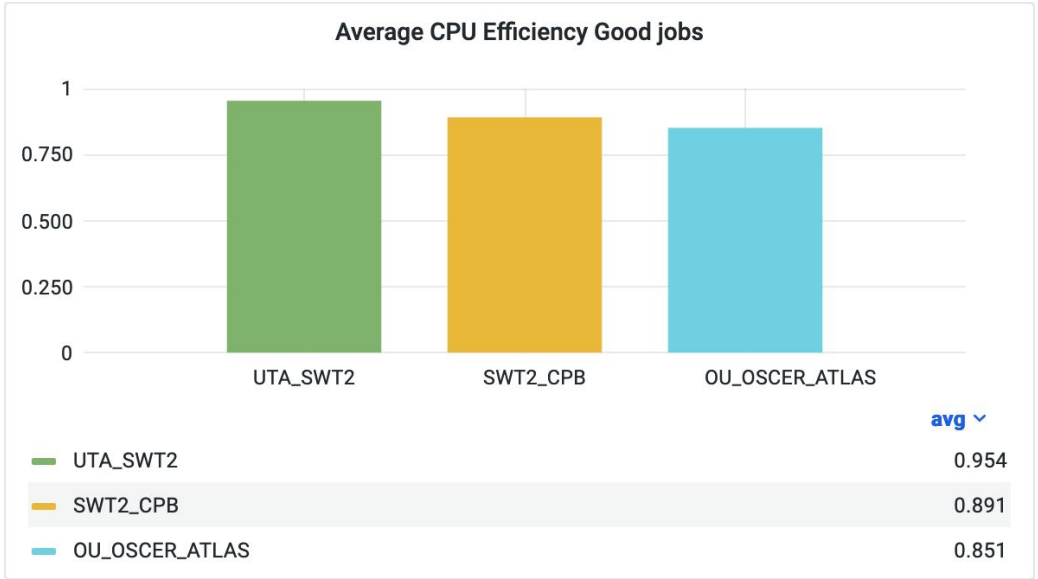

SWT2 within US Context

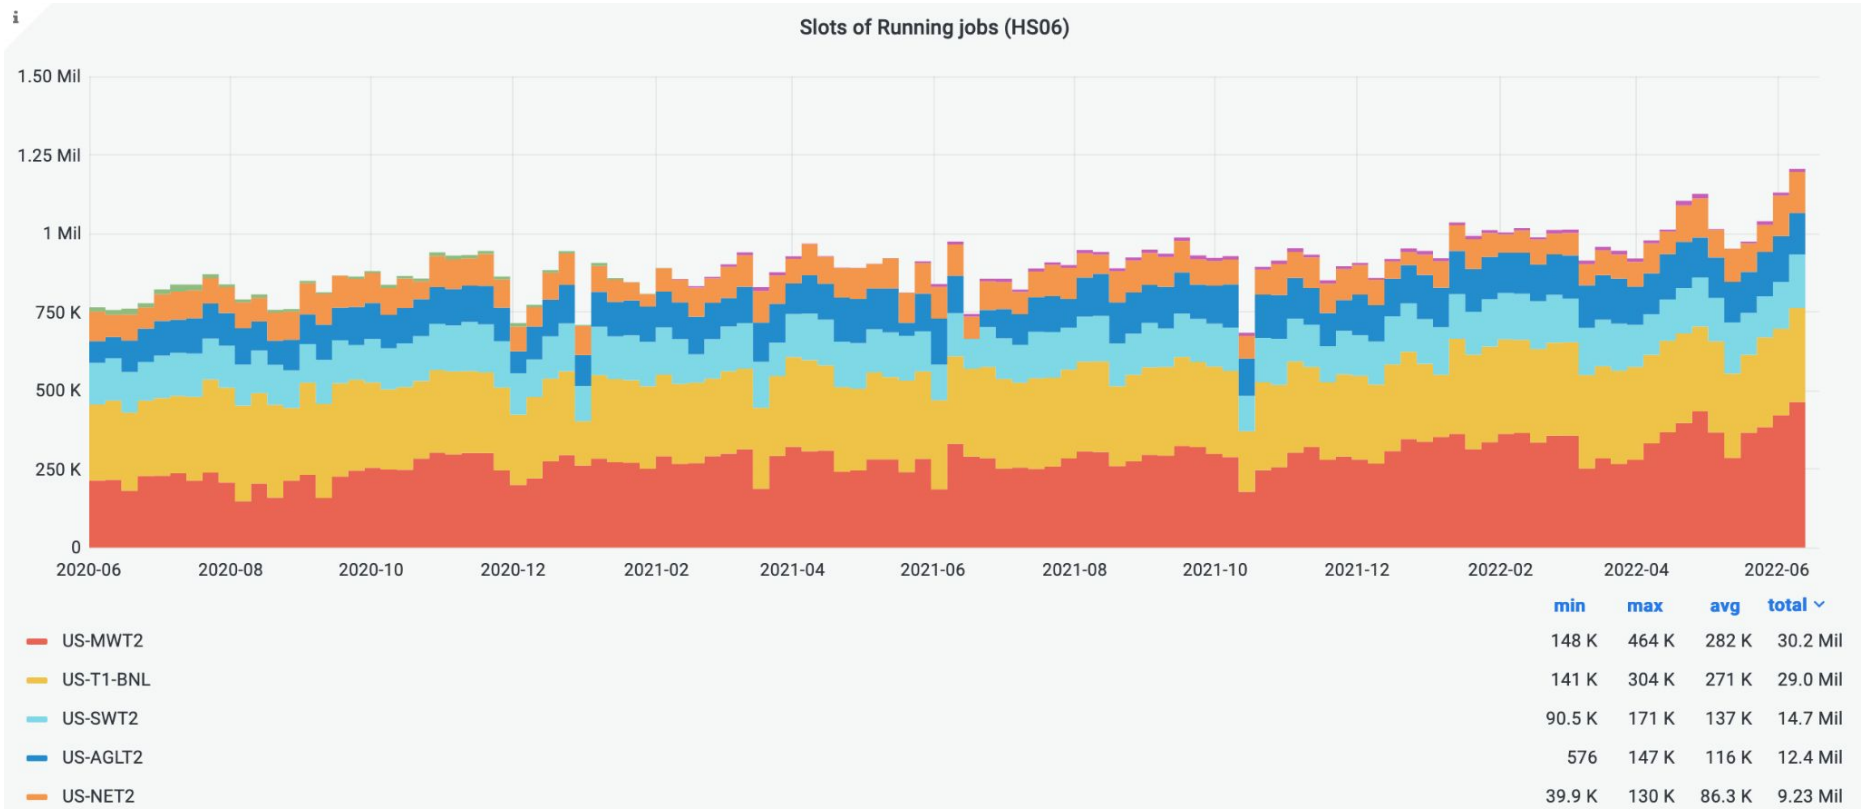

- Delivered CPU at SWT2 in light blue among US ATLAS T1/T2 sites

SWT2 Budget

- FY21 budget:
  - Total $931,305, 43% operating, 57% equipment (large purchases)
- FY22 budget:
  - Total $931,305, 44% operating, 56% equipment
  - Slight increase in operating due to maintenance & service contracts
- Flat FY23 budget request:
  - Request $931,305, expect 44% operating, 56% equipment
- -10% scenario for SWT2 (for one year):
  - Postpone UPS upgrade - hardware failure risk, long outage risk
- -10% scenario for SWT2 (all future years):
  - Buy less storage every year - Run 3 data risk, user analysis risk
  - Reduce equipment 10%, go to 50-50 operating-equipment split - risk ATLAS MC/processing power in the future
  - Overall, we can meet pledge even with equipment cuts, but SWT2 provides high capacity at very low cost, so risk to ATLAS physics