



Monitoring and Metrics for the US-ATLAS Analysis Facilities

Mike Hance (UC Santa Cruz)

Viviana Cavaliere (Brookhaven National Laboratory)

on behalf of the WBS 2.3/5 team



AF Metrics

- ❖ Effort began with a recommendation from US-ATLAS Operations program reviews last year:

Develop metrics and investigate the effectiveness of the T3 analysis facilities and communications tools to identify improvements. Continue to pursue activities to federate the T3 analysis facilities to provide a better user experience and utilization of resources across all T3 centers by the next review.

- ❖ Goal: develop metrics that:
 - Are useful for us to track usage patterns, etc
 - Demonstrate the effectiveness of the facilities to FA's



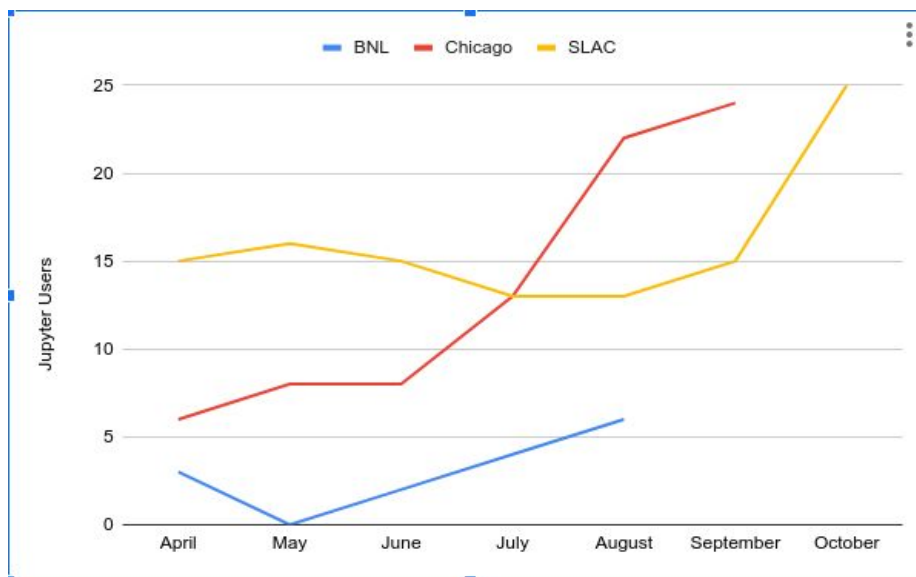

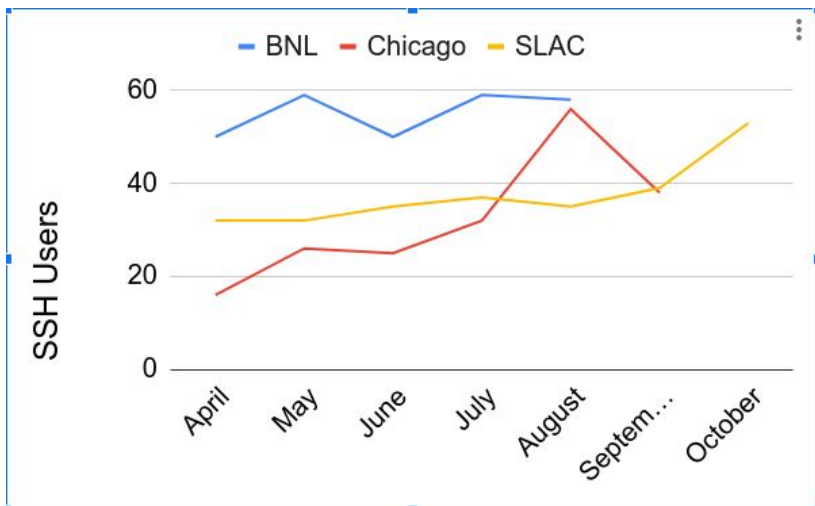
AF Metrics

- ❖ In 2.3/5 meetings we decided to track in “Phase 1”:
 - Number of SSH users
 - Number of Jupyter users
 - Number of users submitting batch jobs
 - Number of batch jobs
 - CPU-hours of batch jobs
- ❖ Originally tracked via a [Google Sheet](#):

	SSH	Jupyter	Batch Users	Batch Jobs	Batch Hours	SSH	Jupyter	Batch Users	Jupyter -Coffea	Batch Jobs	Batch Hours	SSH	Jupyter	Batch Users	Batch Jobs	Batch Hours
	BNL	BNL	BNL	BNL	BNL	Chicago	Chicago	Chicago	Chicago	Chicago	Chicago	SLAC	SLAC	SLAC	SLAC	SLAC
April	50	3	24	762,484	466,170	16	6	9		7,923	2,768	32	15	15	77,475	155,221
May	59	0	24	1,237,176	285,377	26	8	12	4	74372	19552	32	16	12	111,392	285,971
June	50	2	23	827,260	603,795	25	8	18	9	42130	87298	35	15	12	61,372	252,226
July	59	4	22	529,048	177,406	32	13	15	6	181,106	83341	37	13	8	40,154	79,550
August	58	6	19	665,345	468,378	56	22	25	17	73903	22148	35	13	9	47,921	152,338
September						38	24	22	8	149,178	411,720	39	15	12	50,583	115,817
October												53	25	14	117,027	294,184




AF Metrics





AF Metrics

Now using a [kibana dashboard](#) (see more in Ilija's talk).





Next Steps

- ❖ Still some things that need to be captured/fixed for Kibana:
 - Batch users, jobs, hours
 - SLAC ssh user counting is missing login nodes
- ❖ In the meantime, update Google Sheet with October, November, December numbers
- ❖ We have discussed other metrics in the past, agreed to push to “Phase 2”. Can we start adding these?
 - Held batch jobs
 - Storage: disk utilization, files transferred, etc
 - Network, latency
 - GPU usage/availability
 - Others I’m forgetting ([notes](#))



Summary

- ❖ Close to satisfying review comments on AF's
 - Few more small changes to make, but Phase-1 is nearly complete
- ❖ Would like to revisit a few more metrics that will make dashboard even more useful
 - Info that users will find helpful: queue times, available disk, etc.
- ❖ Compiling these metrics was not easy – thanks to the site developers who put real time and energy into this.
- ❖ Thanks also to Ilija for developing the dashboard at UChicago