



# Tier 1 performance

Doug Benjamin, Eric Lancon

ATLAS Pre-Scrubbing Review - June 27, 2022




# Tier-1 WLCG MoU

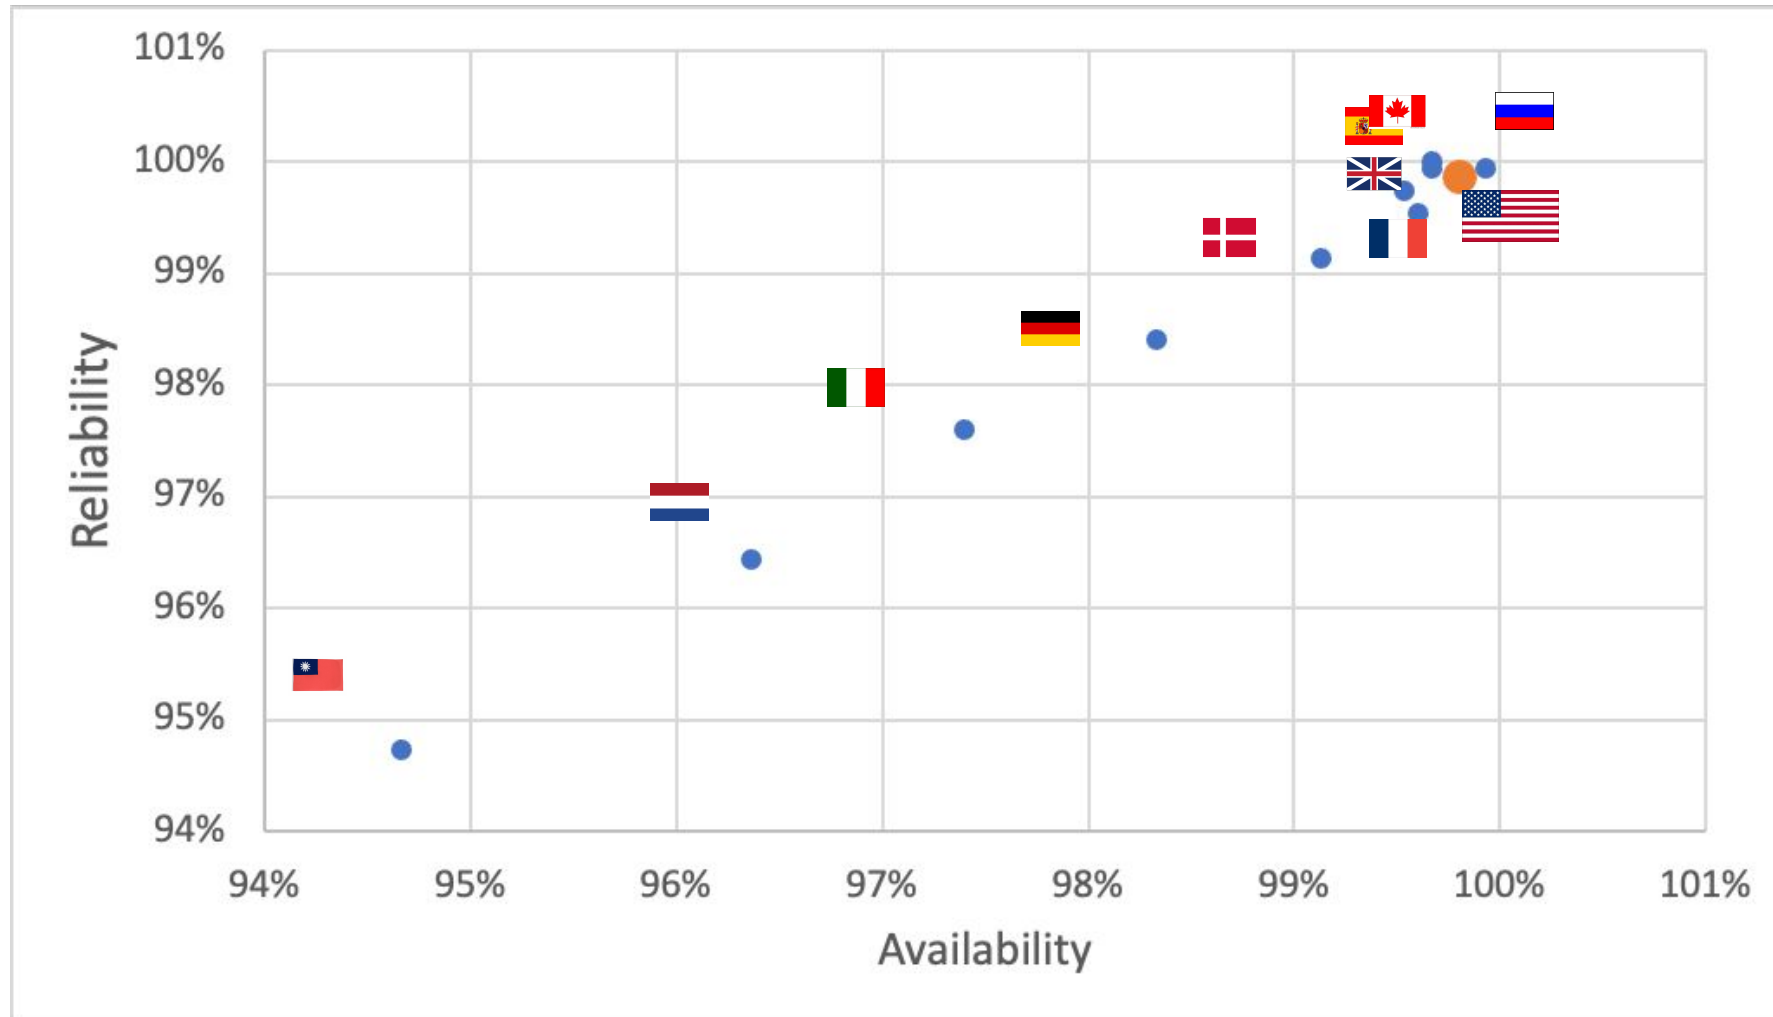
Note  
difference  
between  
accelerator  
ON and OFF

<i>Service</i>	<i>Maximum delay in responding to operational problems</i>			<i>Average availability<sup>6</sup> measured on an annual basis</i>	
	Service interruption	Degradation of the capacity of the service by more than 50%	Degradation of the capacity of the service by more than 20%	During accelerator operation	At all other times
Acceptance of data from the Tier-0 Centre during accelerator operation	12 hours	12 hours	24 hours	99%	n/a
Networking service to the Tier-0 Centre during accelerator operation	12 hours	24 hours	48 hours	98%	n/a
Data-intensive analysis services, including networking to Tier-0, Tier-1 Centres outwith accelerator operation	24 hours	48 hours	48 hours	n/a	98%
All other services <sup>7</sup> – prime service hours <sup>10</sup>	2 hour	2 hour	4 hours	98%	98%
All other services <sup>7</sup> – outwith prime service hours <sup>10</sup>	24 hours	48 hours	48 hours	97%	97%


# Reliability vs Availability all Tier-1s



# Reliability vs Availability all Tier-1s




# Number of CPU slots used over last year




# CPU usage vs pledges

CPU usage  
 2021: 104% pledges  
 2022: 105% pledges

US-T1-BNL: CPU Used Plot




US-T1-BNL: CPU Used Plot



# CPU delivered by ATLAS Tier-1s

CPU delivered at Tier-1 [May 2021-May 2022]




US Tier-1 targets 23% of ATLAS approved requests at Tier-1s

US Tier-1 delivers 20% (some Tier-1s deliver well above pledges)

# CPU by activities

CPU consumption by activity



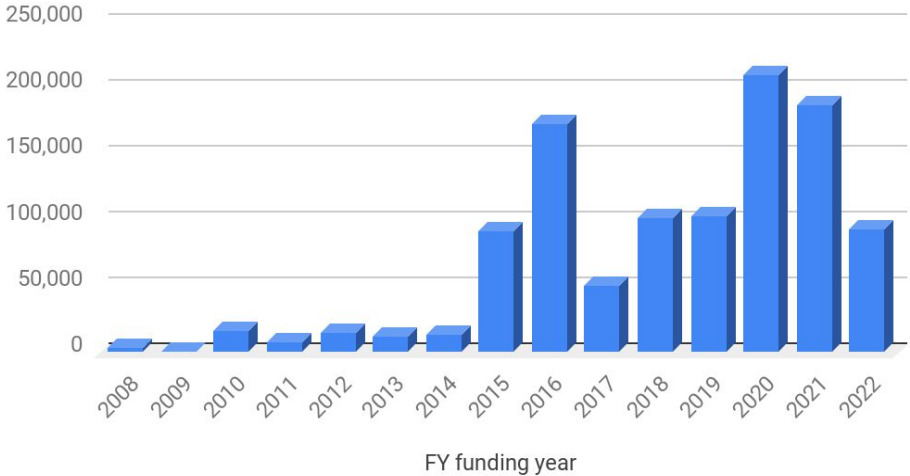
75% IO intensive tasks

**BNL compared to other T1s**  
20% more user & group analysis  
40% more reconstruction  
50% less Simulation

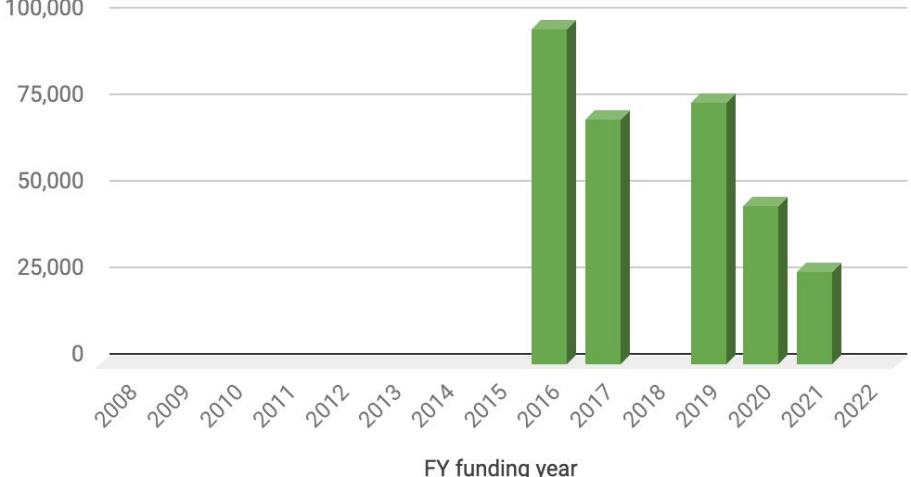


# CPU age at US T1 & T2s


T2s: Online CPU capacity by year of purchase




T1 Online CPU capacity by year of purchase



T2s CPU: Fraction of equipment older than N years

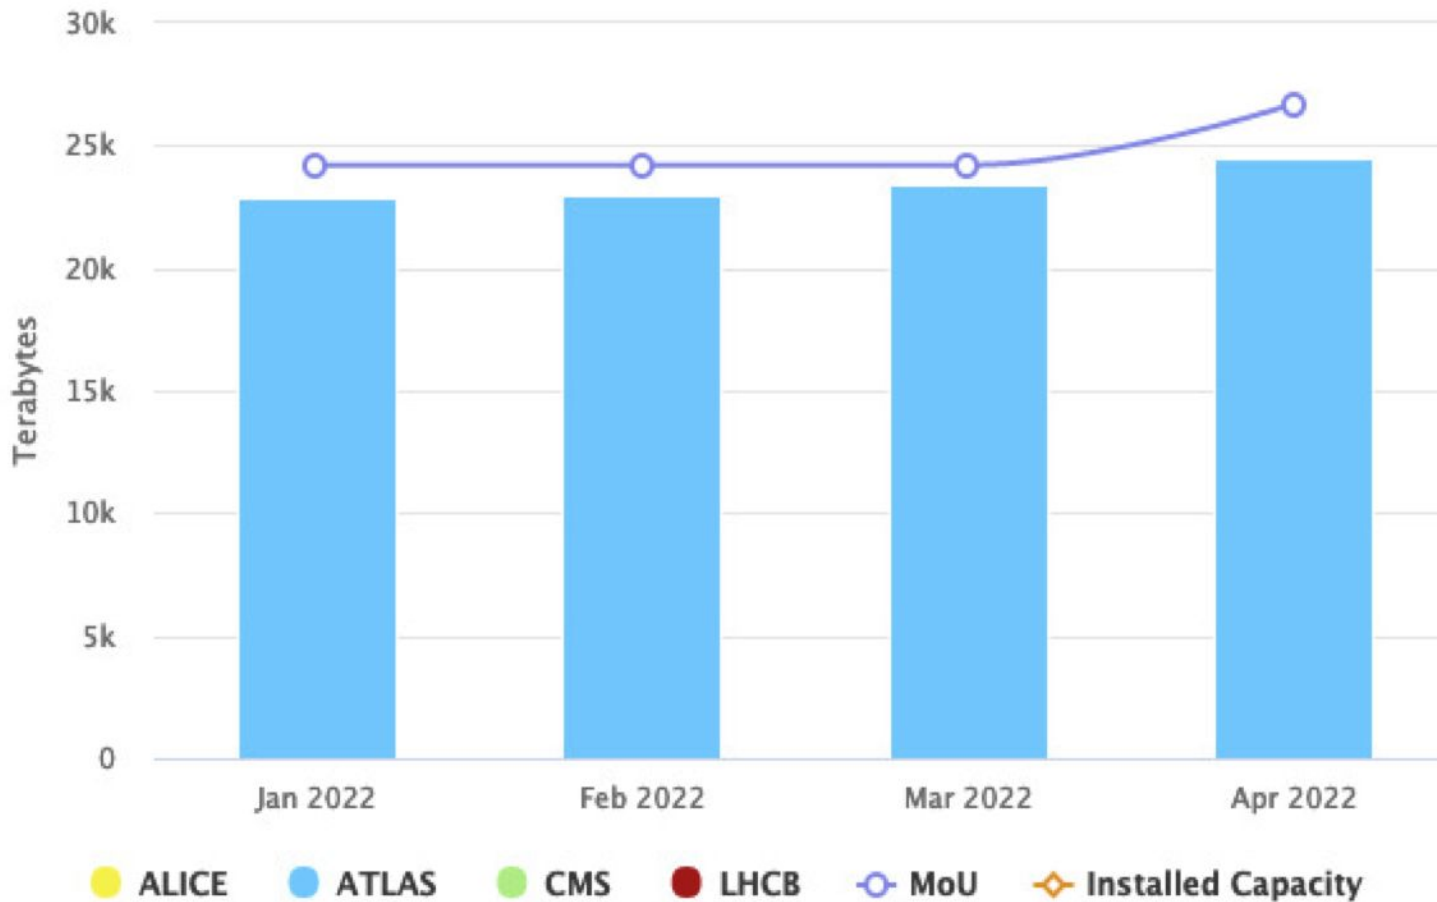


T1 CPU: Fraction of equipment older than N years




# Data Disk storage usage


US-T1-BNL: Disk Used Plot



US-T1-BNL: Disk Used Plot




# Disk (pledged) available at ATLAS Tier-1s




# BNL Tier 1 Datadisk end points.

2 disk copies



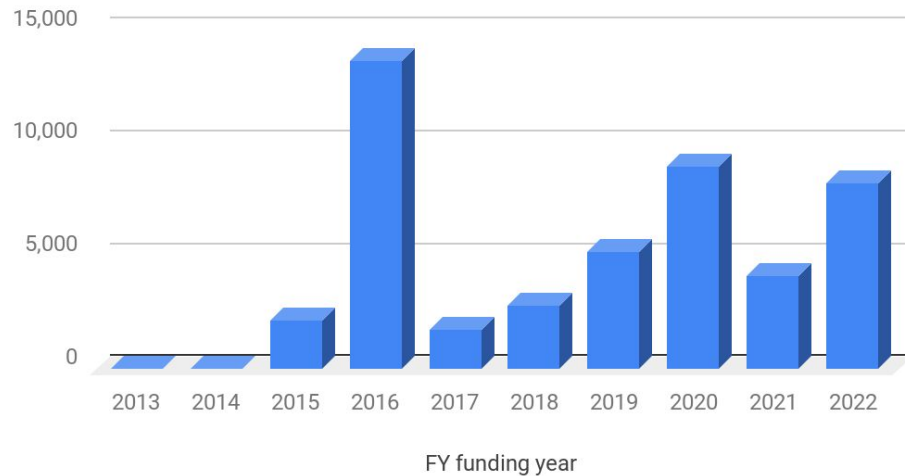
1 disk copy




The two endpoints combined fulfil the disk pledge  
Plan to merge the 2 endpoints

# Disk storage age at US T1 & T2s


T2s: Online Disk usable capacity by year of purchase



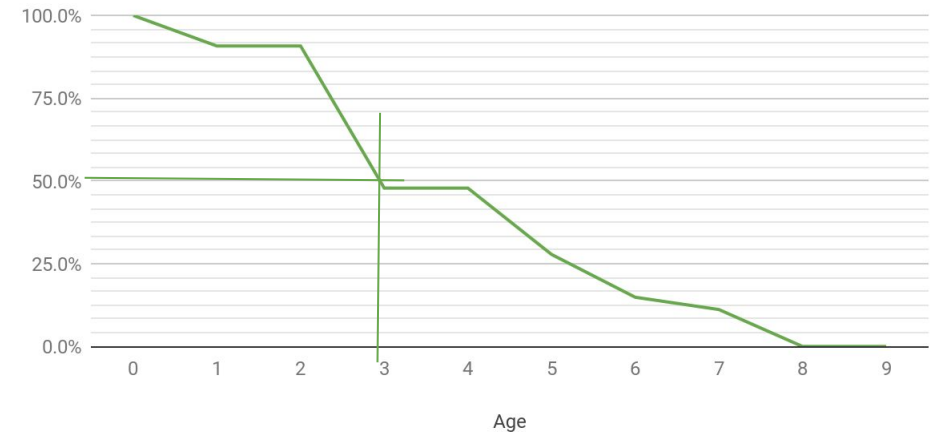
T1: Online Disk usable capacity by year of purchase




T2s: Disk: Fraction of equipment older than N years




T1 Disk: Fraction of equipment older than N years



# Tape Capacity at BNL




# Data Transfer Rates during most recent tape challenge



Tier-1 sites	DT test (tape write)		A-DT test (tape read)		
	*Target (GB/s)	Real rate (GB/s)	Target rate (GB/s)	Max rate (GB/s)	*Stable rate (GB/s)
BNL	3.5	>8	1.9	6.8	5
CC-IN2P3	3.5	Up to 4	1.2	6.9	2.6
CNAF	3.5	>4	0.8	3.0	1.9
KIT	3.5	Up to 3.5	1.0	3.0	1.6
NL-T1	3.5	Up to 7	0.6	1.5	1.2
PIC	3.5	>6	0.3	2.1	1.2
RAL	3.5	Up to 3.5	1.2	2.1	1.7
*TRIUMF	3.5	Up to 4	0.8	2.8	1.6

- \* The DT target rate refers to the peak rate during a run (main stream RAW data)
- \* stable rate refers to a rate sustained for >5 hours.
- \* TRIUMF DT result was from a re-test after the tape challenge

# BNL WAN Traffic in PB/month





# BNL WAN Traffic growth for the past 5 yrs

