ISS operations and plans

ALICE T1-T2 workshop 2019
Bucharest, Romania

Ionel STAN, Adrian SEVCENCO

INSTITUTE OF SPACE SCIENCE

P.O. Box: MG-23, RO 077125 Magurele, ROMANIA

http://www.spacescience.ro

Table of contents

Overview

ISS, ISS_LCG – Institute Of Space Science (ISS)

Overview – Computing Infrastructure

New ISS Computing Infrastructure

➔ Designed for high density computing (Hot Aisle, InRow cooling)

➔ Scalable solution for future investments

➔ UPS Power : 48 kVA (with N+1 redundancy power units)

➔ Cooling capacity : 80 kW installed (2N capacity redundancy)

Overview

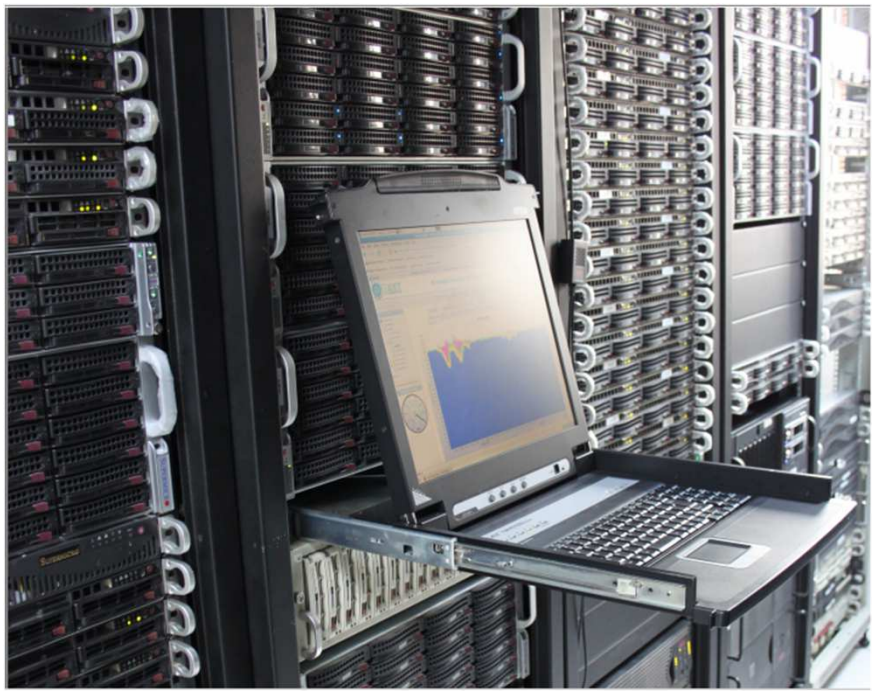

Cluster	Number of servers	Core's
ISS-ALICE	50	1216
RO-13-ISS	8	128
PlanckGrid	16	256
RoSpaceGrid	40	648
Total	114	2248

Overview - ISS Alice Site

Computing resources

- 25x 8 cores (Xeon E5410 – Harpertown) / 16 GB ram
- 4x 24 cores (Opteron 6172) / 64 GB ram

- new computing resources purchased in 2016 and 2017
  - 456 cores (Xeon E5-2650 v4 Broadwell, 14 nm, 2.2 GHz base freq) - 19 nodes
  - memory 5.3 GB/core - DDR4 2400 MHz ECC
  - 2 x 10 Gb network/server; 4x40 QSFP uplinks
  - enclosure expandable to 28 nodes

Storage resources

- 6 storage servers
- storage capacity upgraded from 220 to 460 TB (2017), 835 TB (mid of 2018) , 1.16 PB (end of 2018)

Table of contents

Site capabilities - ISS

HARDWARE AND TOPOLOGY OF COMPUTING FACILITY

- Our hardware is mainly comprised of SuperMicro machines that were chosen for the great resource density/price ratio. For computing nodes we use Twin servers and Blade servers which give us very good densities and for the storage we use servers with 24, 36 drives and JBOD cases with 45 drives in 4U of rack space.
- Generic schematic of ISS computing facility :

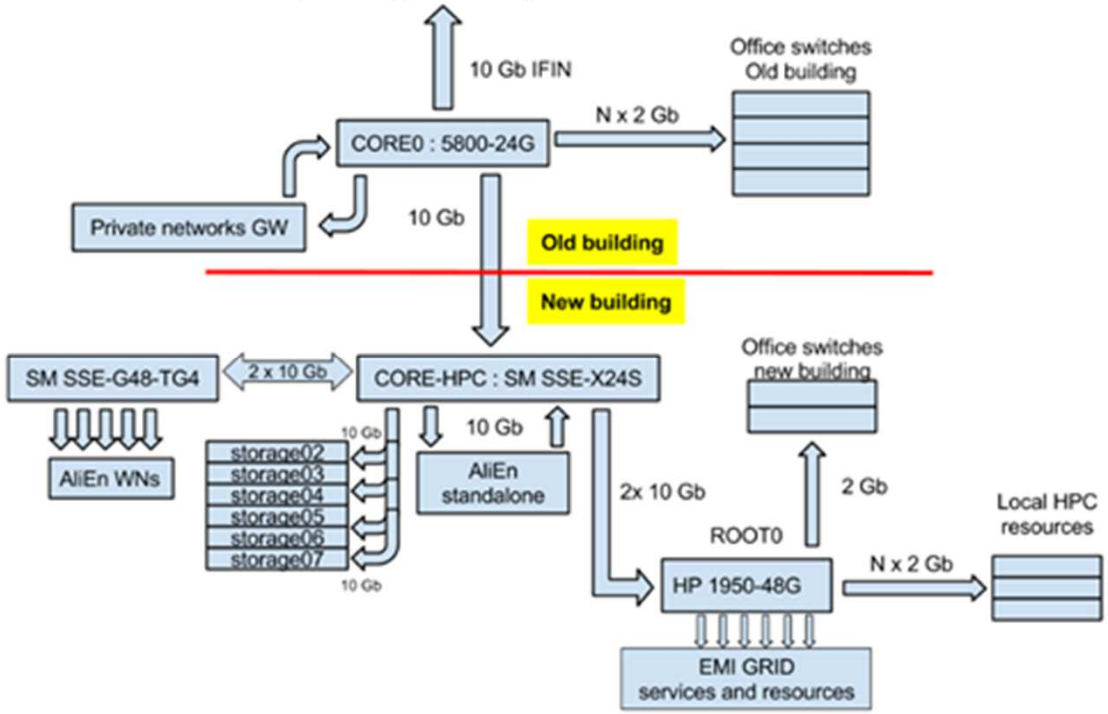

Site capabilities - ISS

HARDWARE AND TOPOLOGY OF COMPUTING FACILITY

- The AliEn cluster has at his core a 10 Gbps aggregating switch which is connected to the top-of-rack switch of the computing nodes. In the aggregating switch are connected the interfaces of the storage node, a topology which give a high bandwidth connection between worker nodes and storage with very little oversubscribing.

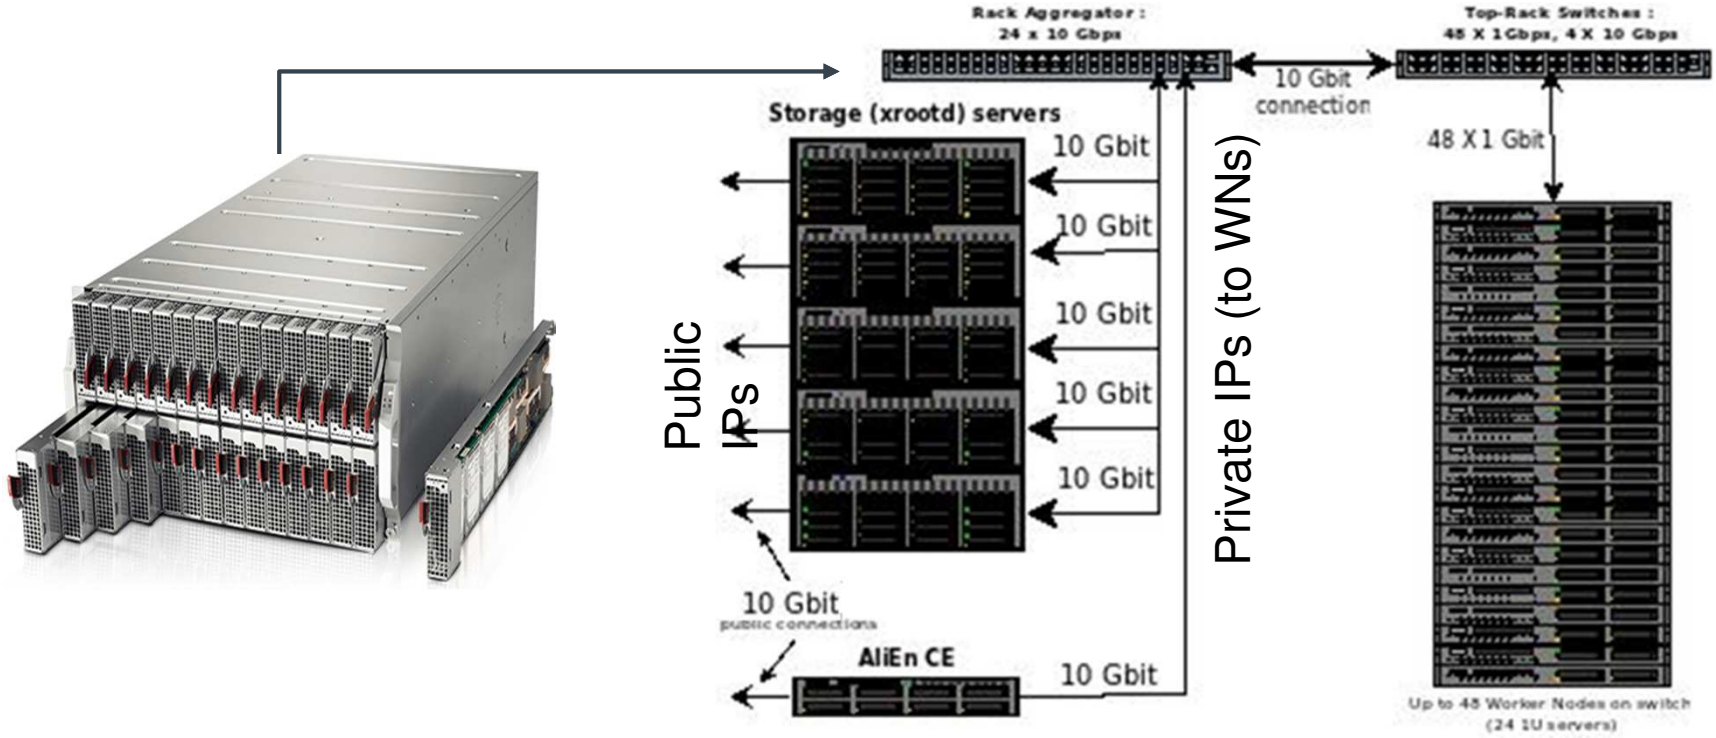

Table of contents

Site status – running jobs profile

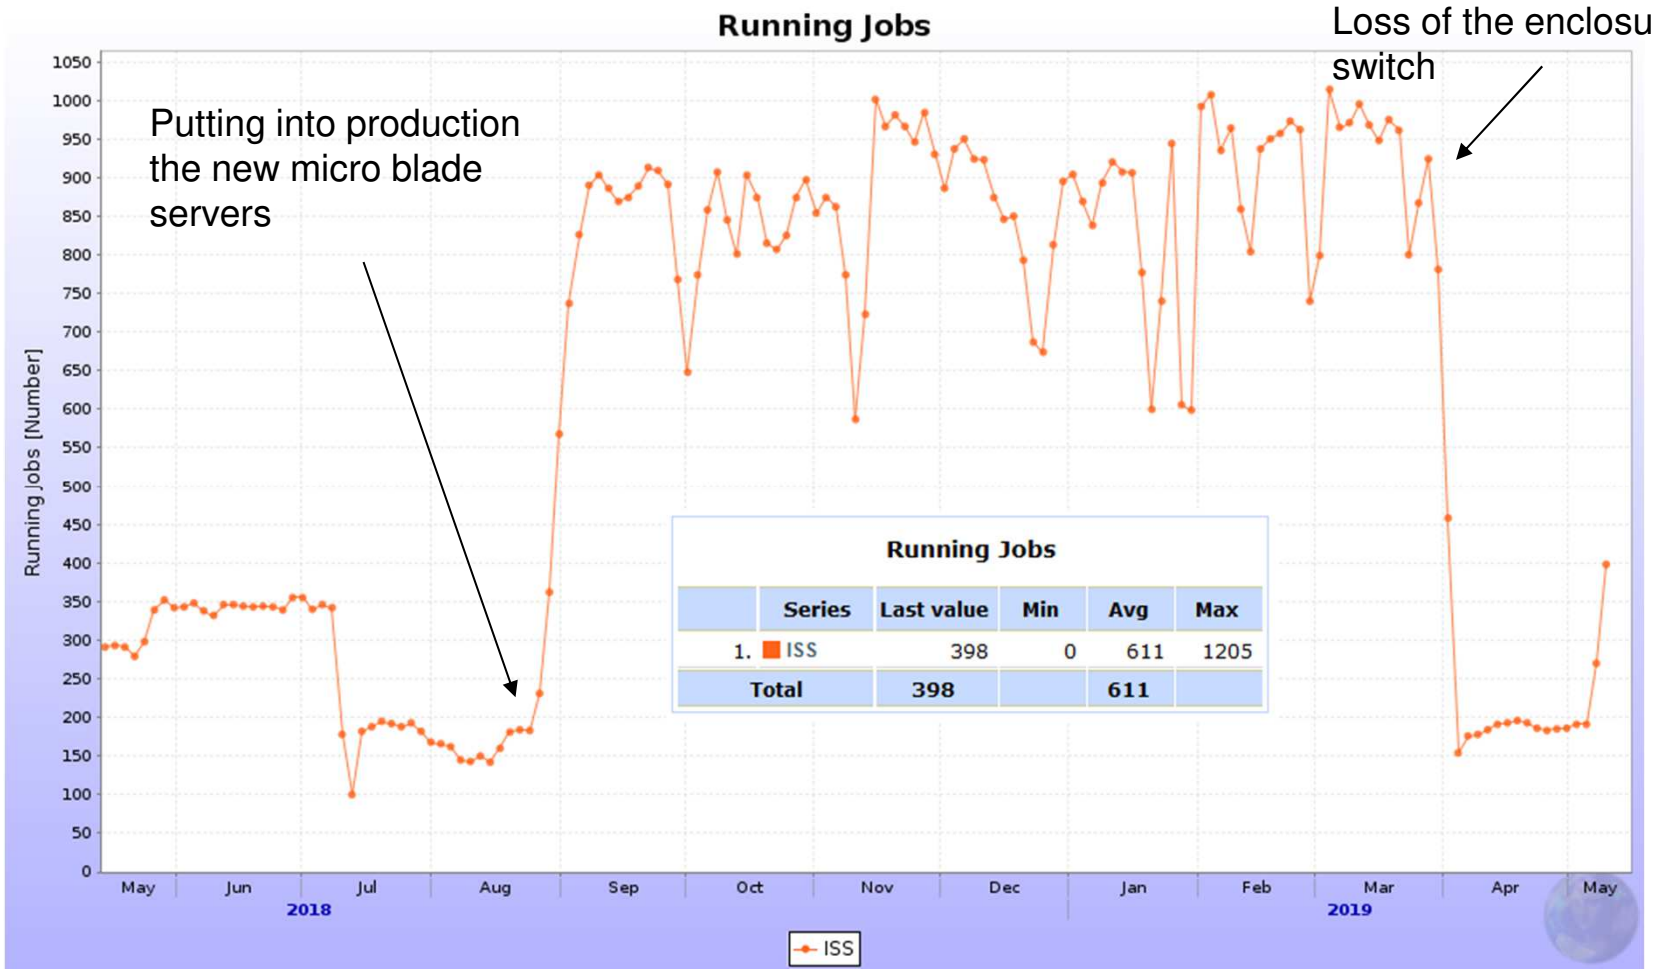

Site status - Job Efficiency

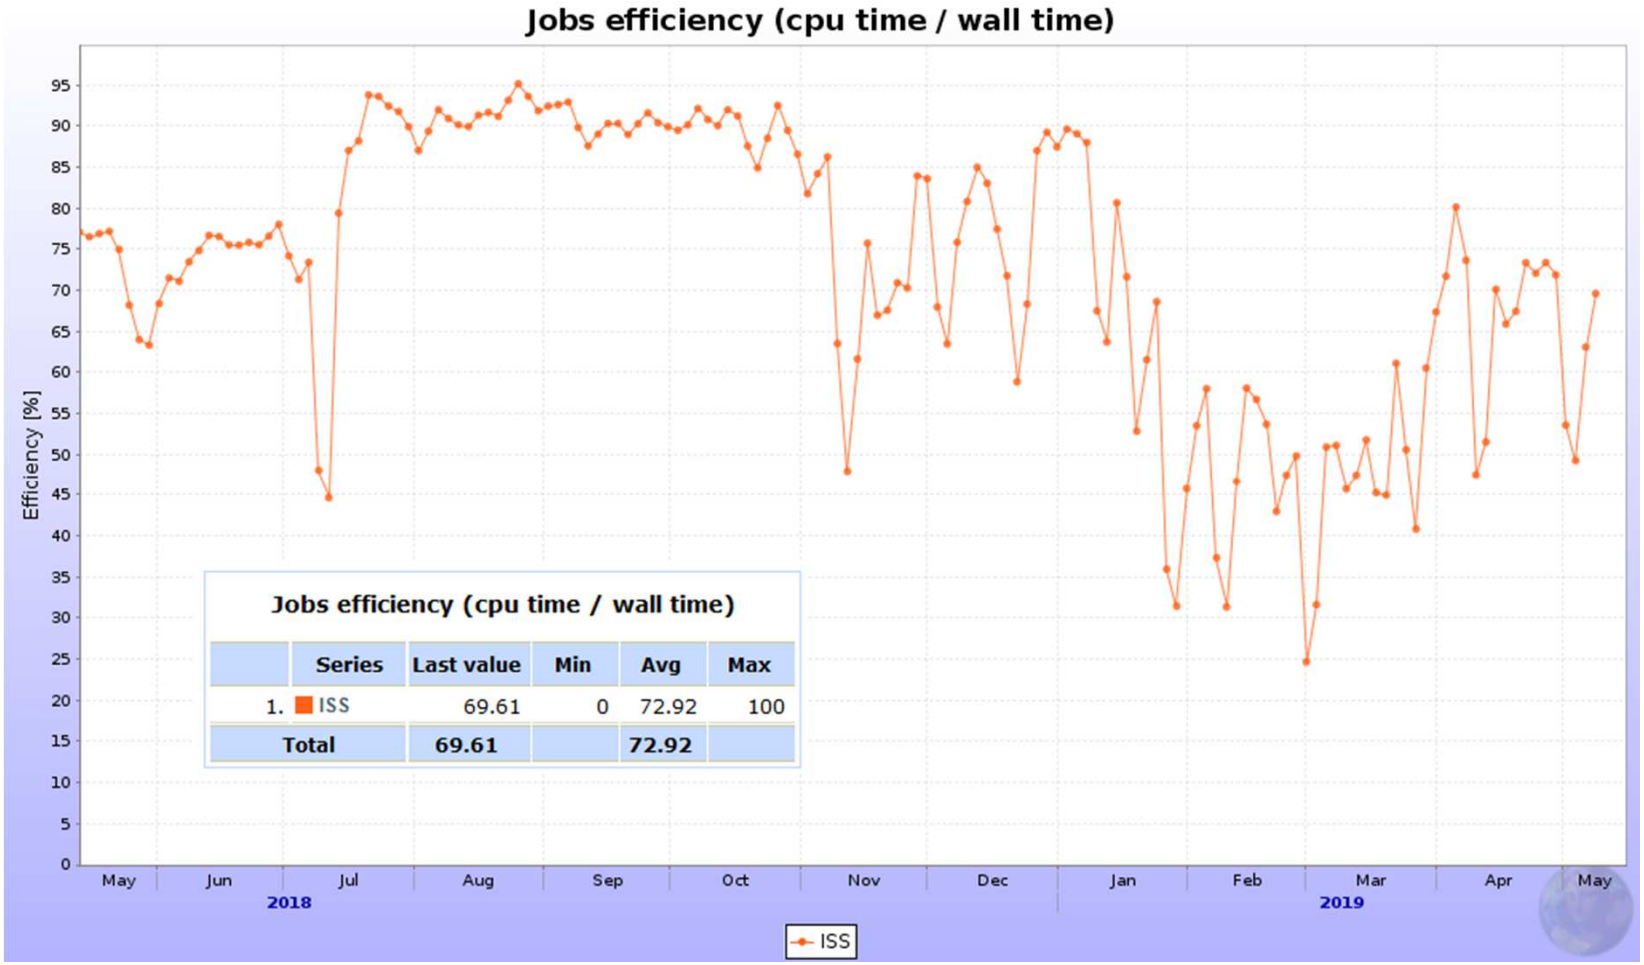

Site status – ISS computing contribution

- over 1 M done jobs
- 4.66 M CPU hours
- 11.48 M kSI2k hours

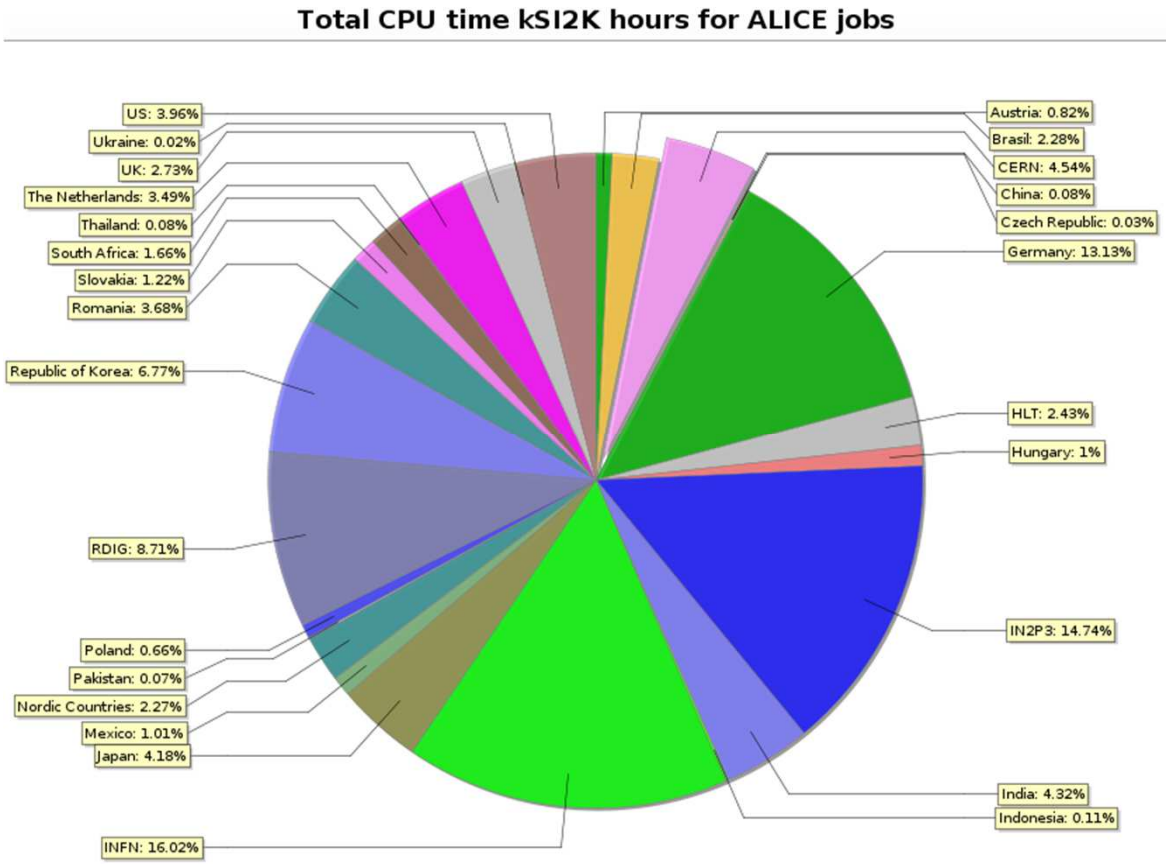

Site status – SE Availability

Disk storage elements

ISS													
	AliEn SE		Catalogue statistics						Storage-provided information				
SE Name	AliEn name	Tier	Size	Used	Free	Usage	No. of files	Type	Size	Used	Free	Usage	Version
1. ISS - FILE	ALICE::ISS::FILE	2	1.163 PB	1020 TB	170.4 TB	85.7%	19,502,984	FILE	1.163 PB	1013 TB	178 TB	85.06%	Xrootd v4.9.1
Total			**1.163 PB**	**1020 TB**	**170.4 TB**		**19,502,984**		**1.163 PB**	**1013 TB**	**178 TB**		

AliEn SEs availability for writing

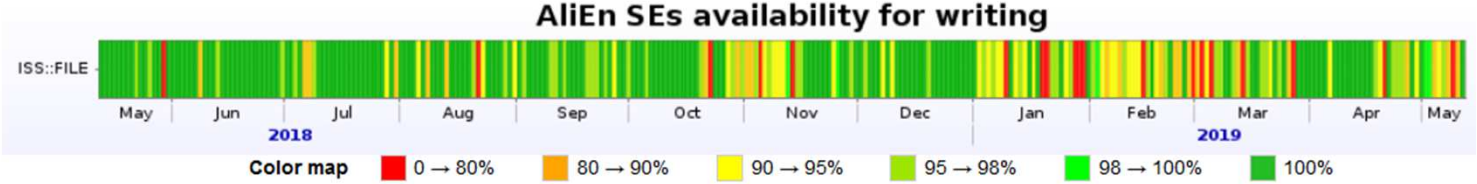

Statistics

Link name	Data		Individual results of writing tests			Overall
	Starts	Ends	Successful	Failed	Success ratio	Availability
ISS::FILE	12 May 2018 10:09	12 May 2019 10:54	8297	410	95.29%	95.32%

AliEn SEs availability for reading

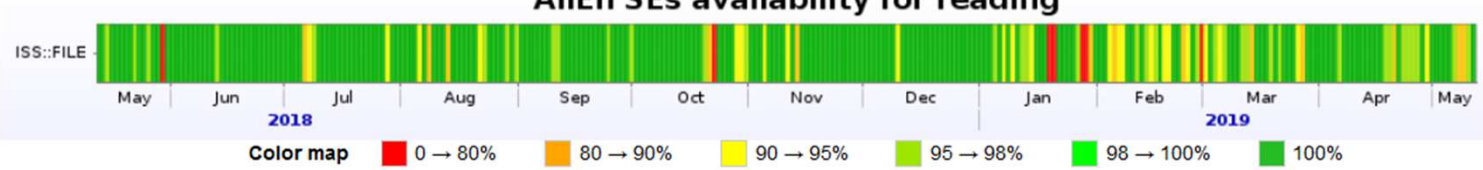

Statistics

Link name	Data		Individual results of reading tests			Overall
	Starts	Ends	Successful	Failed	Success ratio	Availability
ISS::FILE	12 May 2018 10:09	12 May 2019 10:54	8616	228	97.42%	97.40%

Site status - Aggregated network traffic per SE

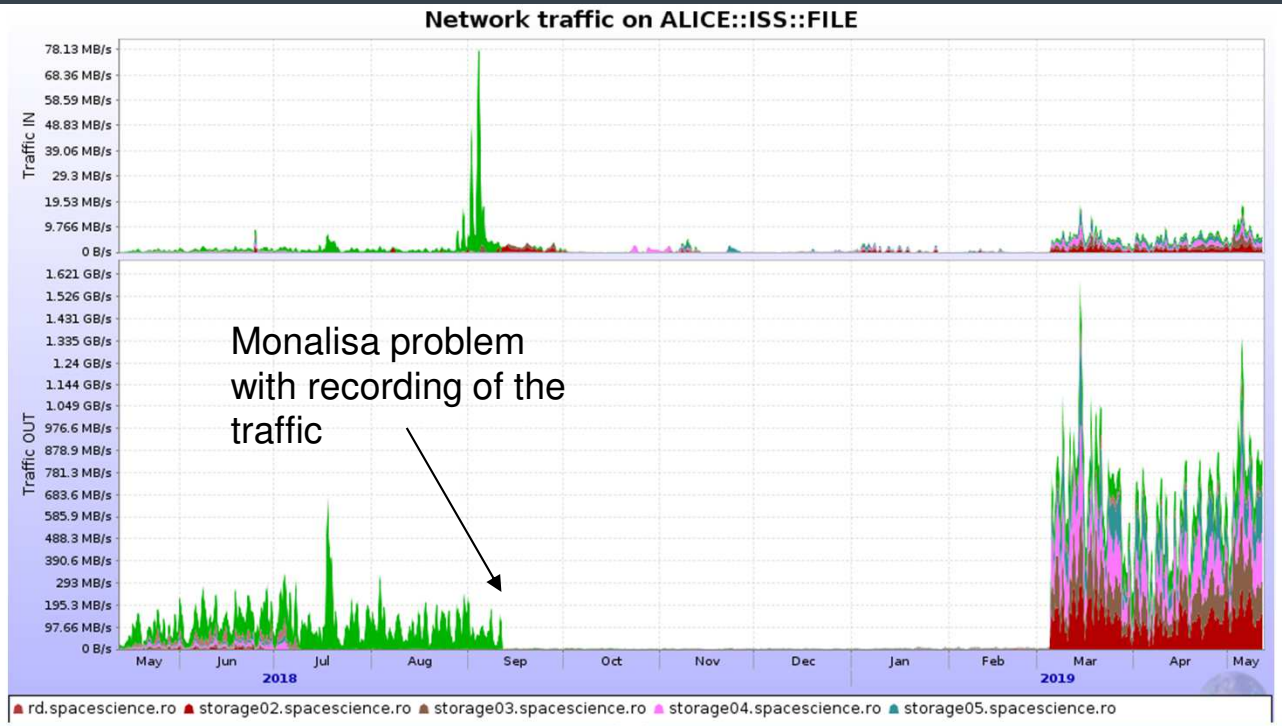

Traffic IN

	Series	Last value	Min	Avg	Max	Total
1.	rd.spacescience.ro	41.58 KB/s	0 B/s	31.93 KB/s	6.658 MB/s	958.2 GB
2.	storage02.spacescience.ro	1.275 MB/s	5.146 B/s	372.2 KB/s	62.87 MB/s	10.58 TB
3.	storage03.spacescience.ro	1.183 MB/s	5.406 B/s	358.4 KB/s	47.61 MB/s	10.5 TB
4.	storage04.spacescience.ro	1.51 MB/s	0 B/s	338.2 KB/s	128.7 MB/s	9.912 TB
5.	storage05.spacescience.ro	1.196 MB/s	0 B/s	329.4 KB/s	706.5 MB/s	9.652 TB
6.	storage06.spacescience.ro	151.3 KB/s	0.373 B/s	82.24 KB/s	104.1 MB/s	2.41 TB
7.	storage07.spacescience.ro	762.8 KB/s	35.53 B/s	1022 KB/s	1.129 GB/s	29.97 TB
	Total	**6.098 MB/s**		**2.475 MB/s**		**73.96 TB**

Traffic OUT

	Series	Last value	Min	Avg	Max	Total
1.	rd.spacescience.ro	10.38 KB/s	0 B/s	10.78 KB/s	349.1 KB/s	323.6 GB
2.	storage02.spacescience.ro	166 MB/s	1.062 B/s	29.42 MB/s	1.03 GB/s	856.1 TB
3.	storage03.spacescience.ro	158.3 MB/s	1.063 B/s	28.54 MB/s	1.285 GB/s	856.4 TB
4.	storage04.spacescience.ro	206.5 MB/s	0 B/s	29.84 MB/s	1.075 GB/s	895.6 TB
5.	storage05.spacescience.ro	178.6 MB/s	0 B/s	24.14 MB/s	1.097 GB/s	724.5 TB
6.	storage06.spacescience.ro	24.31 MB/s	0.053 B/s	8.906 MB/s	957.9 MB/s	267.3 TB
7.	storage07.spacescience.ro	114.8 MB/s	1.57 B/s	54.81 MB/s	1.509 GB/s	1.606 PB
	Total	**848.5 MB/s**		**175.7 MB/s**		**5.122 PB**

Site status - Aggregated network traffic per SE

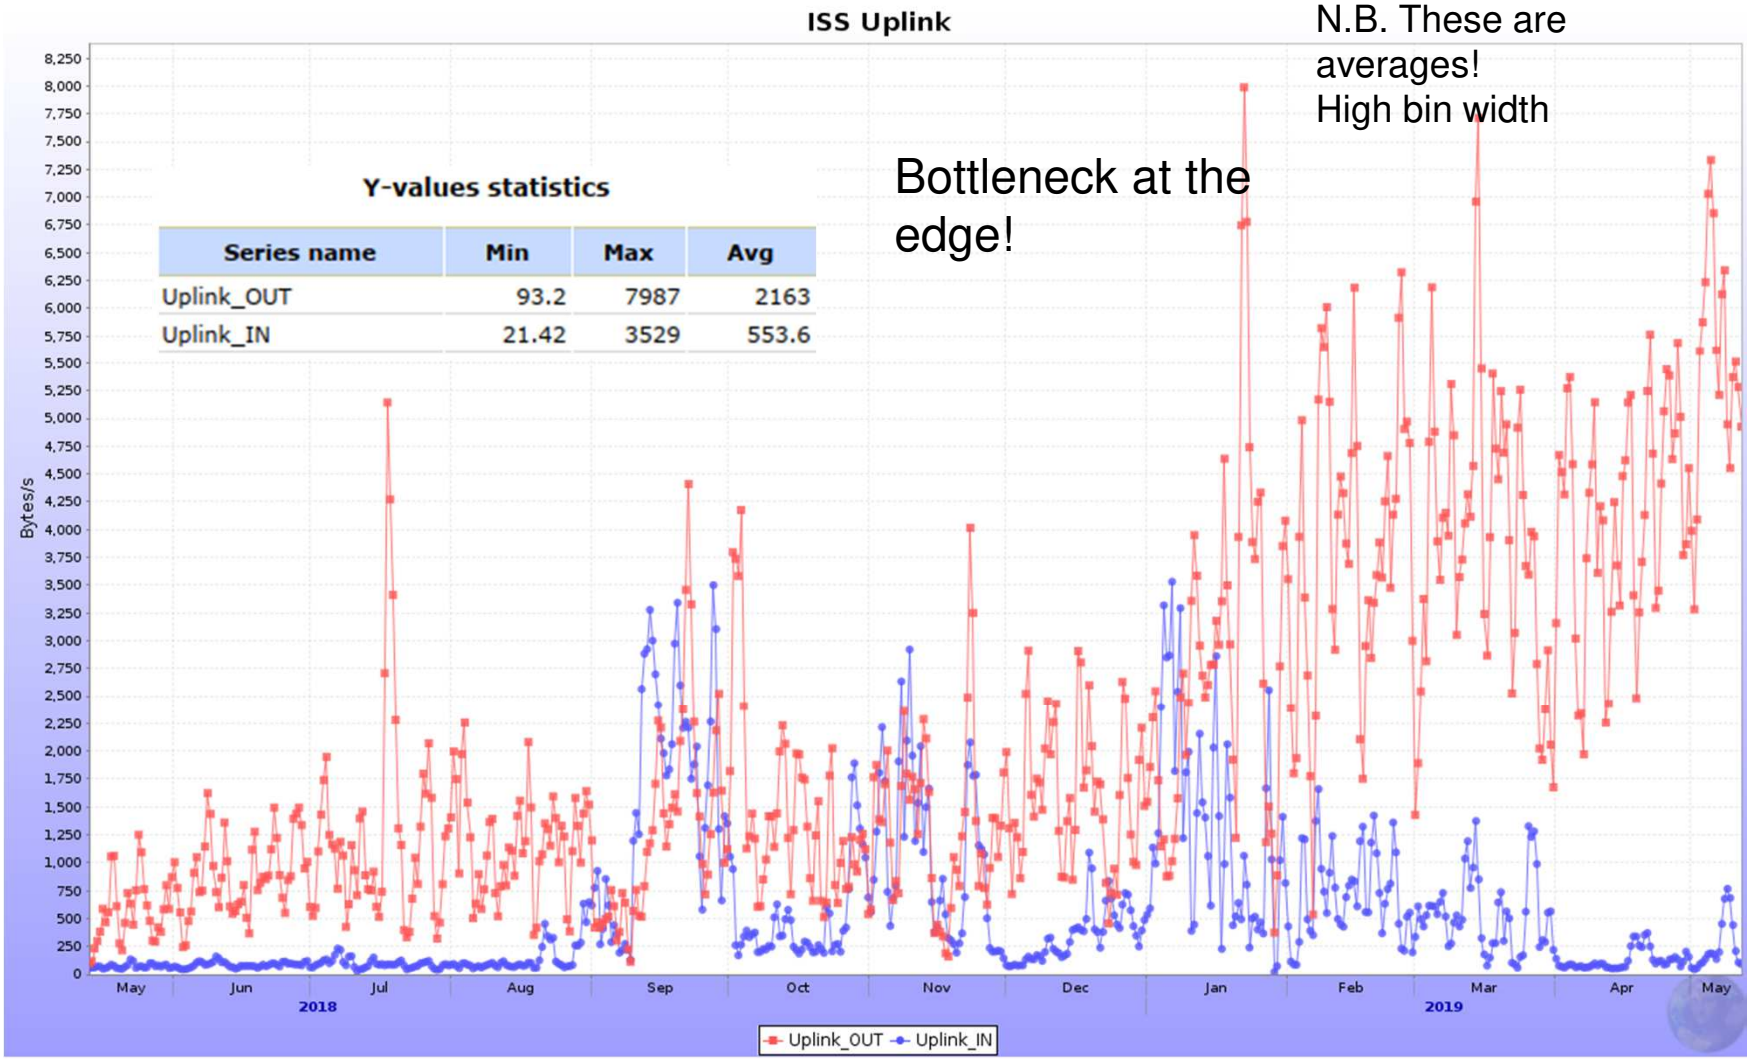

Bottleneck at the edge!

N.B. These are averages!
High bin width

Site status - Aggregated network traffic per SE

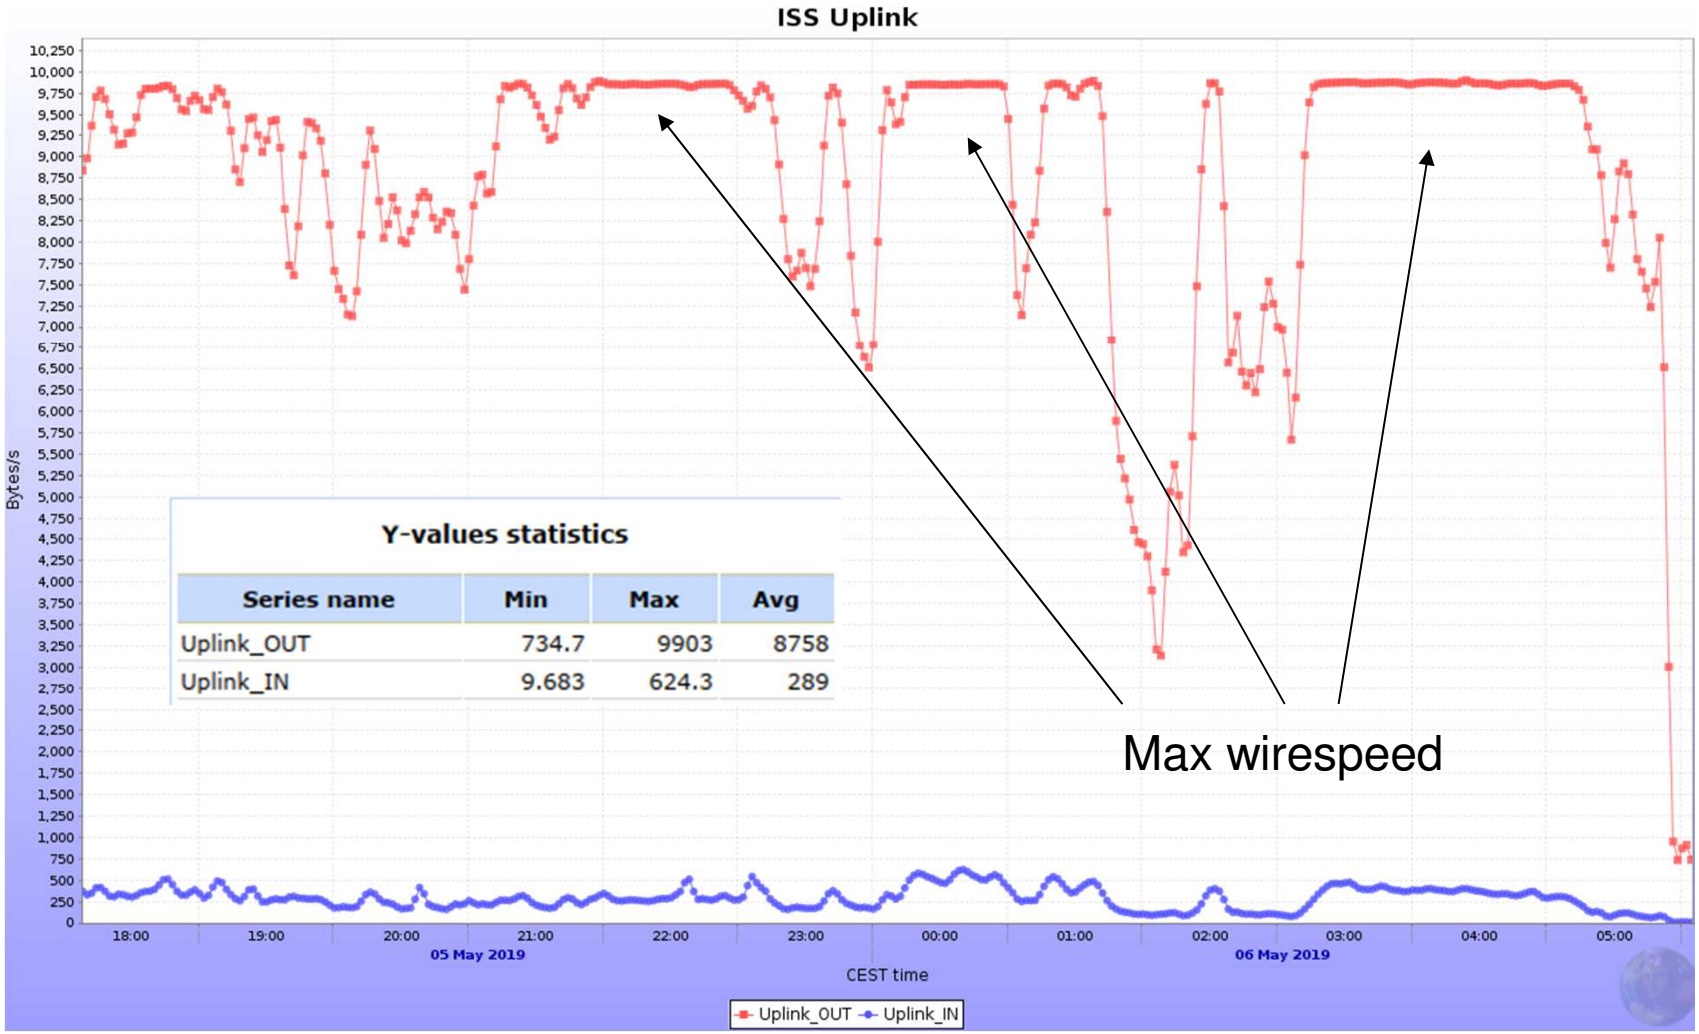

Site status - Ganglia network and cpu reports

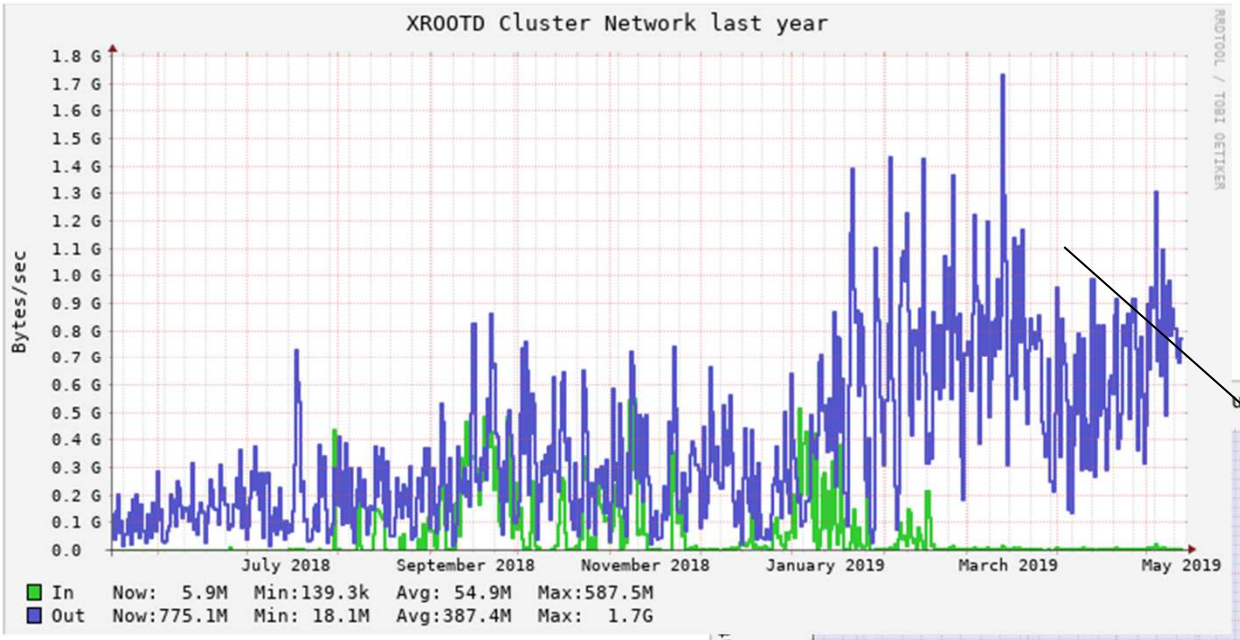

correlated

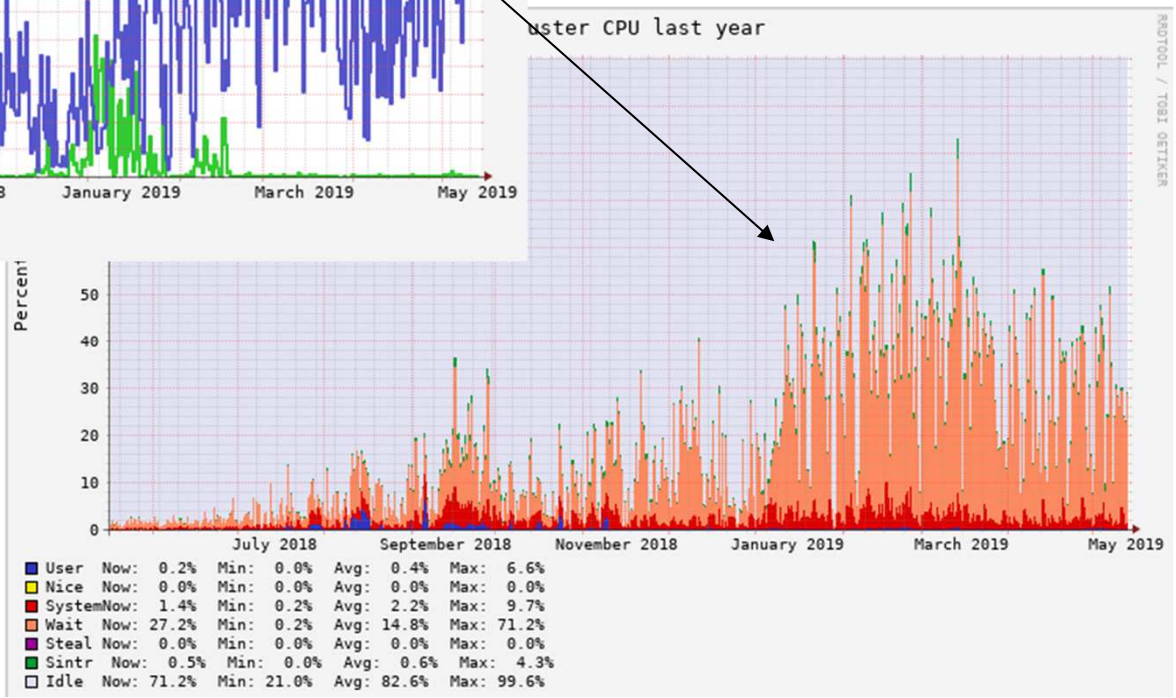

Max of 1.7 GiB/s
private+public transfers

Site status – Summary

- Over 1 M done jobs (11.48 M kSI2k hours)
- 72.92% Jobs efficiency (cpu time/wall time) over last year (77.68% average)
- High availability 97.4%, High success ratio 97.4% of our storages
- Over 5.1 PB data transfer in the last year, 85% storage space occupied
- We reached the maximum bandwidth of our edge connection
- The storage servers also reached the maximum bandwidth of a 10Gb connection

Table of contents

Status of Networking - IPv6

- GGUS Ticket-ID: 132111
- IPv6 dual stack implemented in december 2018 on storages and vobox
- No problems encountered (so far)

Table of contents

EOS

- no plans to migrate to EOS
  - high initial cost effort for new storage cluster
  - upgrading the old and small HDDs have better resource/cost ratio

Table of contents

Site planning

- Increase storage to edge/extern bandwidth to 20 Gb/s (minimum)
  - Further increase difficult
- Increase WNs to storage bandwidth
- Redundant parts for the microblade enclosure needed
  - Two incidents affected the site in the past (both management module and the enclosure switch)
- Complete the entire enclosure (currently 19/28 blades) and retire old servers
- Increase the storage to 2 PB by Run 3

Thank you for your attention!