

# **GridKa Online Storage Monitoring**



# **GridKa & LSDF Online Storage Monitoring**



- Metrics and traps from storage hardware
  - NetApp
  - Seagate
- Infiniband RDMA metrics
- Server metrics (ethernet, ...)
- Spectrum Scale metrics, logs and traps
  - File system performance and statistics
  - Spectrum Scale state
  - Cluster logs















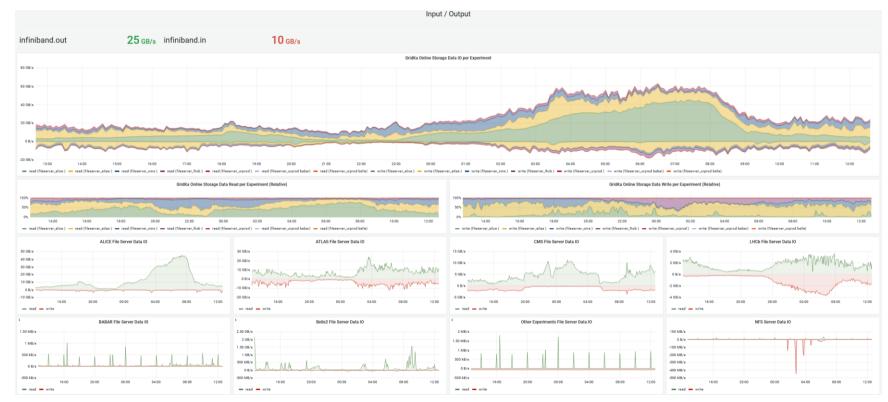
## **Monitoring Platform**

- Karlsruhe Institute of Technology
  - Grafana

- Grafana for metrics and logs
  - Data sources Influx, OpenSearch, MariaDB
- Public and internal dashboards
- Generic and specific dashboards
- Many ad-hoc dashboards to investigate problems
- Never enough metrics ©



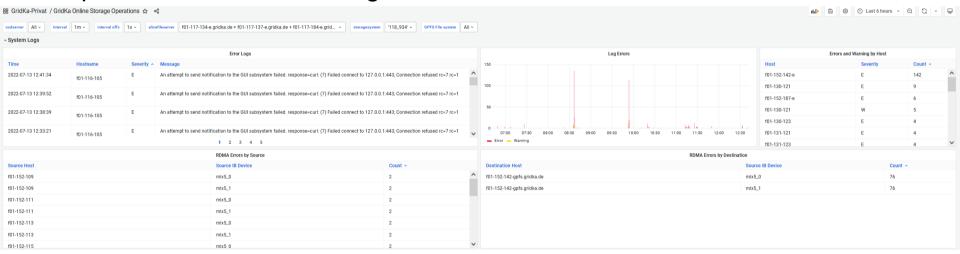
#### https://grafana-sdm.scc.kit.edu/goto/2iKad367k?orgld=1





Spectrum Scale cluster logs, focus on RDMA

GridKa Online Storage Monitoring





## Spectrum Scale cluster logs, focus on waiters



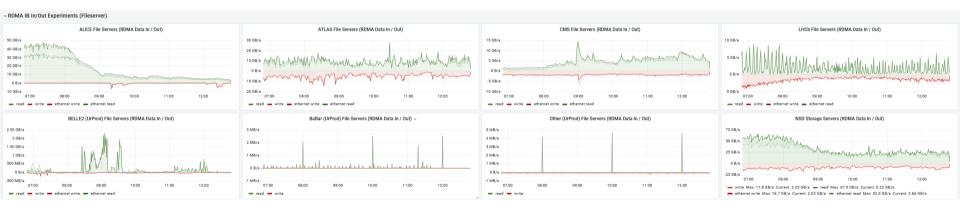


#### Storage system status

v Storage Systems																			
Storage						ige Systems					Drives with Problems								
ID	Manufact.	Model	#Trays	#Drives	Used	#SpareSB	#Spare	Status	OK? ^	ID ^	^	Tray		Drawer	Position	Pool	Status	Cause	
118_912	NetApp	5700	4	240	2.08 PB	0		needsAttn	false	118_	_912	0E500A09800B5843A90000000000000000000	0000	1	6	UNKNOWN	failed	writeOp	^
118_934	NetApp	5700	4	240	2.08 PB	0		needsAttn	false	118_	_934	0E500A09800B584BC100000000000000000000	0000	1	8	DDP04	failed	none	
120_934	NetApp	5600	5	300	1.76 PB	0		needsAttn false		120_	_934	0E50080E520D3F700000000000000000000000000000000000	0000	4	12	DDP06	failed	writeOp	
121_912	.912 NetApp 5600		5	301	1.76 PB			needsAttn	false	121_	_912	0E50080E520D3E80000000000000000000000000000		5	6	DDP03	removed	removed none	
122_912	122_912 NetApp 560		5	300	1.76 PB			needsAttn	false	121_	_912	0E50080E520D3E800000000000000000000000000000000000		5	6	UNKNOWN bypassed bypassU		bypassUnknown	n
122_934	_934 NetApp 5600		5	300	1.76 PB			needsAttn	false	122_	_912	0E50080E520D38700000000000000000000000000000000000		1	3	UNKNOWN	failed	writeOp	
123_912	NetApp	5600	5	300	1.76 PB			needsAttn	false	122_	_912	0E50080E520D387000000000000000000000000		5	12	UNKNOWN	failed	writeOp	
123_934	NetApp	5600	5	300	1.76 PB	0	0	needsAttn	false	122_	_934	0E50080E520D39500000000000000000000000	0000	4	3	DDP02	failed	none	
124_912	NetApp	2806	2	29	11.33 TB	1		needsAttn	false	122_	_934	0E50080E520D395000000000000000000000000	0000	4	5	DDP03	failed	writeOp	
128_912	NetApp	5600	5	300	2.20 PB	0	0	needsAttn	false	123_	_912	0E50080E520D38B00000000000000000000000000000000000	0000	1	4	DDP02	failed	none	~
Storage Systems Pools (Raid Disk Pools)																			
Storage System		Label	Raid Level		pe Total Space			Raid Status	State		Res	erved Drive Count Current Reserved Drive Count		int	Used Reserv		ved Drives 🗸		
122_934		DDP02	raidDiskPool		diskPool	Pool 293.95 TB		optimal	complete		3					3		true	^
122_934		DDP03	raidDiskPool		liskPool	ol 293.95 TB		optimal	complete	complete			0			3		true	
128_912		DDP03	raidDiskPool		fiskPool	(Pool 358.42 TB		optimal	complete		3							true	
123_934		DDP05	raidDiskPool		kPool 287.69 TB			optimal		complete			1			2		true	
118_934		DDP04			diskPool					complete									
123_912										complete									
123_934 E			raidDiskPool			kPool 293.95 TB		optimal											
122_912 DD		DDP02						optimal		complete								true	
128_934		DDP03	raidDiskPool		skPool 366.22 TB			optimal		complete									
123_934		DDP03	DP03 raidDiskPool		fiskPool	kPool 293.95 TB		optimal	complete	complete								true	~
									Storage Syst	ems Pools	s (Raids)								
Storage System ^			Label	Label		Raid Level		Total Space			Free Space		Raid Status		Stat	State		OK?	
120_912		VG01		raid1			22.40 TB				7 TB	optimal		con	complete		true		
124_912		VGC			raid1	raid1		22.40 TB				7 TB	optimal		con	complete		true	
150_101											0 B							true	
150_101																		true	
350,306								<u> </u>				00					*****		



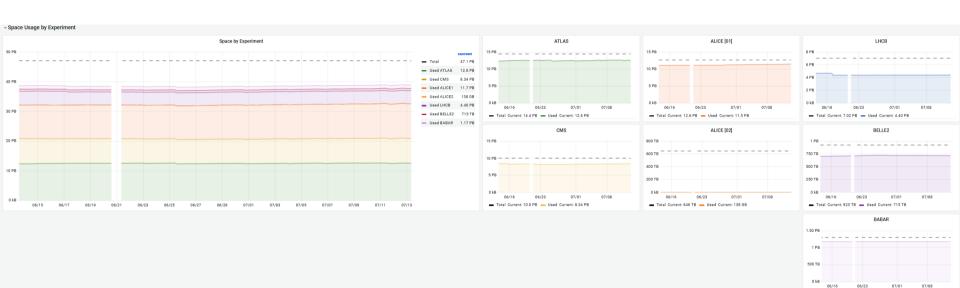
■ IB vs. Ethernet originally to investigate ALICE behavior, but now equally useful to understand tape related activity: checksums etc.





## Experiment Filesystem Usage

GridKa Online Storage Monitoring



■ Total Current: 1.30 PB ■ Used Current: 1.17 PB



## dCache billing





dCache map

