



**ALICE**

A JOURNEY OF DISCOVERY



# System performance monitoring in the ALICE Data Acquisition System with Zabbix

Adriana Telesca  
October 15<sup>th</sup>, 2013  
CHEP 2013, Amsterdam



## ALICE at the CERN LHC

Data Acquisition system requirements:

- 4 GB/s sustained recording rate
- 2.5 GB/s transfer to tape



- For Run 2 (2015-2017):
- ~ 1000 nodes
  - Readout
  - Event building
  - Recording
  - Storage
  - Support (network, PDUs)
  - Operations

For Run 3 (2019-2021):

- ~ 2000 nodes



O2: a new combined online and offline computing for ALICE after 2018

P. Vande Vyvre's talk today at 16:45 – Data Acquisition track

For Run 2 (2015-2017):

- ~ 1000 nodes
- Readout
- Event building
- Recording
- Storage
- Support (network, PDUs)
- Operations

For Run 3 (2019-2021):

~ 2000 nodes

Lemon was used to monitor the DAQ system during Run 1 (2008-2013).


Decision to replace it:

- Lemon future unsure
- Tools with additional/new functionalities
- LHC Long Shutdown 1


Low impact



Extensibility/Flexibility



Scalability




Full administration GUI




Easy access to data



Interface with other components



# Parameters to monitor



CPU  
Memory  
Disk usage  
Network Interfaces  
Processes



Voltage  
Current  
Temperature  
Outlet status  
Disk status



Ethernet:  
CPU utilization  
Memory utilization  
Cards temperature

Fiber Channel:  
RX/TX ports rate

Readout links Bytes In/Out  
DAQ XOFF, HLT XOFF  
Processes CPU  
and memory





## Selection criteria:

1. SNMP
2. Logical grouping
3. Large user community
4. Distributed monitoring

Source: [http://en.wikipedia.org/wiki/Comparison\\_of\\_network\\_monitoring\\_systems](http://en.wikipedia.org/wiki/Comparison_of_network_monitoring_systems)

Name	Agent	SNMP	Syslog	WebApp	Data Storage Method	License
Cacti	No	Yes	Yes	Full Control	RRDtool, MySQL	GPL
Icinga	Supported	Via plugin	Via plugin	Full Control	MySQL, PostgreSQL, Oracle Database	GPL
Zabbix	Supported	Yes	Yes	Full Control	Oracle, MySQL, PostgreSQL, IBM DB2, SQLite	GPL
Zenoss	No	Yes	Yes	Full Control	ZODB, MySQL, RRDtool	GPL
+ Splunk	Supported	Yes	Yes	Full control	Raw files	Commercial


+ MonALISA

# Tools comparison

Name	Data gathering	Graphing	Triggering	Scalability	Data Storage	Extensibility
Icinga	Agent	0	1	1 – up to 1000 hosts	DB	2
Cacti	Server	2	0	1 – up to 1000 hosts	RRDtool – DB	2
Zenoss	Server	1	1	2 – 1000+	RRDtool – DB	1
Zabbix	Agent or Server	2	1	2 – 1000+	DB	2
Splunk	Agent	2	1	2 – 1000+	Raw files	2
MonALISA	Agent	2	1	2 – 1000+	DB	2



0-1 Absent-Present




0-1-2 Absent - Present but not good - Good

# Tools comparison

Name	Data gathering	Graphing	Triggering	Scalability	Data Storage	Extensibility
Icinga	Agent	0	1	1 – up to 1000 hosts	DB	2
Cacti	Server	2	0	1 – up to 1000 hosts	RRDtool – DB	2
Zenoss	Server	1	1	2 – 1000+	RRDtool – DB	1
Zabbix	Agent or Server	2	1	2 – 1000+	DB	2
Splunk	Agent	2	1	2 – 1000+	Raw files	2
MonALISA	Agent	2	1	2 – 1000+	DB	2



0-1 Absent-Present




0-1-2 Absent - Present but not good - Good

# Tools comparison

Name	Data gathering	Graphing	Triggering	Scalability	Data Storage	Extensibility
Icinga	Agent	0	1	1 – up to 1000 hosts	DB	2
Cacti	Server	2	0	1 – up to 1000 hosts	RRDtool – DB	2
Zenoss	Server	1	1	2 – 1000+	RRDtool – DB	1
Zabbix	Agent or Server	2	1	2 – 1000+	DB	2
Splunk	Agent	2	1	2 – 1000+	Raw files	2
MonALISA	Agent	2	1	2 – 1000+	DB	2



0-1 Absent-Present




0-1-2 Absent - Present but not good - Good

# Tools comparison

Name	Data gathering	Graphing	Triggering	Scalability	Data Storage	Extensibility
Icinga	Agent	0	1	1 – up to 1000 hosts	DB	2
Cacti	Server	2	0	1 – up to 1000 hosts	RRDtool – DB	2
Zenoss	Server	1	1	2 – 1000+	RRDtool – DB	1
Zabbix	Agent or Server	2	1	2 – 1000+	DB	2
Splunk	Agent	2	1	2 – 1000+	Raw files	2
MonALISA	Agent	2	1	2 – 1000+	DB	2



0-1 Absent-Present




0-1-2 Absent - Present but not good - Good

# Tools comparison

Name	Data gathering	Graphing	Triggering	Scalability	Data Storage	Extensibility
Icinga	Agent	0	1	1 – up to 1000 hosts	DB	2
Cacti	Server	2	0	1 – up to 1000 hosts	RRDtool – DB	2
Zenoss	Server	1	1	2 – 1000+	RRDtool – DB	1
Zabbix	Agent or Server	2	1	2 – 1000+	DB	2
Splunk	Agent	2	1	2 – 1000+	Raw files	2
MonALISA	Agent	2	1	2 – 1000+	DB	2



0-1 Absent-Present




0-1-2 Absent - Present but not good - Good

# Tools comparison

Name	SNMP	Community	Granularity	Auto Discovery	Free
Icinga	2	2	1 - 1 minute / metric	2	1
Cacti	2	2	1 - 1 minute / metric	1	1
Zenoss	1	1	1- 1 minute / collector	2	1
Zabbix	2	2	2 - No limit / metric	2	1
Splunk	2	2	2 - No limit / metric	2	0
MonALISA	2	1	1 - 1 minute / metric	2	1



0-1 Absent-Present




0-1-2 Absent - Present but not good - Good

# Tools comparison

Name	SNMP	Community	Granularity	Auto Discovery	Free
Icinga	2	2	1 - 1 minute / metric	2	1
Cacti	2	2	1 - 1 minute / metric	1	1
Zenoss	1	1	1- 1 minute / collector	2	1
Zabbix	2	2	2 - No limit / metric	2	1
Splunk	2	2	2 - No limit / metric	2	0
MonALISA	2	1	1 - 1 minute / metric	2	1



0-1 Absent-Present



0-1-2 Absent - Present but not good - Good




# Tools comparison

Name	SNMP	Community	Granularity	Auto Discovery	Free
Icinga	2	2	1 - 1 minute / metric	2	1
Cacti	2	2	1 - 1 minute / metric	1	1
Zenoss	1	1	1- 1 minute / collector	2	1
Zabbix	2	2	2 - No limit / metric	2	1
Splunk	2	2	2 - No limit / metric	2	0
MonALISA	2	1	1 - 1 minute / metric	2	1




0-1 Absent-Present




0-1-2 Absent - Present but not good - Good

# Tools comparison

Name	SNMP	Community	Granularity	Auto Discovery	Free
Icinga	2	2	1 - 1 minute / metric	2	1
Cacti	2	2	1 - 1 minute / metric	1	1
Zenoss	1	1	1- 1 minute / collector	2	1
Zabbix	2	2	2 - No limit / metric	2	1
Splunk	2	2	2 - No limit / metric	2	0
MonALISA	2	1	1 - 1 minute / metric	2	1



0-1 Absent-Present




0-1-2 Absent - Present but not good - Good

# Tools comparison

Name	SNMP	Community	Granularity	Auto Discovery	Free	Total
Icinga	2	2	1 - 1 minute / metric	2	1	12
Cacti	2	2	1 - 1 minute / metric	1	1	12
Zenoss	1	1	1- 1 minute / collector	2	1	11
Zabbix	2	2	2 - No limit / metric	2	1	16
Splunk	2	2	2 - No limit / metric	2	0	15
MonALISA	2	1	1 - 1 minute / metric	2	1	14



0-1 Absent-Present




0-1-2 Absent - Present but not good - Good

# Tools comparison

Name	SNMP	Community	Granularity	Auto Discovery	Free	Total
Icinga	2	2	1 - 1 minute / metric	2	1	12
Cacti	2	2	1 - 1 minute / metric	1	1	12
Zenoss	1	1	1- 1 minute / collector	2	1	11
Zabbix	2	2	2 - No limit / metric	2	1	16
Splunk	2	2	2 - No limit / metric	2	0	15
MonALISA	2	1	1 - 1 minute / metric	2	1	14



0-1 Absent-Present



0-1-2 Absent - Present but not good - Good


- Graphing
- Full configuration GUI
- Many ways of data retrieval  
→ scalability



# Zabbix characteristics



Name	Last check	Last value	Change	History
<b>[-] CPU (16 Items)</b>				
<b>[+] Filesystems / Disks (28 Items)</b>				
<b>[+] General (8 Items)</b>				
<b>[-] Memory (15 Items)</b>				
Available memory	09 Oct 2013 12:34:04	23.29 GB	-48 KB	<a href="#">Graph</a>
Available memory percentage	09 Oct 2013 12:34:08	98.73 %	-	<a href="#">Graph</a>
Buffers memory	09 Oct 2013 12:34:05	172.25 MB	-	<a href="#">Graph</a>
Buffers memory percentage	09 Oct 2013 12:34:00	0.71 %	-	<a href="#">Graph</a>
Cached memory	09 Oct 2013 12:34:06	218.46 MB	-	<a href="#">Graph</a>



Description	Key	Interval	History	Type ↓↑
<a href="#">snmpd cpu</a>	process_memcpu.sh[snmpd,1]	30	90	External check
<a href="#">snmpd mem</a>	process_memcpu.sh[snmpd,2]	30	90	External check
<a href="#">Aggregate CPU load</a>	grpavg["Linux servers","system.cpu.load[,avg1]",last,10]	30	90	Zabbix aggregate
<a href="#">Free disk space on /opt</a>	vfs.fs.size[/opt,free]	30	7	Zabbix agent (active)
<a href="#">Free disk space on /var</a>	vfs.fs.size[/var,free]	30	7	Zabbix agent (active)
<a href="#">Free disk space on /</a>	vfs.fs.size[/,free]	30	7	Zabbix agent (active)
<a href="#">Free disk space on /home</a>	vfs.fs.size[/home,free]	30	7	Zabbix agent (active)
<a href="#">ssh performance</a>	ssh_perf	30	90	Simple check
<a href="#">Boot space used sender</a>	boot_space		90	Zabbix trapper
<a href="#">Local space used sender</a>	local_space		90	Zabbix trapper
<a href="#">Home space used sender</a>	home_space		90	Zabbix trapper
<a href="#">Number of processes SNMP</a>	hrSystemProcesses.0	30	90	SNMPv1 agent
<a href="#">Total number of inodes on /home</a>	vfs.fs.inode[/home,total]	3600	7	Zabbix agent
<a href="#">Total swap space</a>	system.swap.size[,total]	1800	7	Zabbix agent

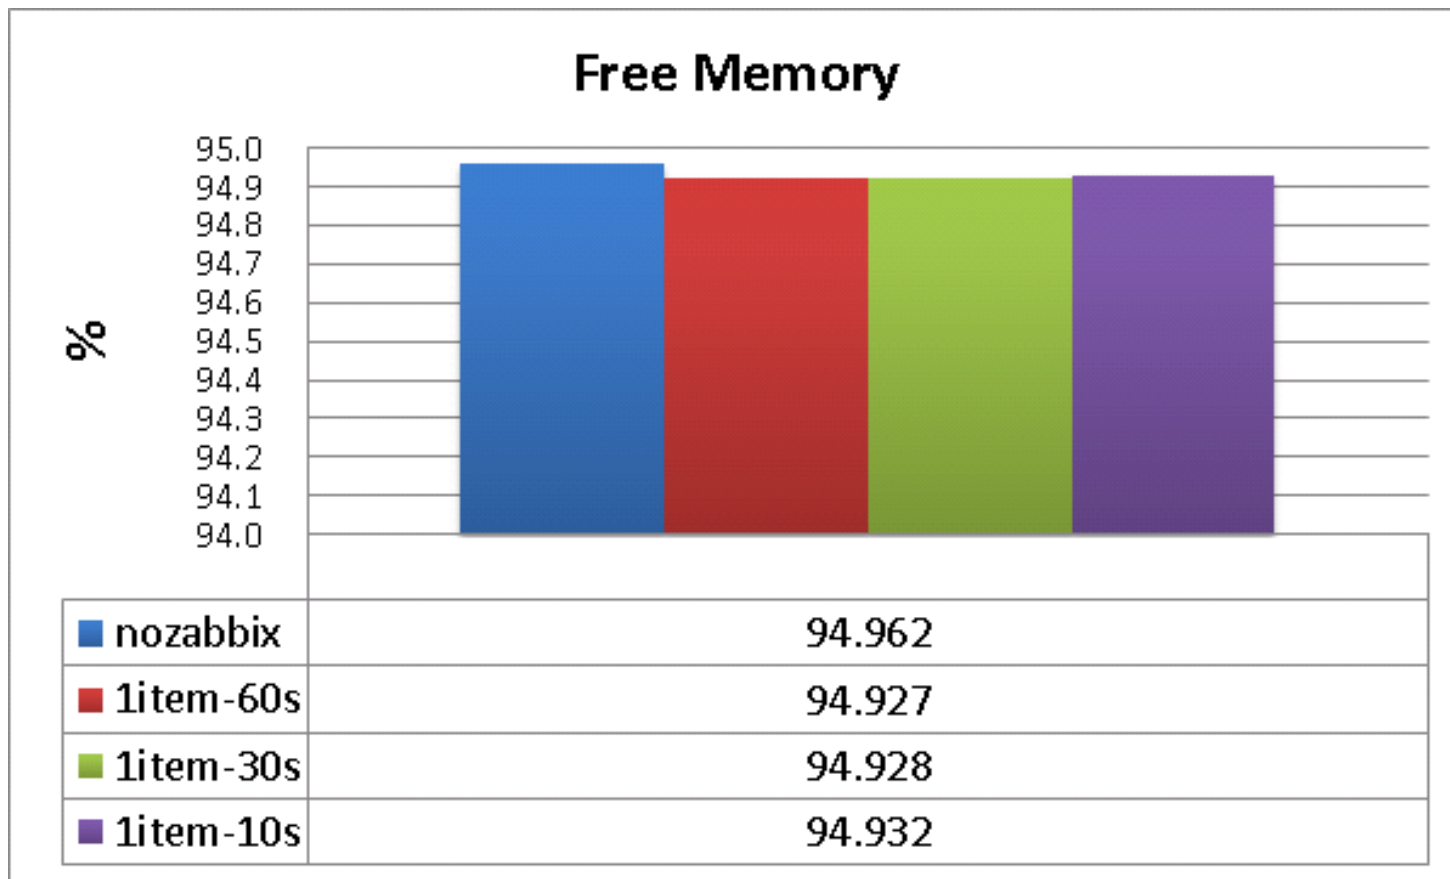
Description	Key	Interval	History	Type ↓
<a href="#">snmpd cpu</a>	process_memcpu.sh[snmpd,1]	30	90	External check
<a href="#">snmpd mem</a>	process_memcpu.sh[snmpd,2]	30	90	External check
<a href="#">Aggregate CPU load</a>	grpavg["Linux servers","system.cpu.load[,avg1]","last,10]	30	90	Zabbix aggregate
<a href="#">Free disk space on /opt</a>	vfs.fs.size[/opt,free]	30	7	Zabbix agent (active)
<a href="#">Free disk space on /var</a>	vfs.fs.size[/var,free]	30	7	Zabbix agent (active)
<a href="#">Free disk space on /</a>	vfs.fs.size[/,free]	30	7	Zabbix agent (active)
<a href="#">Free disk space on /home</a>	vfs.fs.size[/home,free]	30	7	Zabbix agent (active)
<a href="#">ssh performance</a>	ssh_perf	30	90	Simple check
<a href="#">Boot space used sender</a>	boot_space		90	Zabbix trapper
<a href="#">Local space used sender</a>	local_space		90	Zabbix trapper
<a href="#">Home space used sender</a>	home_space		90	Zabbix trapper
<a href="#">Number of processes SNMP</a>	hrSystemProcesses.0	30	90	SNMPv1 agent
<a href="#">Total number of inodes on /home</a>	vfs.fs.inode[/home,total]	3600	7	Zabbix agent
<a href="#">Total swap space</a>	system.swap.size[,total]	1800	7	Zabbix agent




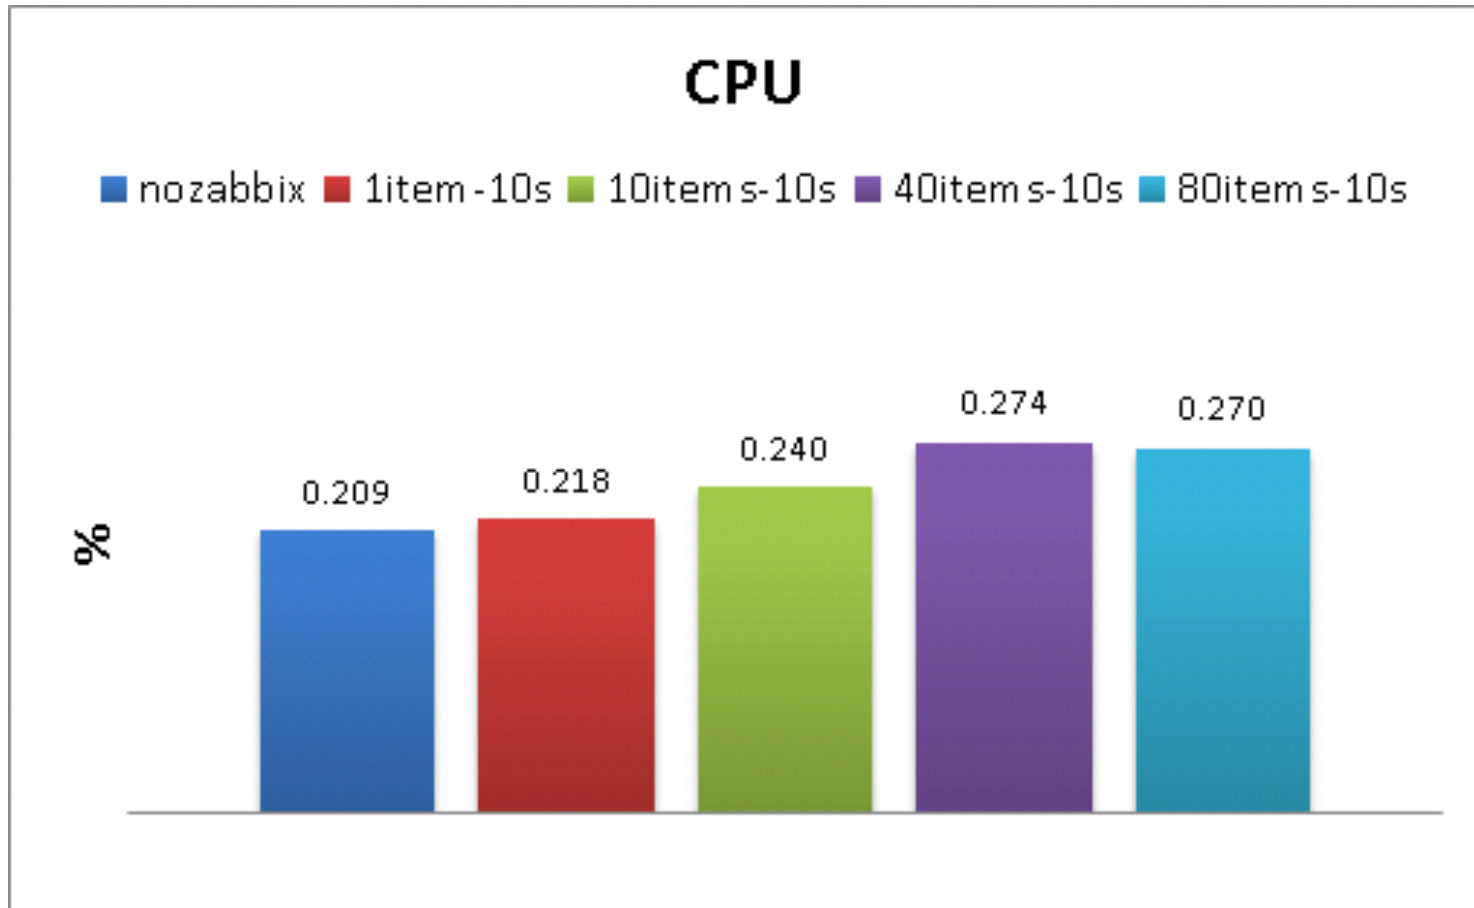
Description	Key	Interval	History	Type ↓↑
<a href="#">snmpd cpu</a>	process_memcpu.sh[snmpd,1]	30	90	External check
<a href="#">snmpd mem</a>	process_memcpu.sh[snmpd,2]	30	90	External check
<a href="#">Aggregate CPU load</a>	grpavg["Linux servers","system.cpu.load[,avg1]",last,10]	30	90	Zabbix aggregate
<a href="#">Free disk space on /opt</a>	vfs.fs.size[/opt,free]	30	7	Zabbix agent (active)
<a href="#">Free disk space on /var</a>	vfs.fs.size[/var,free]	30	7	Zabbix agent (active)
<a href="#">Free disk space on /</a>	vfs.fs.size[/,free]	30	7	Zabbix agent (active)
<a href="#">Free disk space on /home</a>	vfs.fs.size[/home,free]	30	7	Zabbix agent (active)
<a href="#">ssh performance</a>	ssh_perf	30	90	Simple check
<a href="#">Boot space used sender</a>	boot_space		90	Zabbix trapper
<a href="#">Local space used sender</a>	local_space		90	Zabbix trapper
<a href="#">Home space used sender</a>	home_space		90	Zabbix trapper
<a href="#">Number of processes SNMP</a>	hrSystemProcesses.0	30	90	SNMPv1 agent
<a href="#">Total number of inodes on /home</a>	vfs.fs.inode[/home,total]	3600	7	Zabbix agent
<a href="#">Total swap space</a>	system.swap.size[,total]	1800	7	Zabbix agent

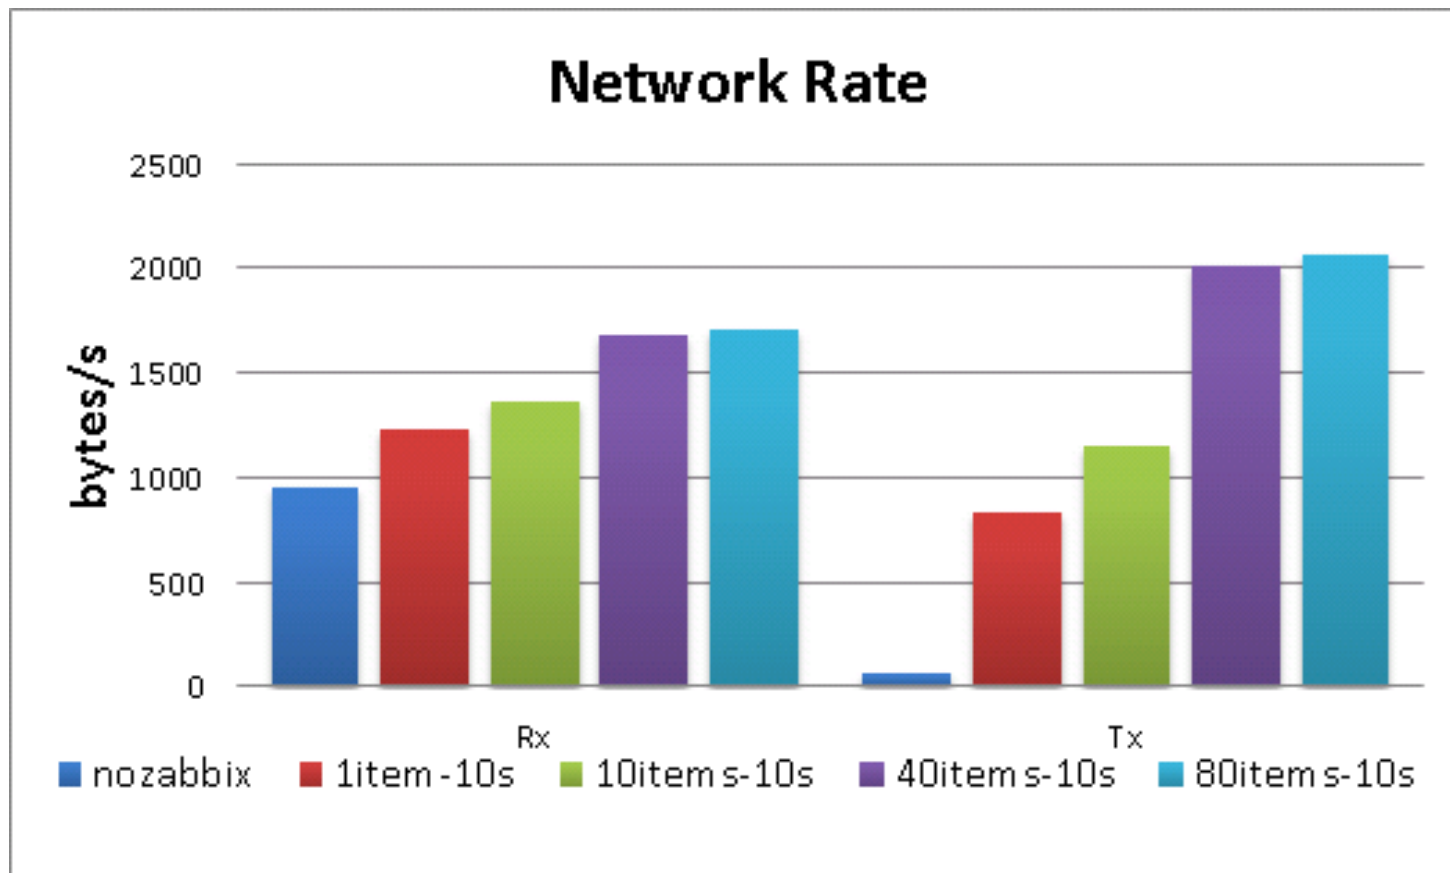
Description	Key	Interval	History	Type ↓↑
<a href="#">snmpd cpu</a>	process_memcpu.sh[snmpd,1]	30	90	External check
<a href="#">snmpd mem</a>	process_memcpu.sh[snmpd,2]	30	90	External check
<a href="#">Aggregate CPU load</a>	grpavg["Linux servers","system.cpu.load[,avg1]",last,10]	30	90	Zabbix aggregate
<a href="#">Free disk space on /opt</a>	vfs.fs.size[/opt,free]	30	7	Zabbix agent (active)
<a href="#">Free disk space on /var</a>	vfs.fs.size[/var,free]	30	7	Zabbix agent (active)
<a href="#">Free disk space on /</a>	vfs.fs.size[/,free]	30	7	Zabbix agent (active)
<a href="#">Free disk space on /home</a>	vfs.fs.size[/home,free]	30	7	Zabbix agent (active)
<a href="#">ssh performance</a>	ssh_perf	30	90	Simple check
<a href="#">Boot space used sender</a>	boot_space		90	Zabbix trapper
<a href="#">Local space used sender</a>	local_space		90	Zabbix trapper
<a href="#">Home space used sender</a>	home_space		90	Zabbix trapper
<a href="#">Number of processes SNMP</a>	hrSystemProcesses.0	30	90	SNMPv1 agent
<a href="#">Total number of inodes on /home</a>	vfs.fs.inode[/home,total]	3600	7	Zabbix agent
<a href="#">Total swap space</a>	system.swap.size[,total]	1800	7	Zabbix agent

Description	Key	Interval	History	Type ↓↑
<a href="#">snmpd cpu</a>	process_memcpu.sh[snmpd,1]	30	90	External check
<a href="#">snmpd mem</a>	process_memcpu.sh[snmpd,2]	30	90	External check
<a href="#">Aggregate CPU load</a>	grpavg["Linux servers","system.cpu.load[,avg1]",last,10]	30	90	Zabbix aggregate
<a href="#">Free disk space on /opt</a>	vfs.fs.size[/opt,free]	30	7	Zabbix agent (active)
<a href="#">Free disk space on /var</a>	vfs.fs.size[/var,free]	30	7	Zabbix agent (active)
<a href="#">Free disk space on /</a>	vfs.fs.size[/,free]	30	7	Zabbix agent (active)
<a href="#">Free disk space on /home</a>	vfs.fs.size[/home,free]	30	7	Zabbix agent (active)
<a href="#">ssh performance</a>	ssh_perf	30	90	Simple check
<a href="#">Boot space used sender</a>	boot_space		90	Zabbix trapper
<a href="#">Local space used sender</a>	local_space		90	Zabbix trapper
<a href="#">Home space used sender</a>	home_space		90	Zabbix trapper
<a href="#">Number of processes SNMP</a>	hrSystemProcesses.0	30	90	SNMPv1 agent
<a href="#">Total number of inodes on /home</a>	vfs.fs.inode[/home,total]	3600	7	Zabbix agent
<a href="#">Total swap space</a>	system.swap.size[,total]	1800	7	Zabbix agent

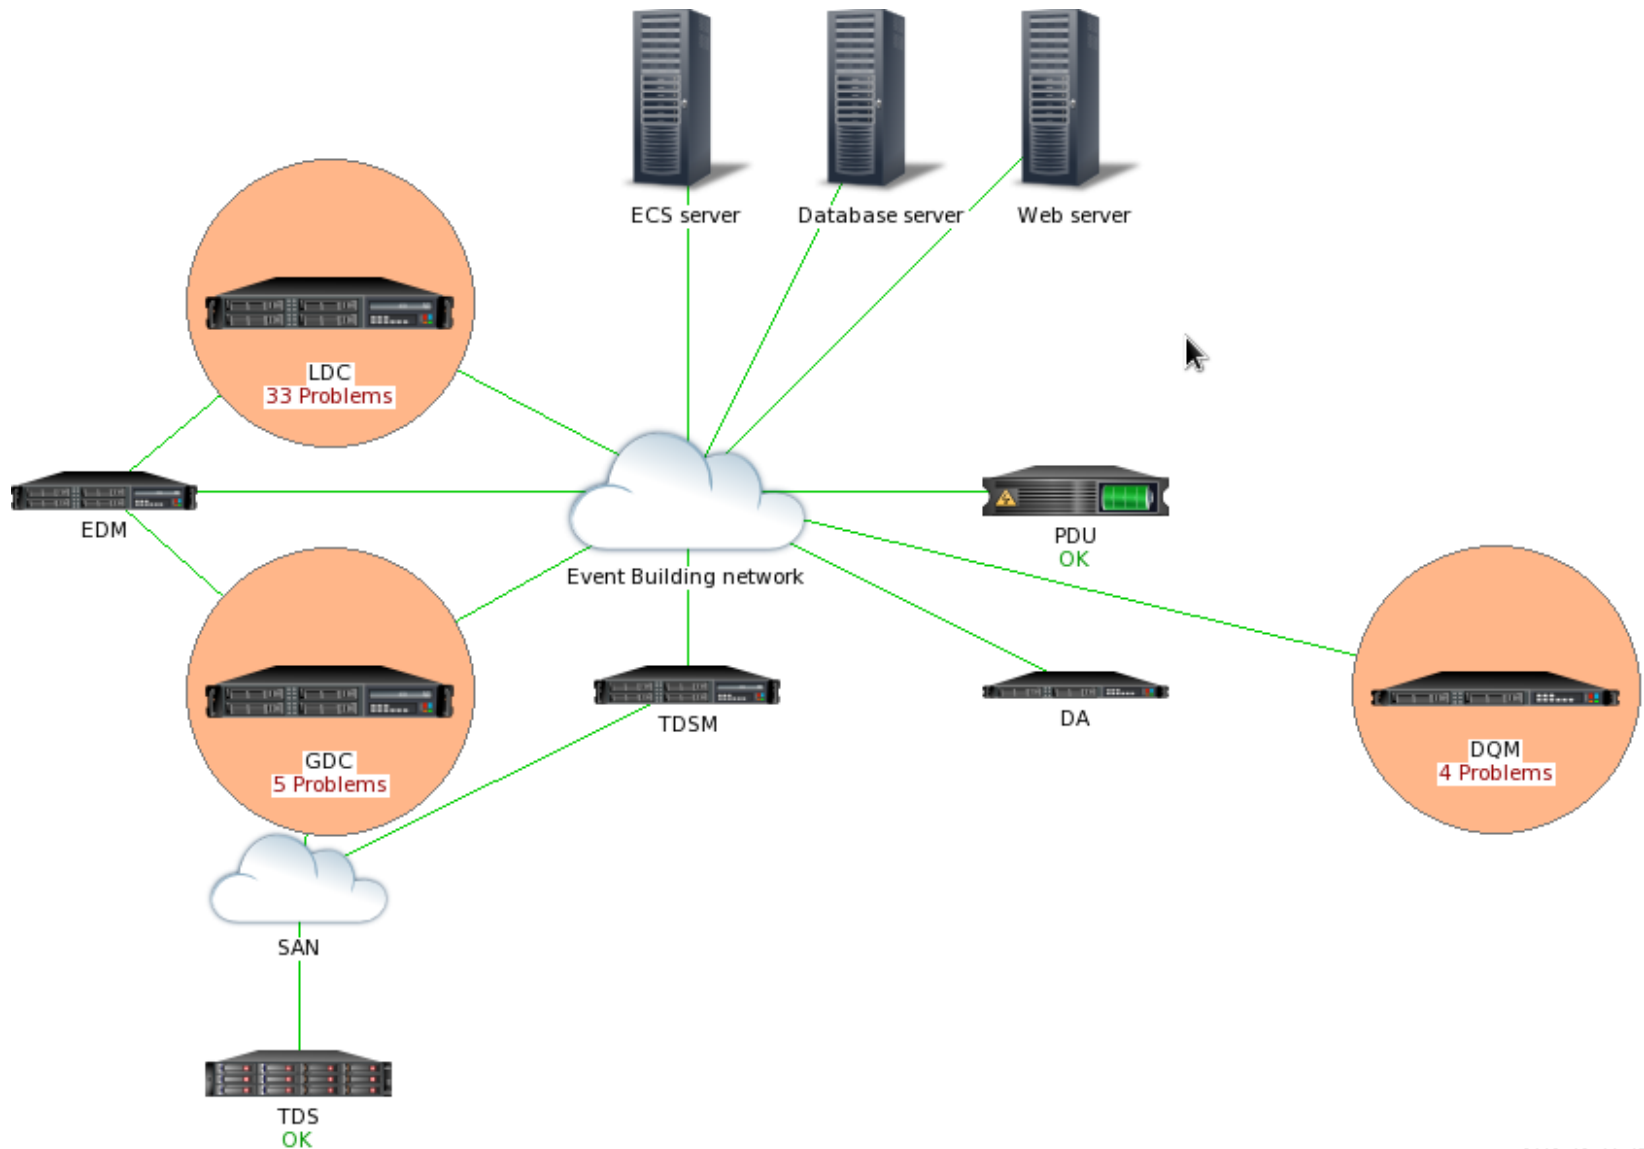








# Zabbix dashboard and usage






# Zabbix dashboard and usage



cpu-system for group LDC-TPC (1h)



		last	min	avg	max
aidrefpc015: CPU system	[avg]	0.02 %	0.01 %	0.02 %	0.14 %
aidrefpc013: CPU system	[avg]	0.02 %	0.01 %	0.03 %	0.14 %
aidrefpc012: CPU system	[avg]	0.03 %	0.01 %	0.03 %	0.13 %
Trigger: CPU system utilization too high on aidrefpc015				[ > 90]	
Trigger: CPU system utilization too high on aidrefpc013				[ > 90]	
Trigger: CPU system utilization too high on aidrefpc012				[ > 90]	

Data from history. Generated in 0.20 sec.

# Zabbix dashboard and usage



Hosts	Agent ping	Available memory	Available memory percentage	Buffers memory	Buffers memory percentage	Cached memory	Cached memory percentage	Checksum of /etc/passwd	Context switches per second	CPU idle	CPU interrupt	CPU iowait
<a href="#">aldrefpc001</a>	<a href="#">Up (1)</a>	<a href="#">23.29 GB</a>	<a href="#">98.74 %</a>	<a href="#">169.58 MB</a>	<a href="#">0.7 %</a>	<a href="#">218.46 MB</a>	<a href="#">0.9 %</a>	<a href="#">1838775862</a>	<a href="#">112.95 sps</a>	<a href="#">99.98 %</a>	<a href="#">0 %</a>	<a href="#">0.002083 %</a>
<a href="#">aldrefpc002</a>	<a href="#">Up (1)</a>	<a href="#">23.29 GB</a>	<a href="#">98.74 %</a>	<a href="#">169.88 MB</a>	<a href="#">0.7 %</a>	<a href="#">218.42 MB</a>	<a href="#">0.9 %</a>	<a href="#">1838775862</a>	<a href="#">113.62 sps</a>	<a href="#">99.97 %</a>	<a href="#">0 %</a>	<a href="#">0.002083 %</a>
<a href="#">aldrefpc003</a>	<a href="#">Up (1)</a>	<a href="#">23.29 GB</a>	<a href="#">98.74 %</a>	<a href="#">170.33 MB</a>	<a href="#">0.71 %</a>	<a href="#">218.42 MB</a>	<a href="#">0.9 %</a>	<a href="#">1838775862</a>	<a href="#">112.94 sps</a>	<a href="#">99.98 %</a>	<a href="#">0 %</a>	<a href="#">0.002083 %</a>
<a href="#">aldrefpc004</a>	<a href="#">Up (1)</a>	<a href="#">23.29 GB</a>	<a href="#">98.74 %</a>	<a href="#">171.04 MB</a>	<a href="#">0.71 %</a>	<a href="#">218.43 MB</a>	<a href="#">0.9 %</a>	<a href="#">1838775862</a>	<a href="#">113.63 sps</a>	<a href="#">99.98 %</a>	<a href="#">0 %</a>	<a href="#">0.003124 %</a>
<a href="#">aldrefpc005</a>	<a href="#">Up (1)</a>	<a href="#">23.28 GB</a>	<a href="#">98.72 %</a>	<a href="#">170.32 MB</a>	<a href="#">0.71 %</a>	<a href="#">218.44 MB</a>	<a href="#">0.9 %</a>	<a href="#">1838775862</a>	<a href="#">108.02 sps</a>	<a href="#">99.98 %</a>	<a href="#">0 %</a>	<a href="#">0.002083 %</a>
<a href="#">aldrefpc006</a>	<a href="#">Up (1)</a>	<a href="#">23.29 GB</a>	<a href="#">98.74 %</a>	<a href="#">170.18 MB</a>	<a href="#">0.7 %</a>	<a href="#">218.46 MB</a>	<a href="#">0.9 %</a>	<a href="#">1838775862</a>	<a href="#">110.07 sps</a>	<a href="#">99.97 %</a>	<a href="#">0 %</a>	<a href="#">0.002083 %</a>

# Zabbix dashboard and usage



## Status of Zabbix 🏠 ⬆

Parameter	Value	Details
Zabbix server is running	Yes	localhost:10051
Number of hosts (monitored/not monitored/templates)	115	63 / 2 / 50
Number of items (monitored/disabled/not supported)	5146	4959 / 52 / 135
Number of triggers (enabled/disabled)[problem/unknown/ok]	1560	1560 / 0 [4 / 0 / 1556]
Number of users (online)	20	1
Required server performance, new values per second	130.49	-

Updated: 11:37:32

## System status 🏠 ⬆

Host group	Disaster	High	Average	Warning	Information	Not classified
<a href="#">DQM</a>	0	0	3	1	0	0
<a href="#">GDC</a>	0	0	5	0	0	0
<a href="#">LDC</a>	0	0	33	0	0	0
<a href="#">MON</a>	0	0	6	0	0	0
<a href="#">PDU</a>	0	0	0	0	0	0
<a href="#">TDS</a>	0	0	0	0	0	0
<a href="#">Zabbix servers</a>	0	0	1	1	0	0

Updated: 15:38:19

The evaluation of different monitoring tools resulted in the selection of Zabbix.

Zabbix meets the ALICE DAQ needs.

Zabbix will be in production for Run 2.



**ALICE**

A JOURNEY OF DISCOVERY



Thanks.  
Questions?