



HYBRILIT MONITORING SYSTEM

MARTIN VALA – INSTITUTE OF EXPERIMENTAL PHYSICS SAS KOSICE

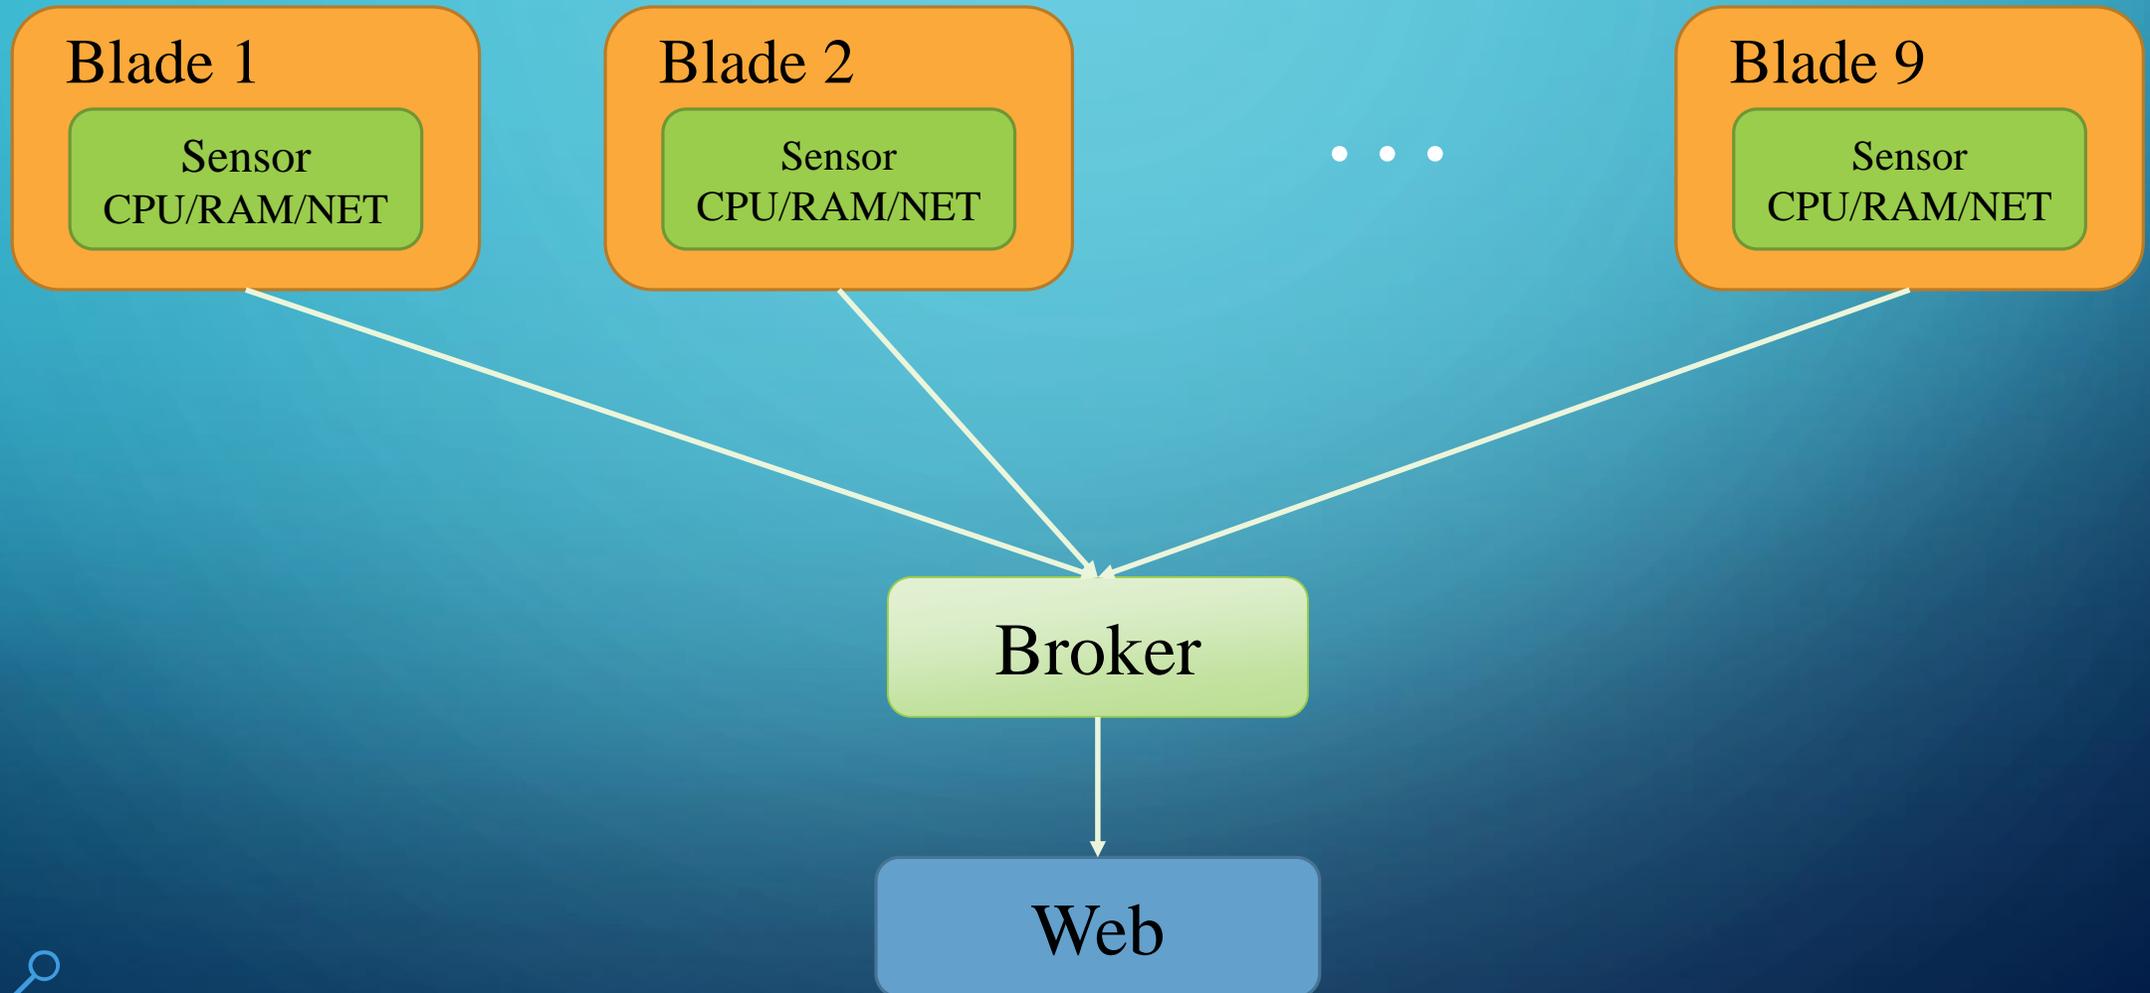
YURI BUTENKO – JOINT INSTITUTE FOR NUCLEAR RESEARCH, LIT

DMITRY BELYAKOV – JOINT INSTITUTE FOR NUCLEAR RESEARCH, LIT

IVAN KASHUNIN – JOINT INSTITUTE FOR NUCLEAR RESEARCH, LIT

HLIT-MON (GRID 2016)

Alexander Mayorov/Yuri Butenko/Gleb Podolyan and Martin Vala



HLIT-MON:SENSOR

Sensor was written on C++ with **libgtop2**

Sensor and Broker was connected via **ZMQ**

- CPU usage (user/idle/system)
- RAM usage (rss/buffer/cache/free)
- NET usage (Eth+IB in/out)

HLIT-MON

- HLIT-MON it's not hardware monitoring, it was monitoring for utilization CPU/MEM/NET by user's apps.
- Information from SENSORS accumulate on BROKER and then sending to WEB

Stat-HybriLit(0.3.0)

Home
Statistic
Total workload
Running jobs
Executed jobs
Distribution of jobs
Cluster monitoring
Organizations

Table

Node	CPU			Memory					Network	
	User	Sys	Idle	RSS	Cached	Buffer	Free	Total	In	Out
blade01	80.4%	0.0%	39.8%	1.3 GB	3.7 GB	4.1 MB	121.2 GB	128.1 GB	128 B/s	2 kB/s
blade02	0.0%	0.0%	100.0%	1.7 GB	108.9 GB	2.4 MB	14.5 GB	128.1 GB	128 B/s	2.3 kB/s
blade03	0.0%	0.0%	100.0%	1.8 GB	79.1 GB	3.8 MB	45.5 GB	128.1 GB	248 B/s	2.6 kB/s
blade04	3.8%	0.8%	95.8%	1.7 GB	119.5 GB	4 MB	4.9 GB	128.1 GB	248 B/s	2 kB/s
blade05	0.0%	0.0%	100.0%	1.4 GB	107.2 GB	1.7 MB	17.8 GB	128.1 GB	248 B/s	2 kB/s
blade06	0.0%	0.0%	100.0%	1.4 GB	102.5 GB	3.1 MB	22.3 GB	128.1 GB	188 B/s	2 kB/s
blade07	0.0%	0.0%	100.0%	1.3 GB	5.7 GB	4 MB	119.2 GB	128.1 GB	248 B/s	2 kB/s
blade08	83.9%	0.0%	16.0%	249.6 GB	253.7 GB	24 kB	1.5 GB	504.8 GB	572 B/s	2.2 kB/s
blade09	0.0%	0.0%	100.0%	4.2 GB	207.2 GB	44 kB	293.5 GB	504.8 GB	572 B/s	2.2 kB/s
stat-hybrilit	1.0%	1.0%	98.0%	544.7 MB	1.2 GB	1.8 MB	5.9 GB	7.6 GB	111 kB/s	158 kB/s
TOTAL	-	-	-	264.7 GB	989.5 GB	25 MB	646 GB	1.9 TB	113.5 kB/s	175.4 kB/s
AVG	14.9%	0.2%	84.9%	26.5 GB	99 GB	2.5 MB	64.6 GB	190 GB	11.3 kB/s	17.5 kB/s

Chart

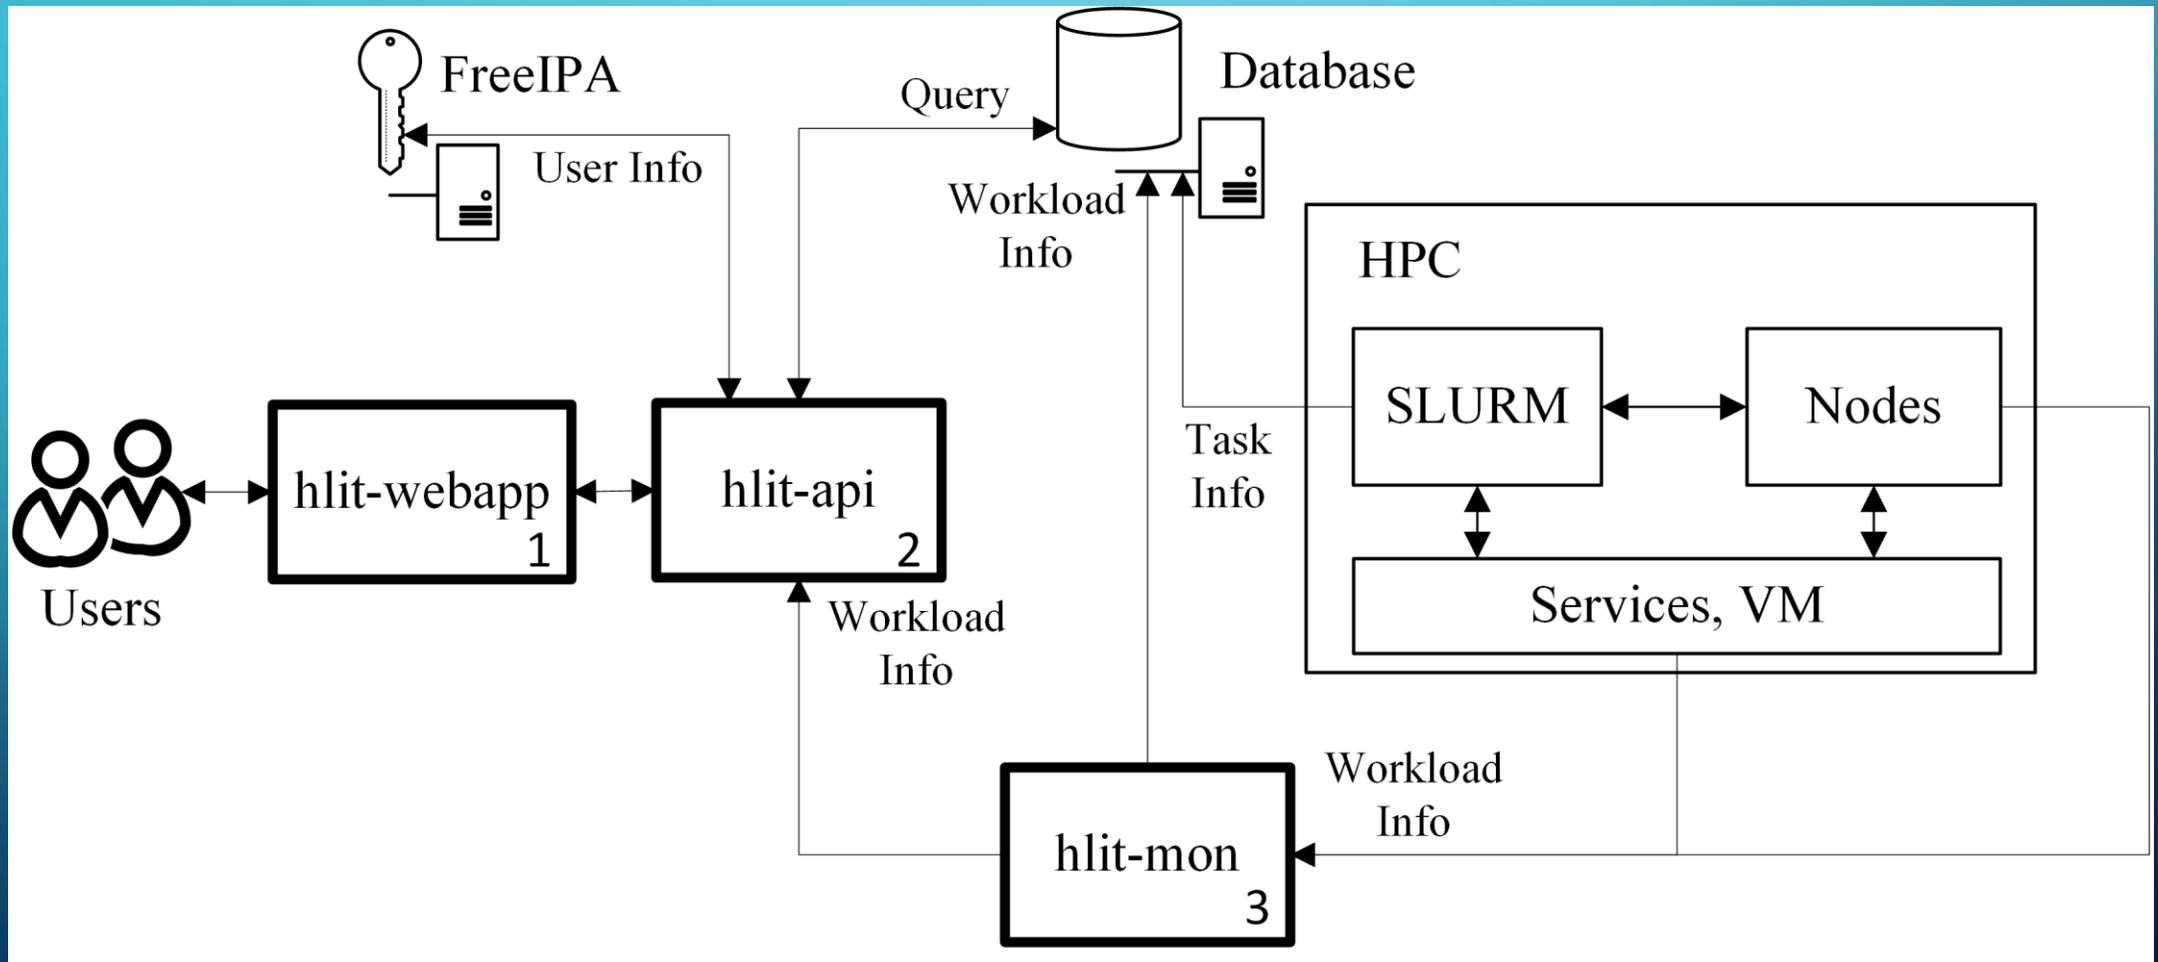
- After that user start job via SLURM on CLUSTER we can see how app utilize resources and if something went wrong (SYS load bigger than USER load) we can help user for improve his code if it needed.

HLIT-WEBAPP (MMCP2017, PFUR ITMM2017)

We used hlit-mon for build our web app for monitoring workload on HybriLIT

- Auth via FreeIPA (Kerberos+LDAP).
- Monitoring blades, vm, services, storages via hlit-mon.
- Info was represented via hlit-web app written on Angular.
- Archiving workload info in DB.

HLIT-WEBAPP (MMCP2017, PFUR ITMM2017)



Stat-HybriLIT

Statistics of cluster "HybriLIT"

Stat-HybriLIT - is a web-service for information support of users of cluster "HybriLIT". The service provides information about user, user tasks, load of hardware resourcess, usage statistics.

About cluster

Account



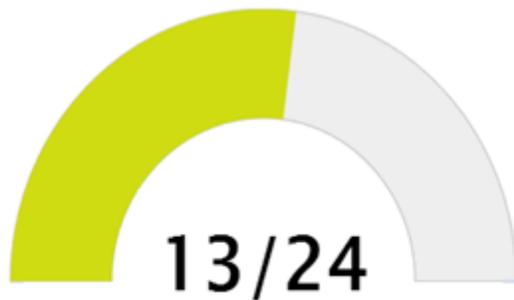
Yurii Butenko
Admin

[Profile](#)

[Logout](#)

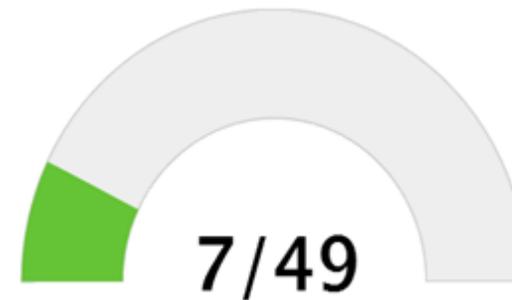
Running jobs

The ratio of number of running tasks to the total number of started tasks



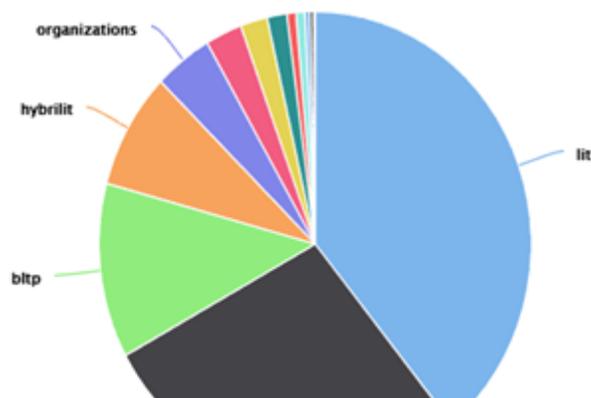
Active users

The ratio of number of users using cluster at the moment relative to active users over past week



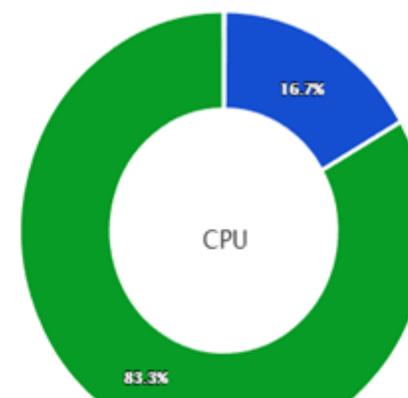
Top active organizations

Organizations most active using cluster HybriLIT for the last 6 months



Total workload of cluster HybriLIT

The average utilization of processors by system and users on all nodes.



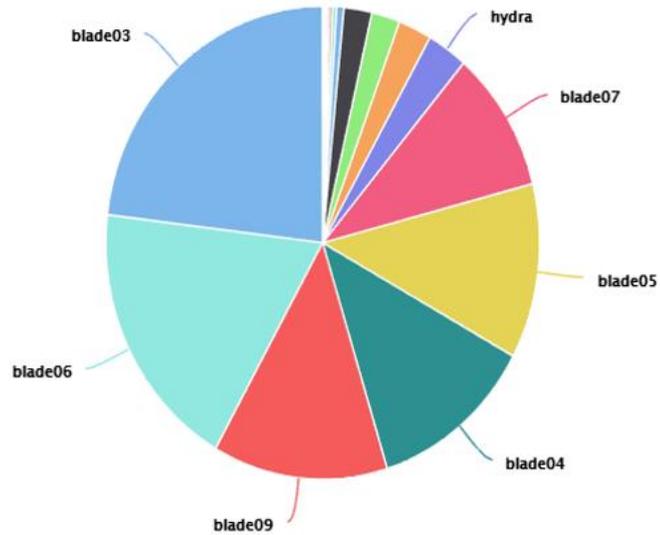
Analysis of cluster workload

Information of cluster workload over a given period of time

Date start: 22-10-2016 | Date end: | Type of nodes: | Data: CPU: User | Scale: Month | Build:

Analysis of cluster workload

Information of cluster workload over a given period of time



Usage of nodes

Information of usage for each node

Nodename ↑	Max	Min	AVG	MaxPoss
blade02	23.88 %	0.00 %	5.87 %	100.00 %
blade03	67.01 %	38.41 %	54.35 %	100.00 %
blade04	42.74 %	16.61 %	28.78 %	100.00 %
blade05	43.15 %	15.10 %	28.15 %	100.00 %
blade06	67.33 %	22.31 %	43.91 %	100.00 %
blade07	40.54 %	7.14 %	22.87 %	100.00 %

PLANS FOR 2018

- ✓ IMPROVE code. Fix memory leaks, wrong use **ZMQ**, not dynamic list of subscribers and we send info all time, and when no subscribers too (We was students in 2016 😊)
- ✓ Added new features for CPU sensor (IOwait, Load by cores)
- ✓ C++11 compatibility (We was students in 2016 😊)
- ✓ Support GPU monitoring via NVML. Thanks for Branislav Beke.
- ✓ Support multi cluster or Supercomputers

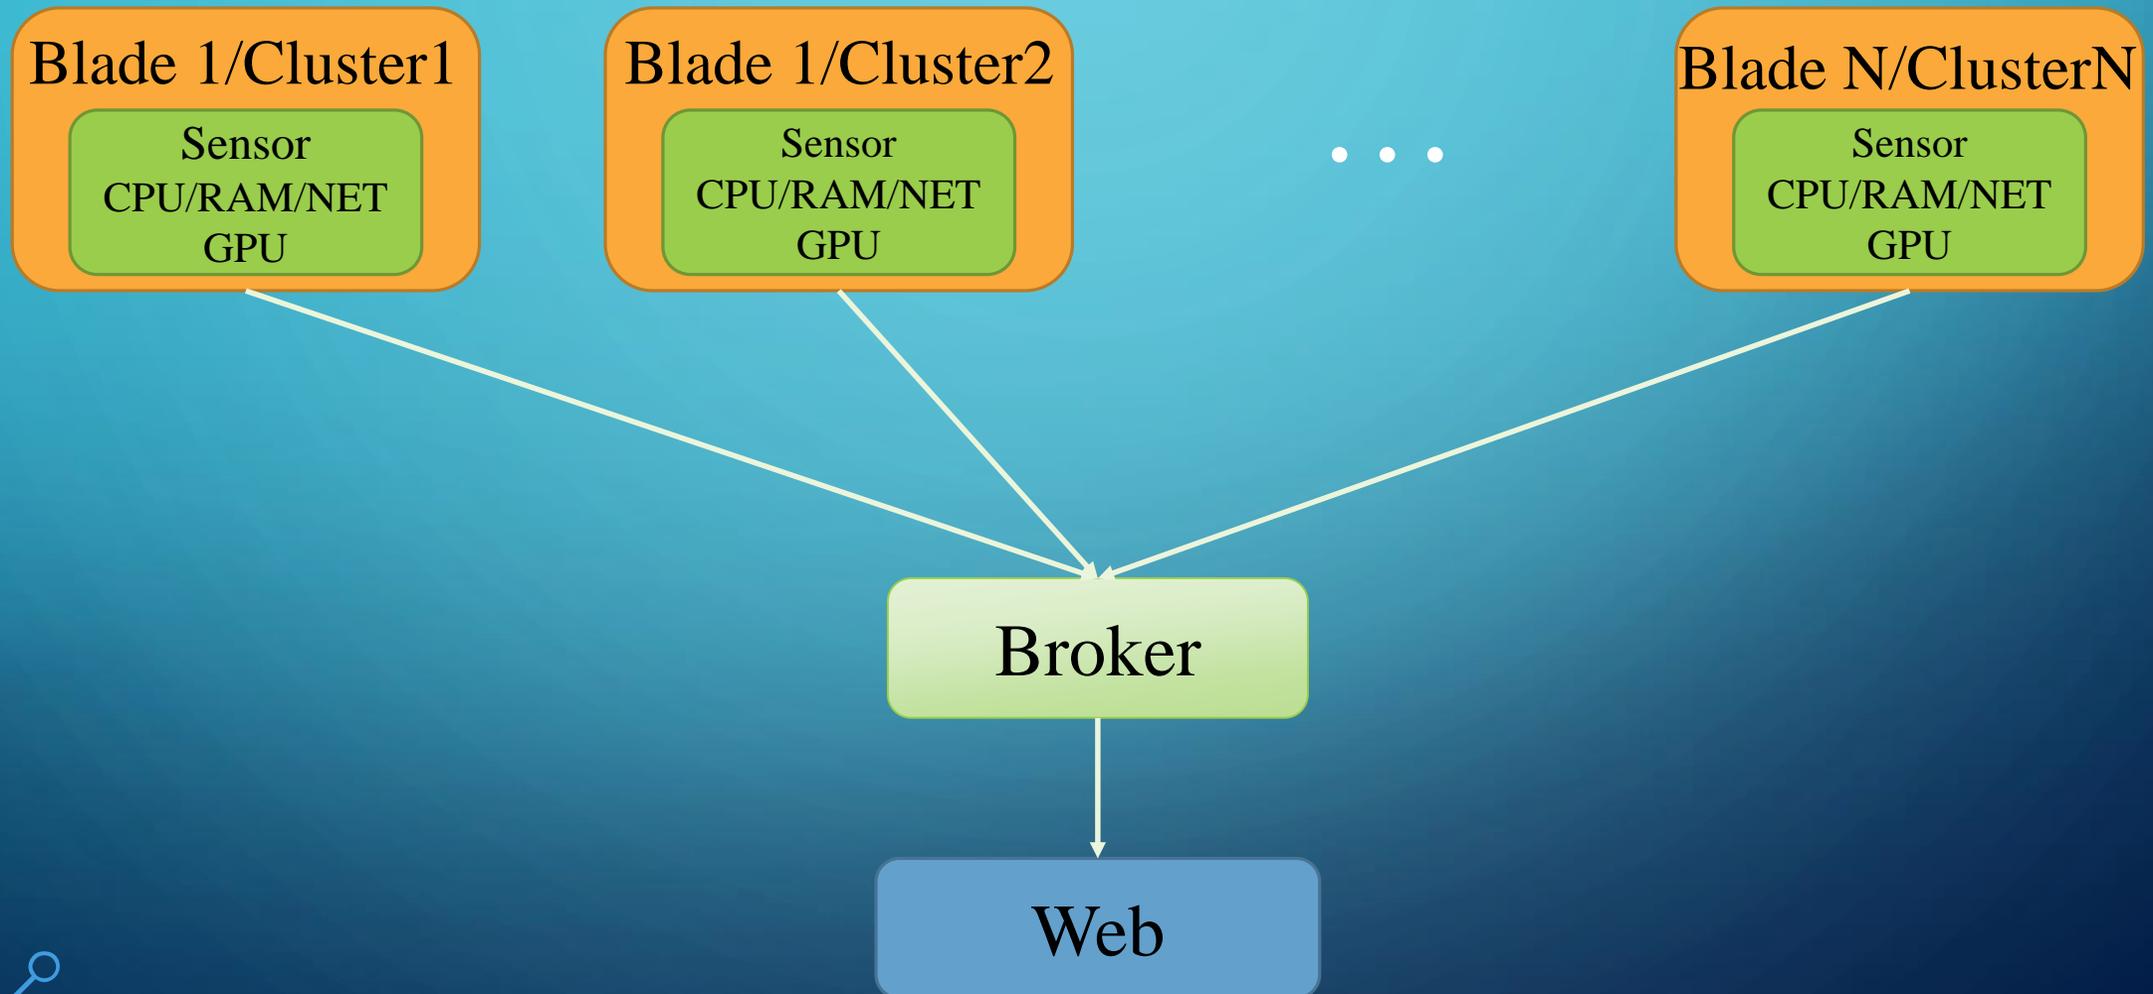
All of it done (hotovo) because now we develop our system with OB team 😊

Marian Skalský, Dmitry Podgainy, Oksana Streltsova,

Michal Dančo, Dominik Matis, Branislav Beke, Matej Fedor, Andrey Bulatov.

Thanks for ALL!!

OB-MON STRUCTURE



WORK STILL IN PROGRESS....

- Now we implement our monitoring app in SALSA and now it can be installed not only on HybriLIT or SC Govorun.
- We are moving to React. Because every new major version of Angular brings many problems and we solve problems with versions and syntax. We want develop new features! not solve 3rd party library problems...
(Thanks for Michal Dančo, Dominik Matis)
- Everyday we want to do something cool and useful

<https://salsa.openbrain.sk/salsa-webapp/>

hlit

saske

eos

gvr:dgx

gvr:kn1

gvr:skylake

gvr:wn

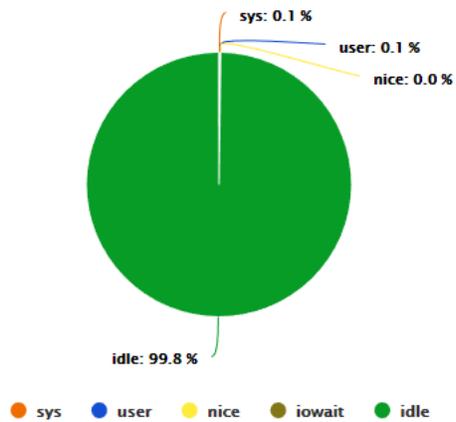
hlit:wn

service

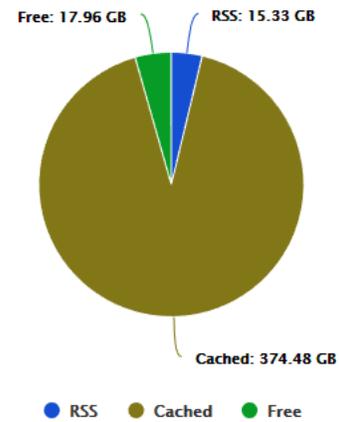
vm

Sliced Diagram 

eos CPU



eos Memory

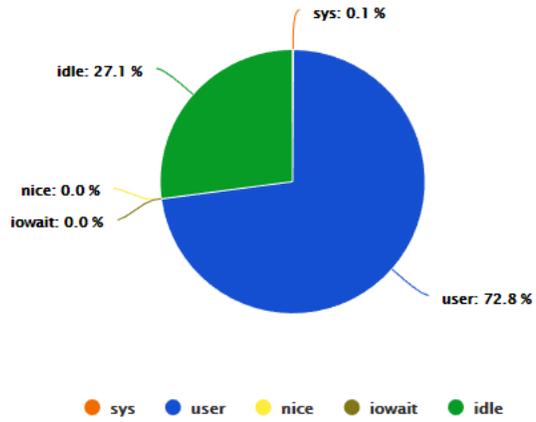


- eos
- gvr:dgx
- gvr:kn1
- gvr:skylake**
- gvr:wn
- hlit
- hlit:wn
- service
- vm
- saske

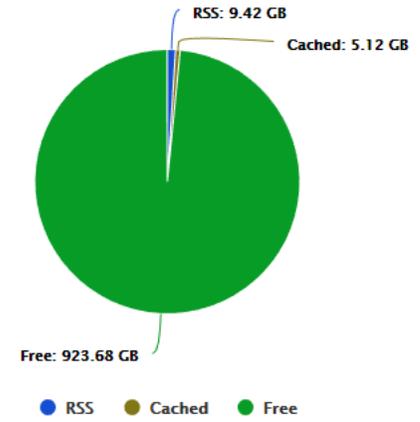
🔄 📊

Sliced Diagram

gvr:skylake CPU



gvr:skylake Memory

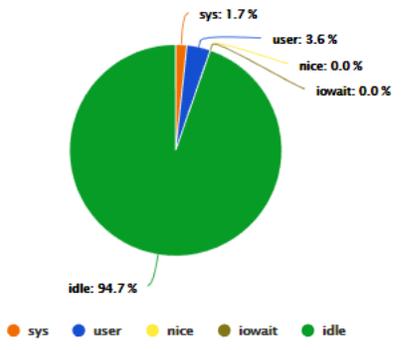


- hlit
- saske
- gvr:dgx
- gvr:kn1
- gvr:skylake
- gvr:wn
- hlit.wn
- service
- vm
- eos

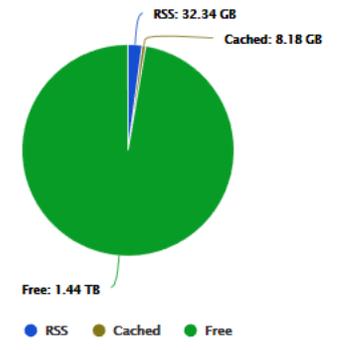
Sliced Diagram



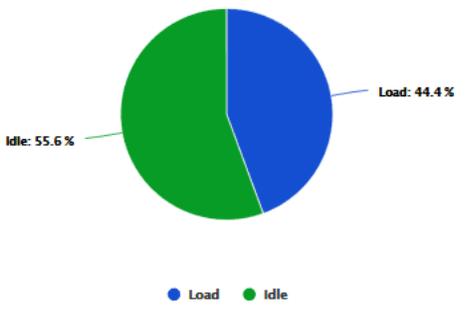
gvr:dgx CPU



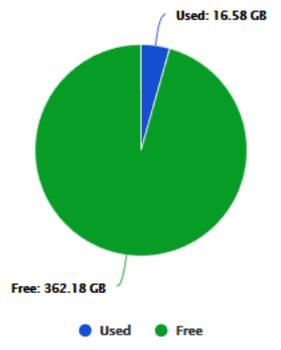
gvr:dgx Memory



gvr:dgx GPU



gvr:dgx GPU Memory



eos
gvr:dgx
gvr:kn1
gvr:skylake
gvr:wn
hlit
saske
hlit:wn
service
vm

🔄
🏠

Sliced Diagram

				CPU					Memory			Network		GPU			
id	name	cores	load	sys	user	nice	iowait	idle	used	cached	total	In	Out	GPU's	Usage	Used mem.	Total mem.
1	dgx01	80	4.15	1.6 %	3.6 %	0.0 %	0.0 %	94.8 %	10.83 GB	2.92 GB	503.80 GB	556 B/s	5.23 KB/s	8	43.9 %	4.81 GB	126.25 GB
2	dgx02	80	4.11	1.8 %	3.6 %	0.0 %	0.0 %	94.6 %	11.01 GB	2.61 GB	503.80 GB	5.90 KB/s	10.06 KB/s	8	44.9 %	5.38 GB	126.25 GB
3	dgx03	80	4.14	1.8 %	3.5 %	0.0 %	0.0 %	94.7 %	10.49 GB	2.65 GB	503.80 GB	498 B/s	5.44 KB/s	8	44.4 %	6.39 GB	126.25 GB
Total		240	12.40	-	-	-	-	-	32.34 GB	8.18 GB	1.48 TB	6.93 KB/s	20.72 KB/s	-	-	-	-
Average		80.0	4.13	1.7 %	3.6 %	0.0 %	0.0 %	94.7 %	10.78 GB	2.73 GB	503.80 GB	2.31 KB/s	6.91 KB/s	-	-	-	-

hlit saske
alice-eos service

🔄 📊

Sliced Diagram

				CPU					Memory			Network		GPU			
id	name	cores	load	sys	user	nice	iowait	idle	used	cached	total	In	Out	GPU's	Usage	Used mem.	Total mem.
1	eosf02-iep-grid	6	0.26	1.0%	0.2%	0.0%	0.0%	98.3%	1016.03 MB	14.12 GB	15.48 GB	2.42 MB/s	62.91 MB/s	-	-	-	-
2	eosf03-iep-grid	6	1.03	1.5%	0.2%	0.0%	1.2%	96.8%	2.49 GB	12.73 GB	15.48 GB	2.82 MB/s	71.95 MB/s	-	-	-	-
3	eosm01-iep-grid	6	0.25	0.2%	0.3%	0.0%	0.0%	99.5%	20.21 GB	10.78 GB	31.39 GB	15.95 KB/s	3.44 KB/s	-	-	-	-
Total		18	1.54	-	-	-	-	-	23.70 GB	37.63 GB	62.35 GB	5.25 MB/s	134.86 MB/s	-	-	-	-
Average		6.0	0.51	0.9%	0.2%	0.0%	0.4%	98.2%	7.90 GB	12.54 GB	20.78 GB	1.75 MB/s	44.95 MB/s	-	-	-	-

BUT WE WANT NOT ONLY MONITORING FOR USER APPS ON HLIT AND GOVORUN

- We work on implementation with Ivan Kashunin's monitoring (litmon) for HybriLIT and SC Govorun
- We already install all NRPE packages and configure all blades on HybriLIT for work with litmon

Tier-1 tape library



changer operational status **ok**

media access operational status **ok**

library computer status **ok**

CICC COMMON SERVERS

t2-s730-734

t2-s735-737

JINR TIER-2 SERVICE

LCGS SERVERS

lcvob0[2-3]

lcsedc01

lcsqd10

PANDA SERVERS

lxpanda

lxvm-h1[0-2]

TIER 1 COMMON SERVERS

t1-s154-158

t1-s159-163

t1-s164-168

t1-s169-173

t1-s174-174

CICC COMMON SERVERS

l4u-00

lit-n

lxadm

lxafs0[0-4]

lxafs0[5-6]

lxbkp01

lxdbs-n

lxdevel

lxmail

lxmpd-s1

lxs-s0[3,9]

n3w4lit

rocmon

topmon

CICC COMMON SERVERS

rda01-04

rda05-09

rda10-13

JINR CLOUD

COMMON SERVERS

cloudfs1-3

cwn41-44

cwn45-49

cwn50-54

cwn57-60

cwn61-63

cwn66-70

cwn72-76

cwn77-81

csn021-024

csn026-027

statistics

lt-qt1-vm1-4

localhost

lt-vm4

COMMON SERVERS

JINR HYBRLIT

COMMON SERVERS

hltmon_space20

store1-4

dgx01-05

blade1-5

blade6-10

titan

unit1-2

HYPERVISORS

saber01-05

saber06-08

ovirt1-4

UPS



OUR PLANS

- HARDWORK
- More deeply integration with SALSA-webapp and full join forces to develop
 - Add new features and optimize current
 - Present a universal product that can be customized for your targets/ideas/needs

Hardwork. Hardwork. Hardwork. Hardwork. Hardwork. Hardwork. Hardwork. Hardwork. Hardwork.

The image features a blue gradient background with white circuit-like lines in the corners. The lines consist of straight segments connected by small circles, resembling a network or data flow diagram. The text "Thank you for attention!" is centered in a white, sans-serif font.

Thank you for attention!