Networking & perfSONAR
Task 2.4

Rizart Dona

CERN

November 20, 2019 - WP2 Fortnightly Meeting, CERN

ESCAPE - The European Science Cluster of Astronomy & Particle Physics ESFRI Research Infrastructures has received funding from the European Union's Horizon 2020 research and innovation programme under the Grant Agreement n° 824064.

Overview

- The perfSONAR system
- Datalake perfSONAR status
- Configuration tools
- Monitoring tools
- Future work
- References

Funded by the European Union's Horizon 2020 - Grant N° 824064

The perfSONAR system

- perfSONAR is a collection of open source software for performing and sharing end-to-end network measurements
- It consists of multiple tools brought together (iperf, ping, traceroute, nuttcp, owping, and other..)
- WLCG and OSG (Open Science Grid) jointly operate a network of perfSONAR agents deployed world-wide, which provides an open platform that can be used to baseline network performance and debug any potential issues

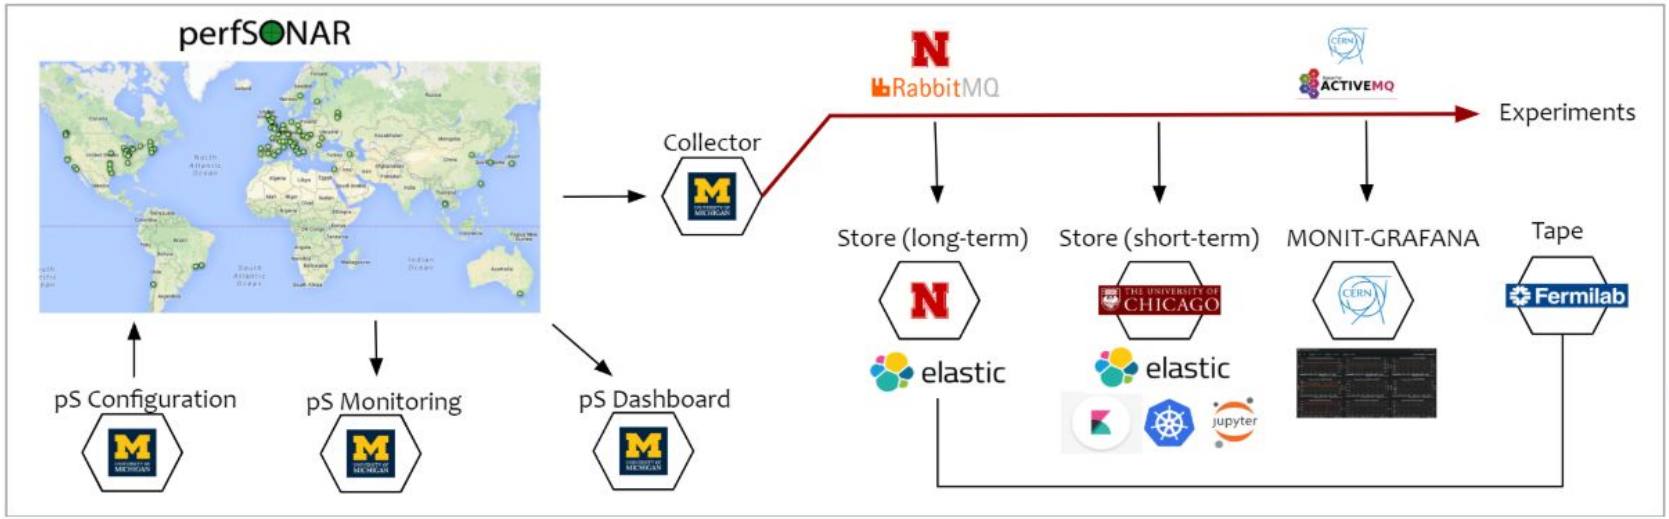

Funded by the European Union's Horizon 2020 - Grant N° 824064

Datalake perfSONAR status

Institute	Technology	Version	Quota	XROOT	HTTPS	GSIFTP	ESCAPE VO enabled (X509)	Token based auth/z	perfSONAR Enabled	perfSONAR Hosts (bandwidth & latency)
CERN	EOS	?	?	YES	YES	YES	WiP	NO	YES	psb01-gva.cern.ch psl01-gva.cern.ch
LAPP	DPM (DOME)	1.13.1	?	YES	YES	YES	WiP	?	YES	lapp-ps01.in2p3.fr lapp-ps02.in2p3.fr
SURFsara	dCache	?	?	YES	YES	YES	YES	?	YES	perfsonar-bandwidth.grid.surfsara.nl perfsonar-latency.grid.surfsara.nl
LAPP	dCache	?	?	NO	YES (http)	NO	YES	?	YES	lapp-ps01.in2p3.fr lapp-ps02.in2p3.fr
GSI	xRootD	?	?	?	?	?	?	?	NO	? ?
INFN-CNAF	StoRM	?	?	NO	YES	NO	YES	YES	YES	perfsonar-ps.cnaf.infn.it perfsonar-ow.cnaf.infn.it
INFN-ROMA1	DPM	?	?	WiP	WiP	WiP	WiP	?	YES	perfsonar1.roma1.infn.it perfsonar2.roma1.infn.it
DESY	dCache	?	?	YES	YES	YES	YES	?	YES	perfsonar-ps-04.desy.de perfsonar-ps-03.desy.de
CC-IN2P3	dCache	?	60 TB	YES	YES	YES	YES	?	YES	ccperfsonar1.in2p3.fr ccperfsonar2.in2p3.fr
PIC	dCache	?	?	YES	YES	?	YES	?	YES	psb01.pic.es psl01.pic.es

Funded by the European Union's Horizon 2020 - Grant N° 824064

perfSONAR - Central Configuration (Hosts)

- Some **hosts** are dedicated for <u>latency</u> testing and others for <u>bandwidth</u> testing

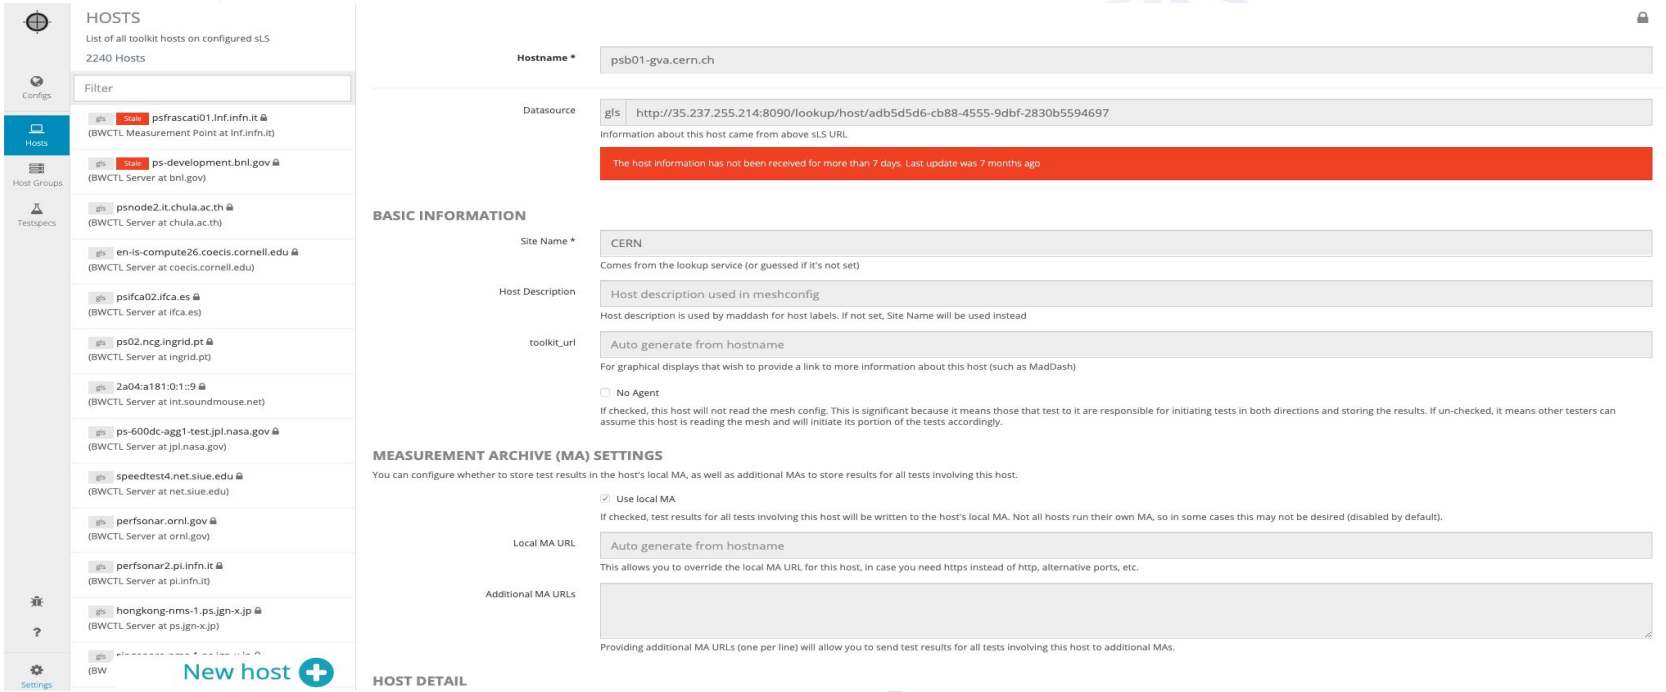

Funded by the European Union's Horizon 2020 - Grant N° 824064

perfSONAR - Central Configuration (Hostgroups)

- Each **hostgroup** contains some **hosts**

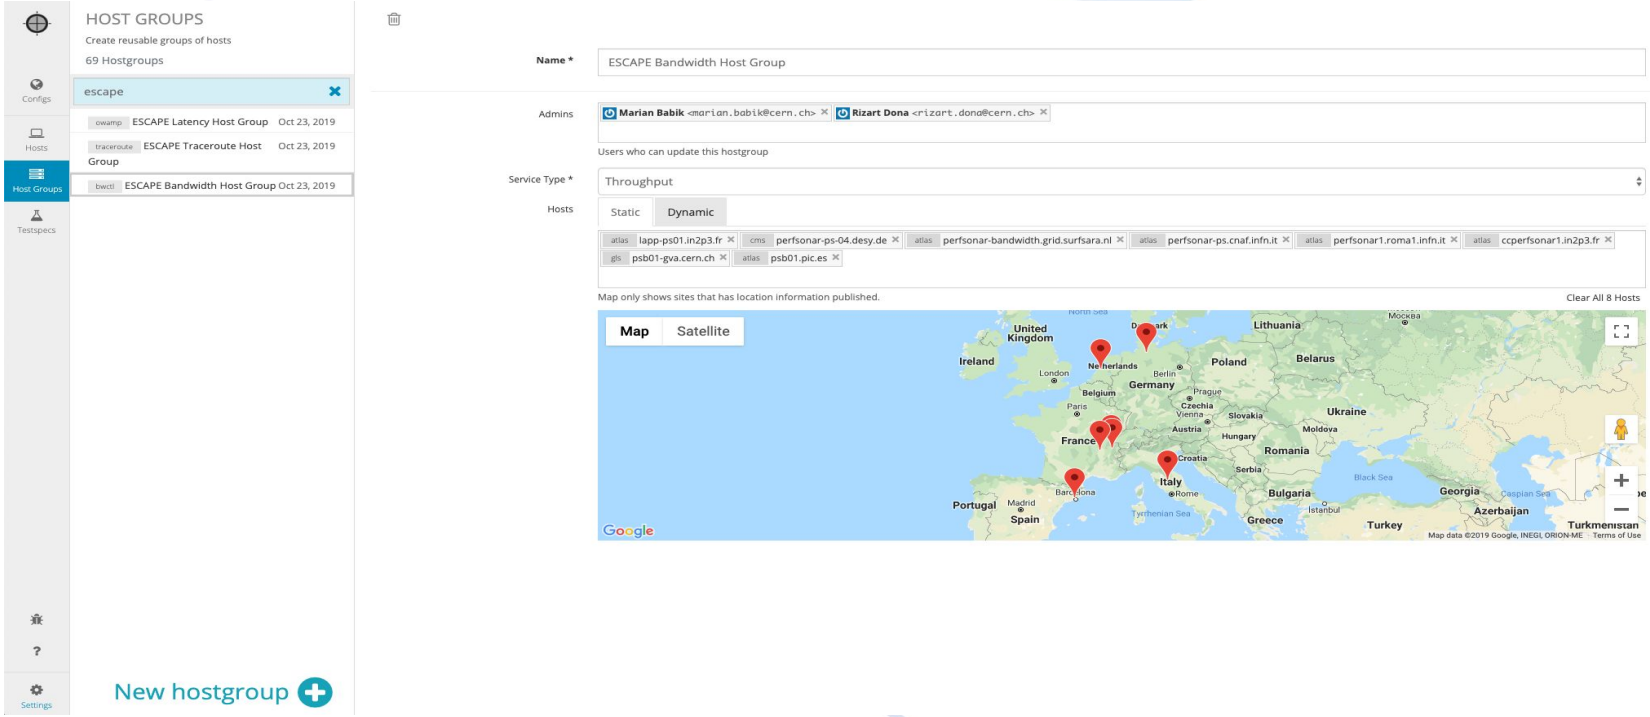

Funded by the European Union's Horizon 2020 - Grant N° 824064

perfSONAR - Central Configuration (Mesh)

- Each **mesh** configuration can have different **tests** (latency, traceroute, throughput, ping)
- Each **test** is connected with a **hostgroup**

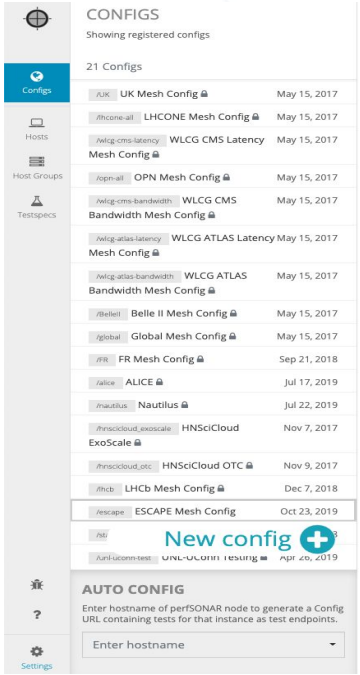

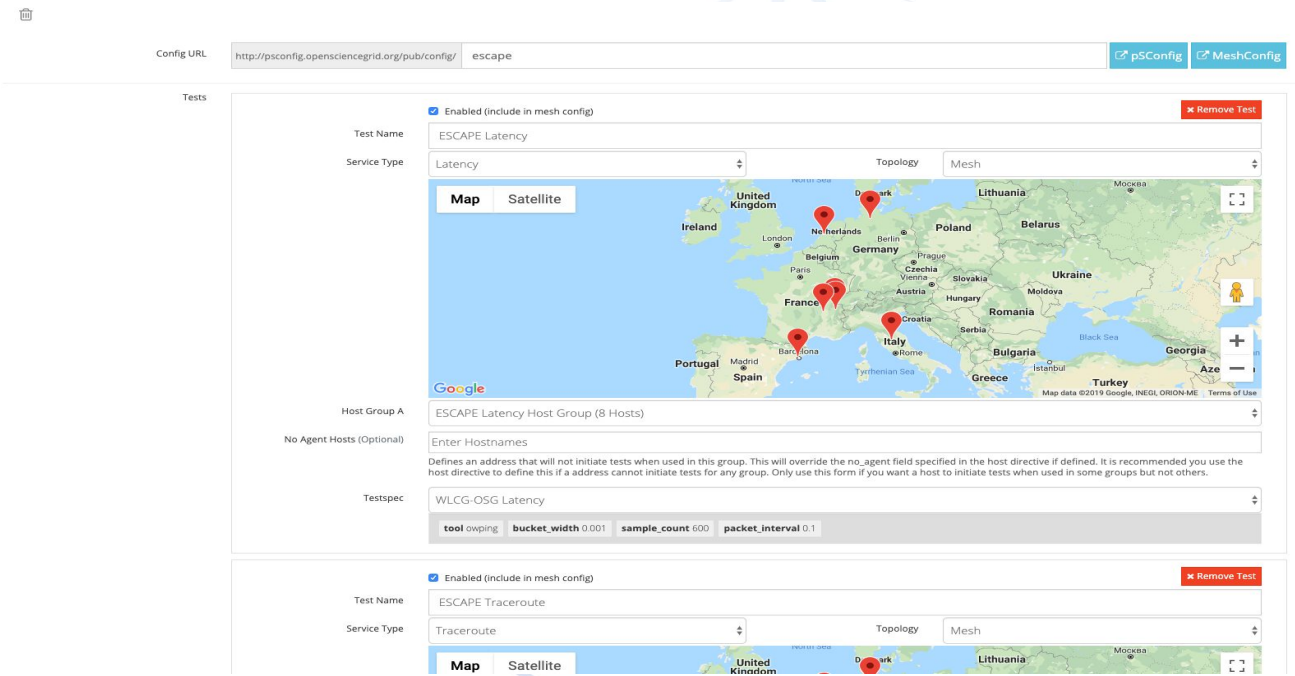

Funded by the European Union's Horizon 2020 - Grant N° 824064

Monitoring - Check_MK (Infrastructure Monitoring)

- A portal to check individual perfSONAR boxes and the status of theirs services

state	Host	Icons	OK	Wa	Un	Cr	Pd
UP	perfsonar2.roma1.infn.it		2	0	8	4	0
UP	cmsrm-perfsonar1.roma1.infn.it		13	0	2	1	0
UP	ccperfsonar1.in2p3.fr		14	1	0	0	0
UP	ccperfsonar2.in2p3.fr		14	0	0	0	0
UP	lapp-ps01.in2p3.fr		14	1	0	0	0
UP	lapp-ps02.in2p3.fr		13	1	0	0	0
UP	perfsonar1.roma1.infn.it		14	1	0	0	0
UP	perfsonar-bandwidth.grid.surfsara.nl		14	1	0	0	0
UP	perfsonar-latency.grid.surfsara.nl		13	1	0	0	0
UP	perfsonar-ow.cnaf.infn.it		13	1	0	0	0
UP	perfsonar-ps-03.desy.de		14	0	0	0	0
UP	perfsonar-ps-04.desy.de		15	0	0	0	0
UP	perfsonar-ps.cnaf.infn.it		14	1	0	0	0
UP	psb01-gva.cern.ch		14	1	0	0	0
UP	psb01.pic.es		13	2	0	0	0
UP	psl01-gva.cern.ch		13	1	0	0	0
UP	psl01.pic.es		13	1	0	0	0

psb01-gva.cern.ch

State	Service	Icons	Status detail	Age	Checked	Perf-O-Meter
OK	perfSONAR configuration: contacts		OK - Contact and organization found	2019-11-15 10:35:23	59 m	
OK	perfSONAR configuration: location		OK - Location: 6.0457234/46.2324045	2019-11-15 10:35:23	59 m	
OK	perfSONAR configuration: meshes		OK - Auto-URL configured correctly	2019-11-15 10:35:23	59 m	
WARN	perfSONAR esmond freshness: bwctl		WARNING - esmond 63.25 % complete for event type: throughput	35 h	68 m	
OK	perfSONAR esmond freshness: trace		OK - esmond 97.93 % complete for event type: packet-trace	43 h	64 m	
OK	perfSONAR hardware check		OK - CPU:2/16cores/3000.0000Mhz RAM:126GB NIC:40Gbps/1500MTU/IPv6 enabled	2019-11-15 10:35:23	59 m	
OK	perfSONAR json summary		OK - Toolkit metadata successfully retrieved	2019-11-15 10:35:18	59 m	
OK	perfSONAR services: bwctl		connect to address psb01-gva.cern.ch and port 4823: Connection refused	2019-11-15 10:40:19	7 m	
OK	perfSONAR services: http/https		OK - Toolkit homepage reachable	2019-11-15 10:35:23	59 m	
OK	perfSONAR services: ndt/npad disabled		OK - NDT/NPAD disabled	2019-11-15 10:35:23	59 m	
OK	perfSONAR services: ntp		OK - NTP synchronized	2019-11-15 10:35:23	59 m	
OK	perfSONAR services: pscheduler		OK - pscheduler stats retrieved	2019-11-15 10:15:17	18 m	
OK	perfSONAR services: regular testing/pscheduler		OK - pscheduler is enabled and running	2019-11-15 10:35:23	59 m	
OK	perfSONAR services: versions		OK - Toolkit version found: 4.2.2-1.el7	2019-11-15 10:35:23	59 m	
OK	perfSONAR services: web/https IPv6		TCP OK - 0.111 second response time on psb01-gva.cern.ch port 443	2019-11-15 10:20:25	12 m	

Funded by the European Union's Horizon 2020 - Grant N° 824064

Monitoring - MadDash

- The Monitoring and Debugging Dashboard (MaDDash) software is aimed at collecting and presenting two-dimensional monitoring data as a set of grids referred to as a dashboard.

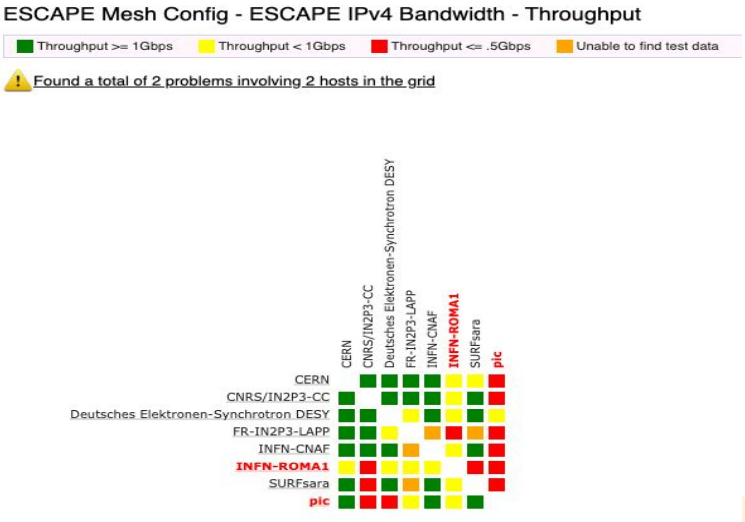

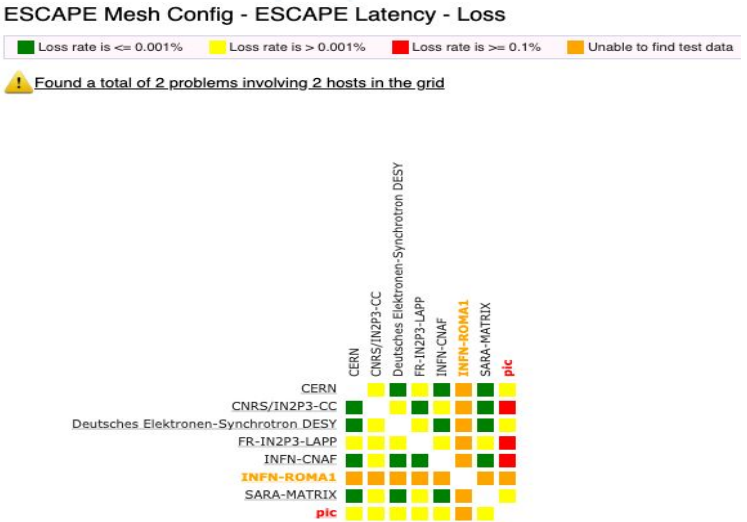

Funded by the European Union's Horizon 2020 - Grant N° 824064

Monitoring - Monit Grafana (I)

- CERN Monit Grafana instance to visualize network statistics concerning LHCOPN, LHCONE and perfSONAR sites that participate in the OSG (Open Science Grid)
- Already provides some statistics for datalake hosts

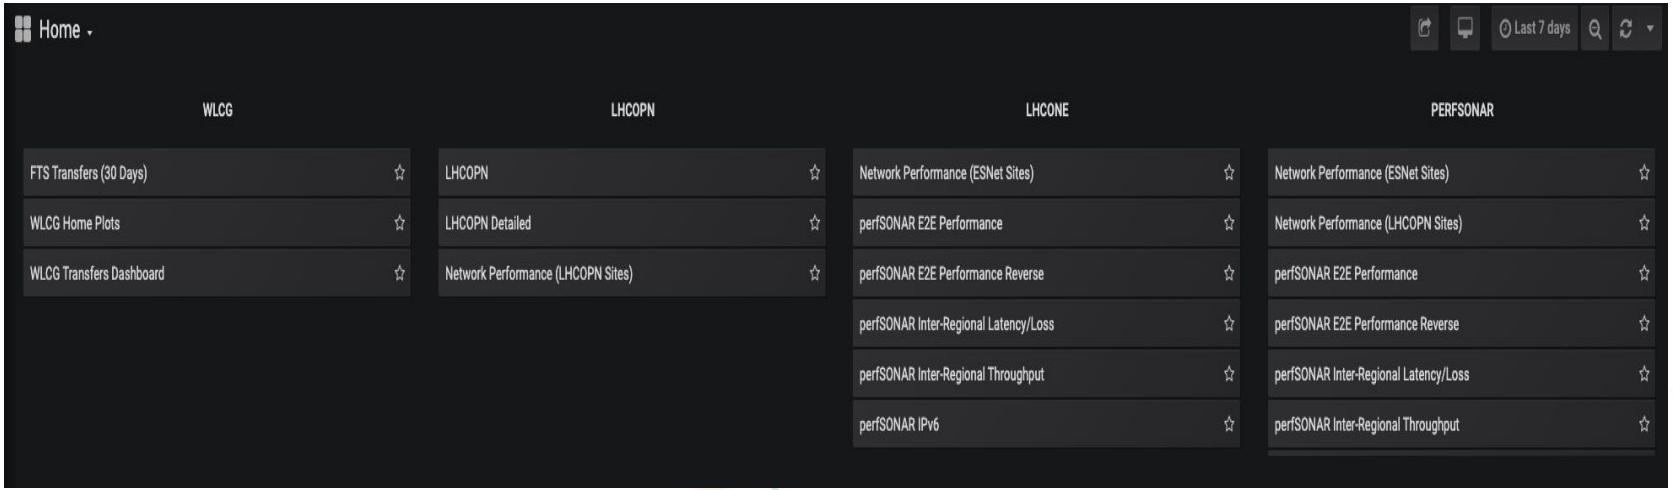

Funded by the European Union's Horizon 2020 - Grant N° 824064

Monitoring - Monit Grafana (II)

Rizart.Dona@cern.ch

Funded by the European Union's Horizon 2020 - Grant N° 824064

Future Work

- Datalake sites need to confirm that the current perfSONAR hosts are correct
  - See Datalake status
- Remaining datalake sites need to setup the perfSONAR hosts in order to join the ESCAPE mesh tests (GSI, etc..)
- Grant admin access to relevant parties in the perfSONAR Central Configuration platform for the ESCAPE mesh configuration
  - SKA & GSI
- Aggregated real time visualization of network monitoring data
  - We need to specify what and how will be visualized

Funded by the European Union's Horizon 2020 - Grant N° 824064

References (I)

- perfSONAR (https://docs.perfsonar.net/index.html)
- perfSONAR Web Admin (https://psconfig.opensciencegrid.org)
- Open Science Grid Netowking (https://opensciencegrid.org/networking/)
- HEPiX Fall 2019 OSG/WLCG Networking Update (https://docs.google.com/presentation/d/1DnpVpwuOyx464nmtvdIQoxreT7-qCqGERkuu7PQ9F20/edit#slide=id.p)
- LHCONE Wiki (https://twiki.cern.ch/twiki/bin/view/LHCONE/WebHome)
- LHCOPN Wiki (https://twiki.cern.ch/twiki/bin/view/LHCOPN/WebHome)
- CHECK_MK Monitoring (https://psetf.aglt2.org/etf/check_mk/index.py)
- Grafana Network Monitoring (http://monit-grafana-open.cern.ch)

Funded by the European Union's Horizon 2020 - Grant N° 824064

References (II)

- ESCAPE Datalake status (https://wiki.escape2020.de/index.php/WP2_-_DIOS#Datalake_Status)
- ESCAPE Hosts MadDash Monitoring (http://maddash.aglt2.org/maddash-webui/index.cgi?dashboard=ESCAPE%20Mesh%20Config)
- ESCAPE Hosts CHECK_MK Monitoring (https://psetf.aglt2.org/etf/check_mk/view.py?hostgroup=ESCAPE&opthost_group=ESCAPE&view_name=hostgroup)

Funded by the European Union's Horizon 2020 - Grant N° 824064