



CERN IT Monitoring Service

Nikolay Tsvetkov

26.10.2022

Service Catalogue

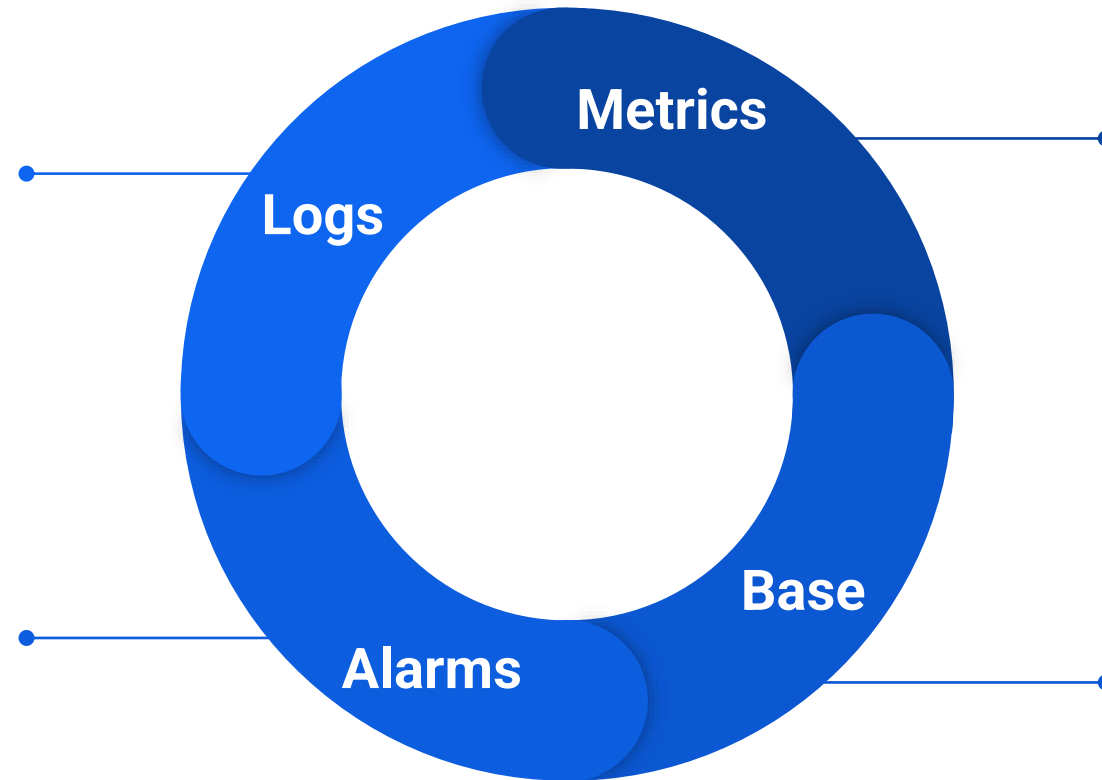
Service Logs

Integration of service logs via a central HTTP endpoint from any compatible client. Support for public and private logs.

Service Alarms

Integration of service alarms from Collectd, Prometheus, or custom tools.

Multiple targets provided: ServiceNow, Grafana, email



Service Metrics

Integration of service metrics from Collectd, Prometheus, or custom tools.

Remote probes, service availability, and service SLIs.

Base Monitoring

Out-of-the-box monitoring for Puppet managed nodes and Openstack K8s clusters: Collectd metrics and alarms, NoContact alarms, and Syslog.

Service Catalogue

Out-of-the-box

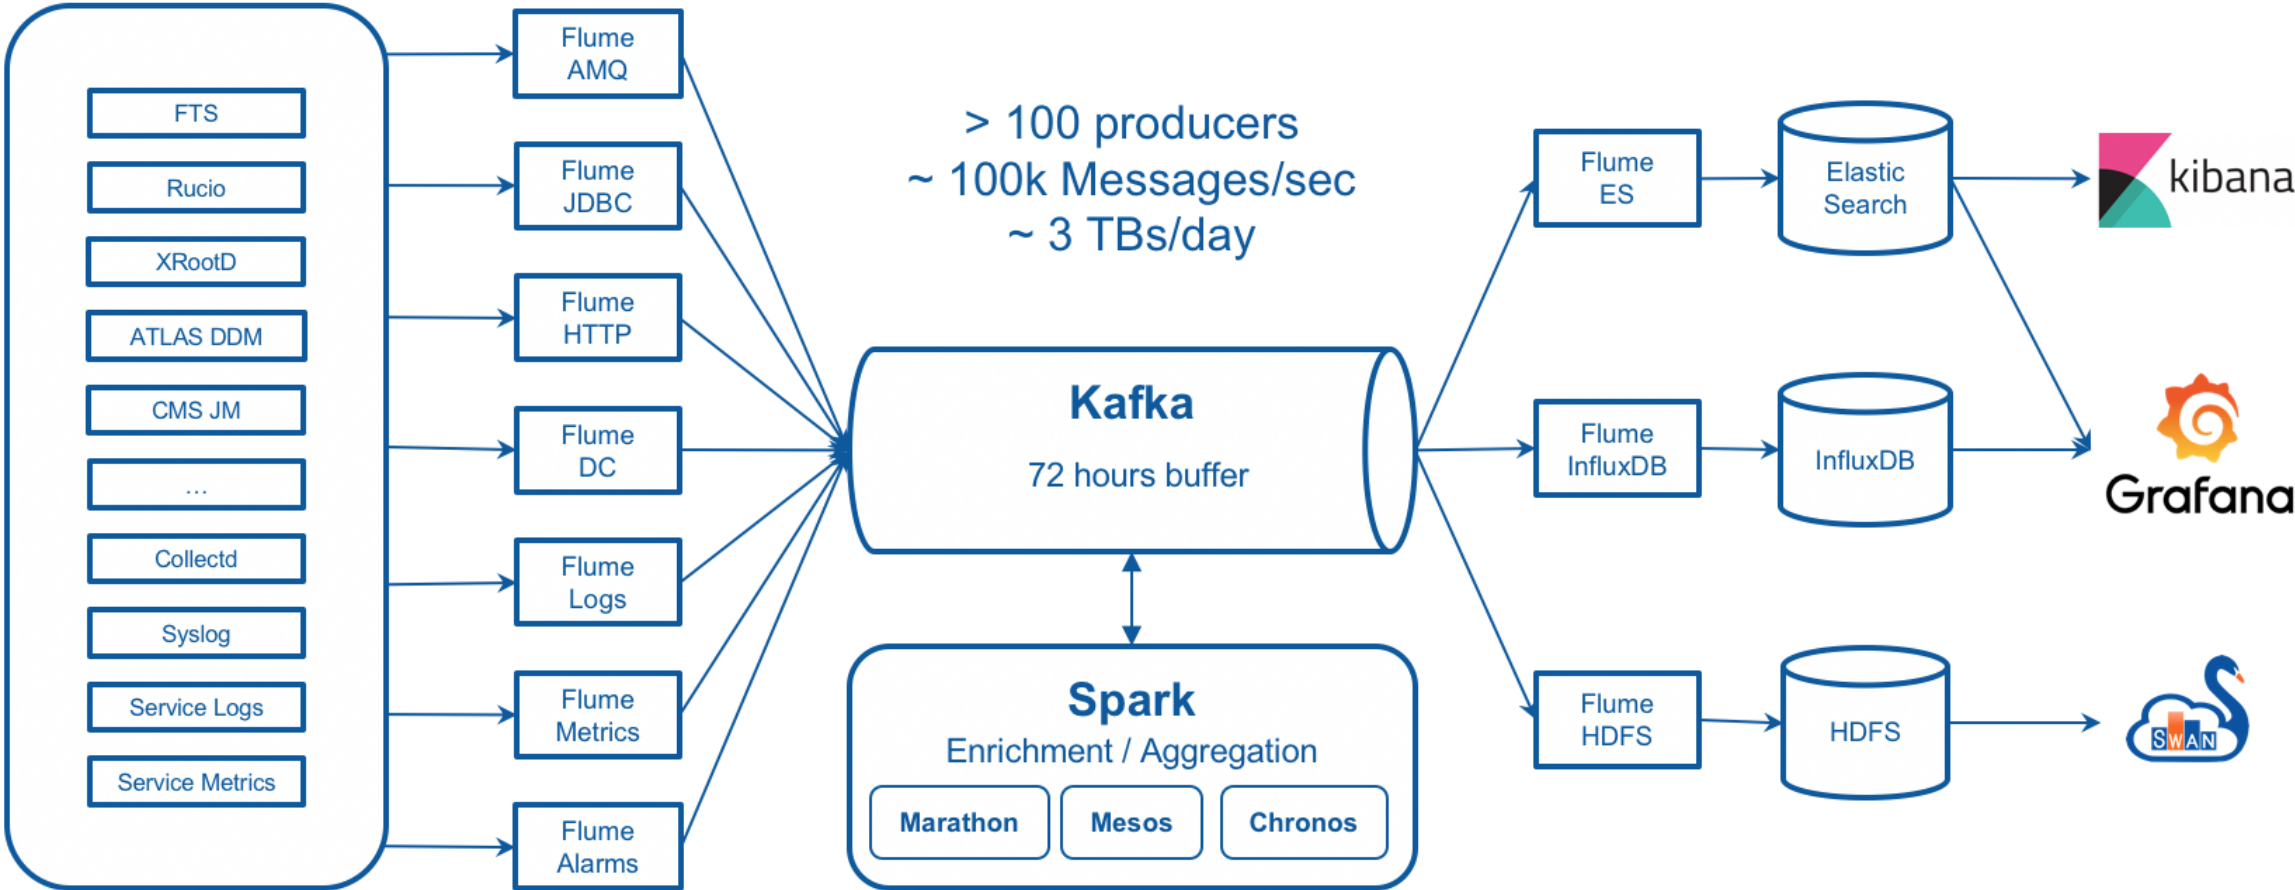
Base Monitoring	Integration	Backend	Access
Puppet nodes metrics	Collectd	InfluxDB	Grafana
Puppet nodes alarms	Collectd, NoContact	InfluxDB, Cortex	Grafana, SNOW, Email
Puppet nodes logs	Syslog	ElasticSearch	Grafana, Kibana
Kubernetes metrics	Prometheus	InfluxDB, Cortex	Grafana

Service Setup

Service Monitoring	Integration	Backend	Access
Service Metrics	Collectd, Prometheus, Custom	InfluxDB, ES, Cortex	Grafana
Service Alarms	Collectd, Prometheus, Grafana	InfluxDB, ES	Grafana, SNOW, Email
Service Logs	Flume, Logstash, FluentBit (HTTP)	ElasticSearch	Grafana, Kibana
Service Availability	Custom tools (HTTP)	InfluxDB	Grafana
Service Level Indicators	Gitlab repo	InfluxDB	Grafana
Service Remote Probes	Gitlab repo	Cortex	Grafana

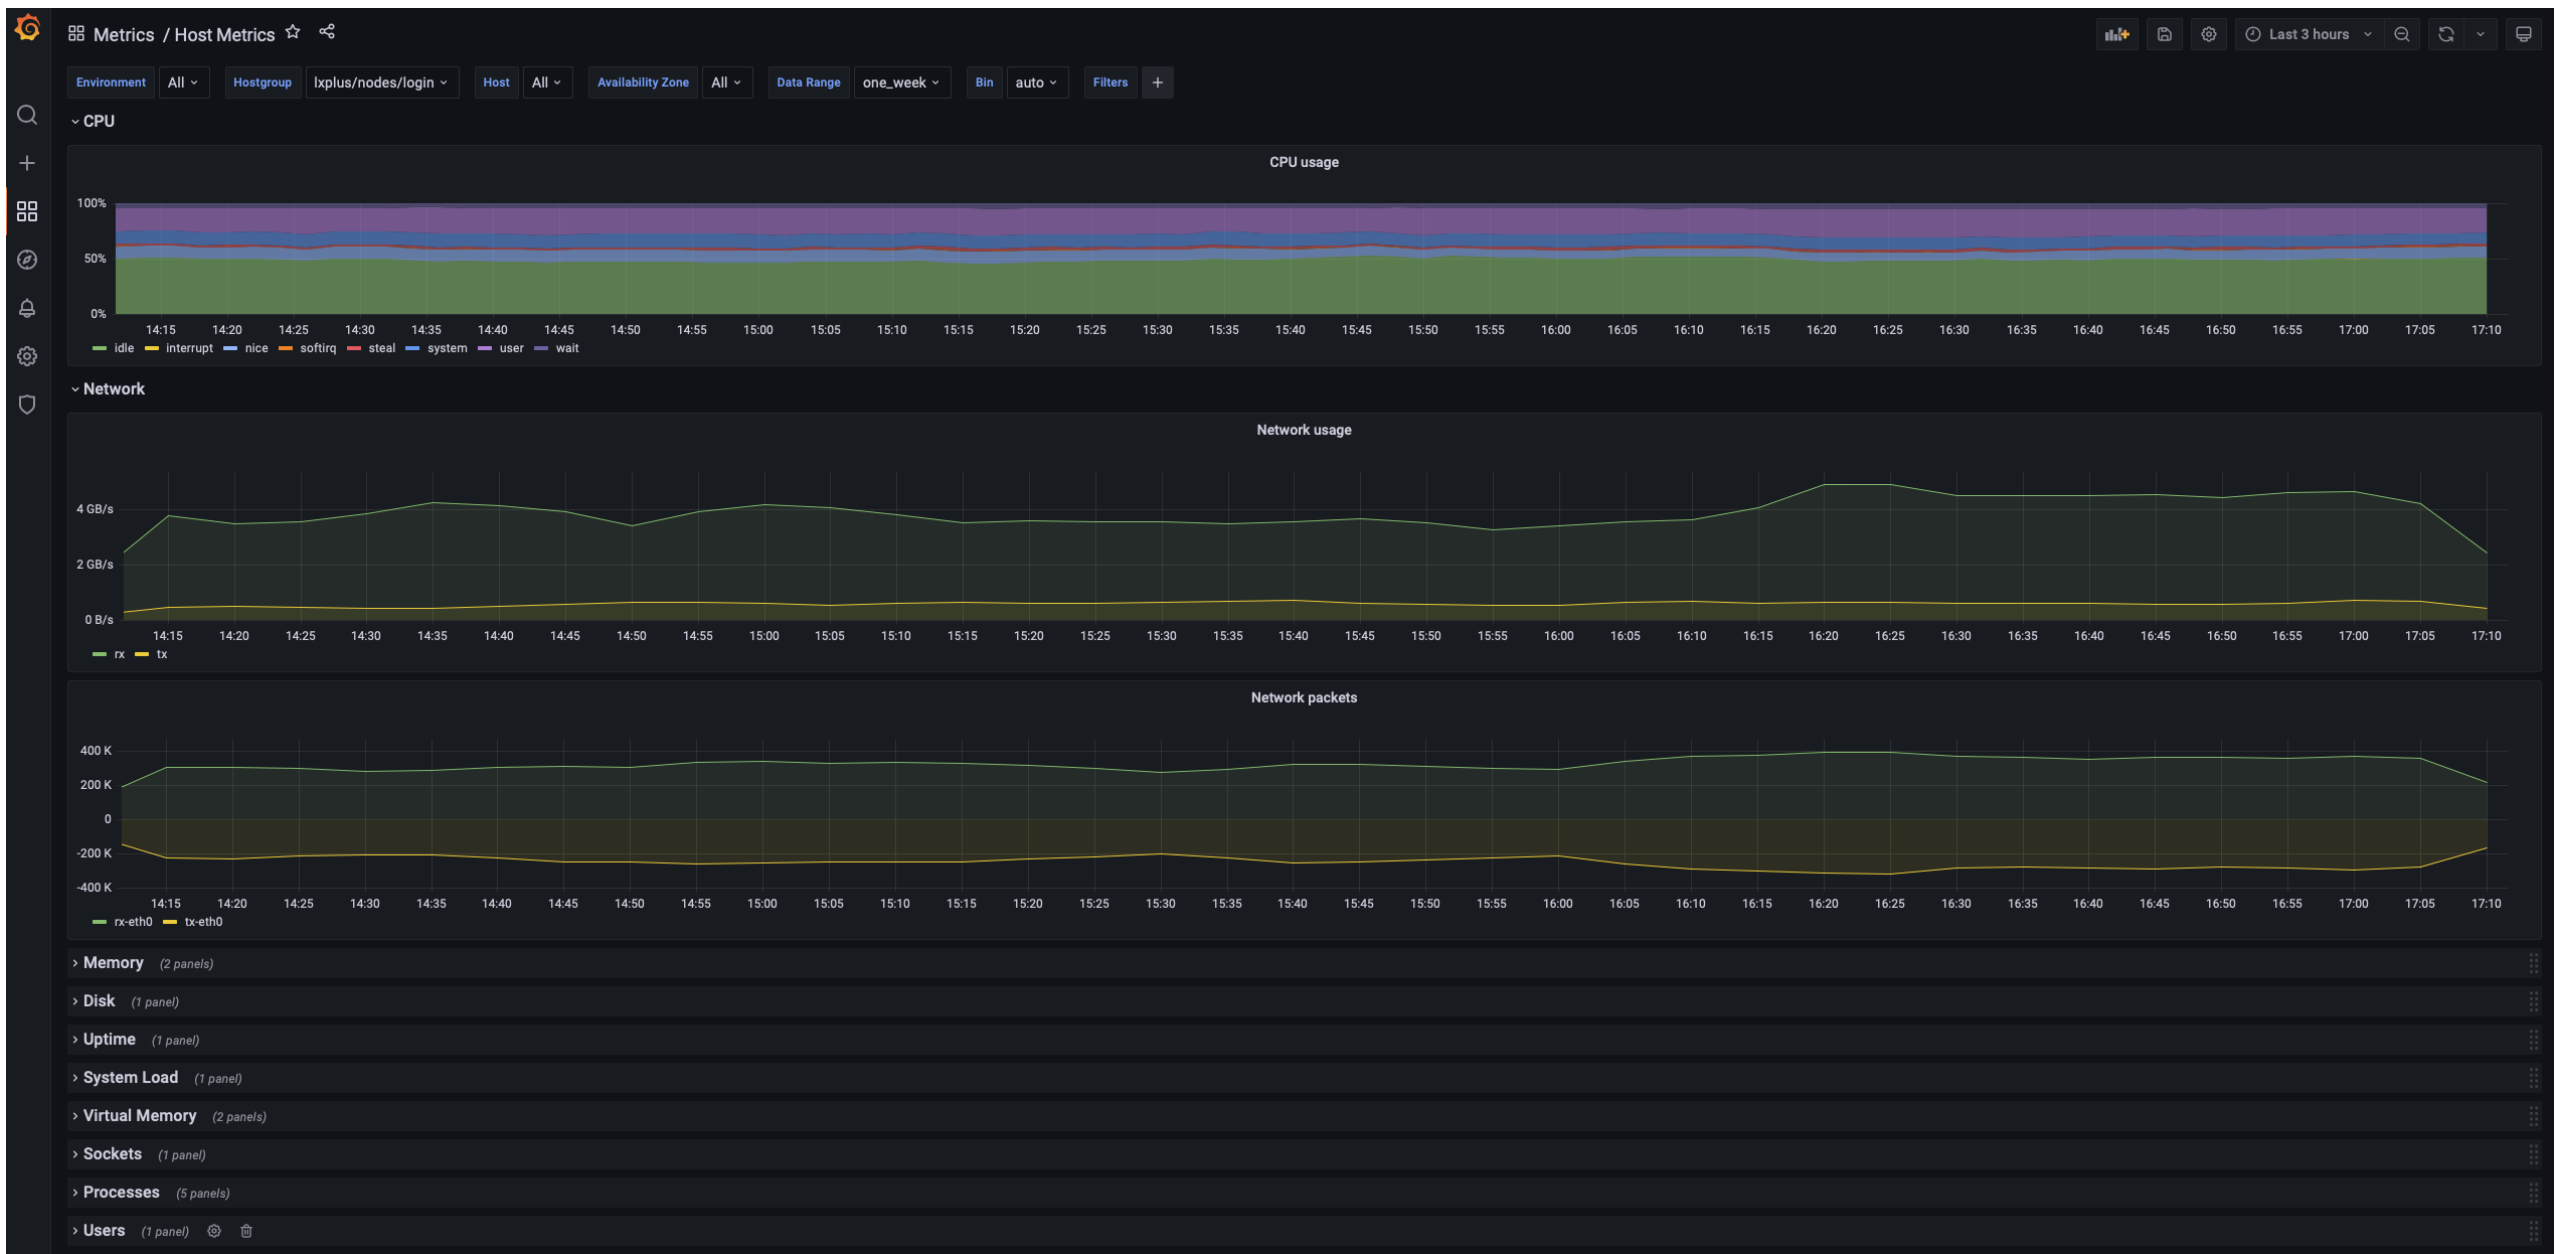
MONIT Architecture

Sources > Transport > (Processing) > Storage > Access



MONIT DC Monitoring

- **Local DC host agents:**
 - Collectd: Lightweight daemon for metrics collection
 - Flume: Java agent to forward documents outside the node (metrics, syslog, alarms)
 - NoContact heartbeat: Small Python script to generate heartbeat to MONIT
- **Data storage:**
 - InfluxDB with different resolution based on retention policies (1 min/1 week, 5 min/1 month, 1 hour/5 years)
 - HDFS for long term archive (data stored in raw format)
- **Visualization:**
 - Grafana for dashboards and real-time monitoring
 - SWAN for long-term data analysis from HDFS



MONIT Service Metrics & Logs Monitoring

- **Data producers:**

- Producers are under the user responsibility. MONIT requires JSON documents
 - ATLAS and CMS are main users among other IT services (e.g. Network)
- Logstash, FluentD, FluentBit are the most used for logs export

- **Data storage:**

- InfluxDB/Elasticsearch for metrics and ES (OpenSearch) for logs
 - Different ES clusters depending on the retention and privacy policies
- HDFS for long term archive

- **Visualization:**

- Grafana for dashboards and real-time monitoring
- Kibana for log analysis
- SWAN for long-term data analysis from HDFS

Servers
9.5 K

Cores
447.4 K


Disks
102.3 K

Tape Drives
187


Routers
296

Wifi Points
5.0 K


Batch Jobs Running




EOS Active Data Transfers




File Transfer Throughput




Cloud Virtual Machines Created



Databases Activity



LHCOPN and LHCONE Total traffic



Details

Compute

Metric	Current
Servers	10 K
Cores	447 K
Disks	102 K

Storage

Metric	Current
Tape Drives	187
Tape Cartridges	41539
Tape Used Space (TB)	460706
Disk Total Space (TB)	740148


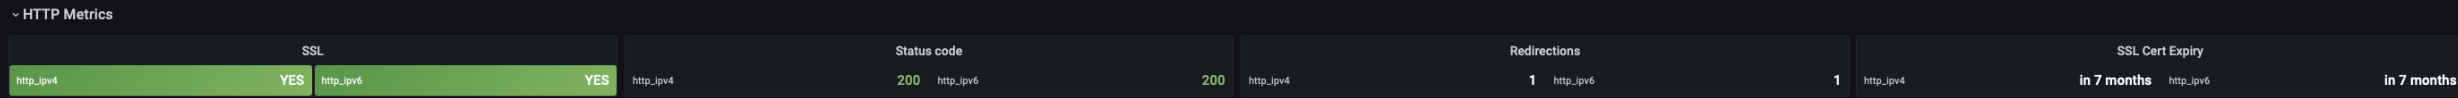
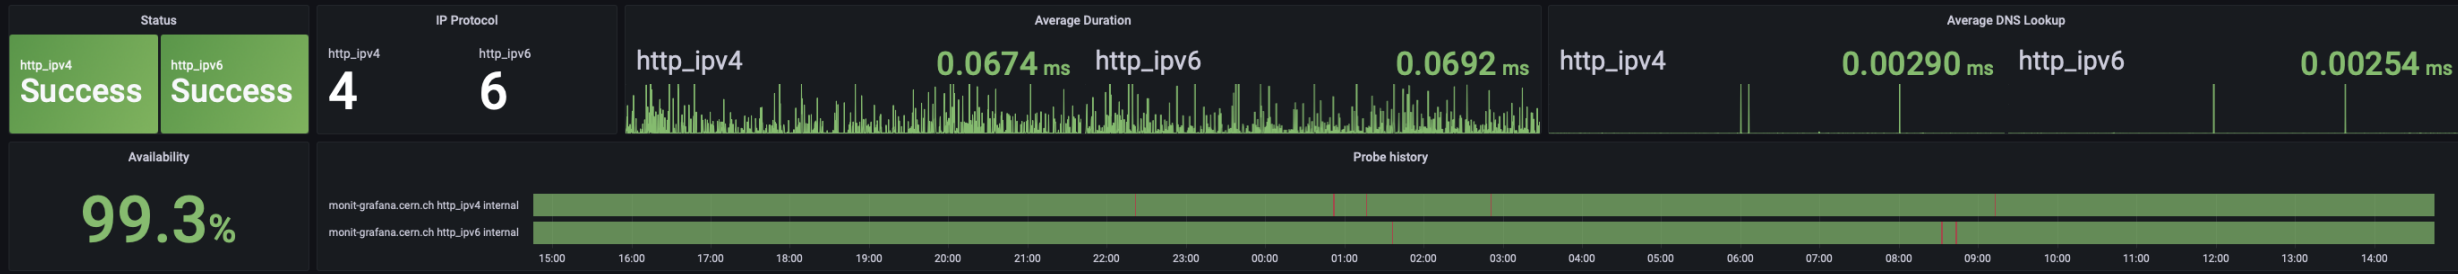
Network

Metric	Current
Routers	296
Start Points	736
Switches	4 K
Wifi Points	5 K
UTP Outlets	75 K
Devices	275 K

MONIT Remote Probes

- **MONIT service for probing endpoints**
 - Based on the Prometheus stack (Blackbox exporter)
 - Fully deployed in Kubernetes
 - CERN for internal network probes
 - OracleCloud for external probes
- **Remote Probe features**
 - Supports HTTP and ICMP over IPv4 and **IPv6**
 - Only IPv4 supported from the OracleCloud cluster
 - Trusted CERN/GRID certificates
 - Puppet probes (probing all hosts under same hostgroup)
- **Data stored in Prometheus and available for dashboards in Grafana**

Remote probe source: monit_prometheus_remote | Producer: monitoring | Instances: http://monit-grafana.cern.ch | Module: All | Job: All | Preferred IP Protocol: All | IP Protocol Fallback: All | Probe Type: All | Module Group: All | Mode: All



ICMP Metrics (1 panel)

Probes Logs

Timestamp	Raw Logs
2022-10-25 14:45:26.491	ts=2022-10-25T12:45:26.491Z caller=main.go:180 module=http_ipv6 target=http://monit-grafana.cern.ch level=debug msg="Probe succeeded" duration_seconds=0.119789039
2022-10-25 14:45:26.491	ts=2022-10-25T12:45:26.491Z caller=main.go:180 module=http_ipv6 target=http://monit-grafana.cern.ch level=debug msg="Response timings for roundtrip" roundtrip=1 start=2022-10-25T12:45:26.40613957Z dnsDone=2022-10-25T12:45:26.411030016Z connectDone=2022-10-25T12:45:26.411465594Z gotConn=2022-10-25T12:45:26.415185421Z responseStart=2022-10-25T12:45:26.49051534...
2022-10-25 14:45:26.491	ts=2022-10-25T12:45:26.491Z caller=main.go:180 module=http_ipv6 target=http://monit-grafana.cern.ch level=debug msg="Response timings for roundtrip" roundtrip=0 start=2022-10-25T12:45:26.403359013Z dnsDone=2022-10-25T12:45:26.403359013Z connectDone=2022-10-25T12:45:26.403840667Z gotConn=2022-10-25T12:45:26.403879765Z responseStart=2022-10-25T12:45:26.4046434...
2022-10-25 14:45:26.490	ts=2022-10-25T12:45:26.490Z caller=main.go:180 module=http_ipv6 target=http://monit-grafana.cern.ch level=debug msg="Received HTTP response" status_code=200
2022-10-25 14:45:26.405	ts=2022-10-25T12:45:26.405Z caller=main.go:180 module=http_ipv6 target=http://monit-grafana.cern.ch level=debug msg="Making HTTP request" url=https://monit-grafana.cern.ch/login host=
2022-10-25 14:45:26.405	ts=2022-10-25T12:45:26.405Z caller=main.go:180 module=http_ipv6 target=http://monit-grafana.cern.ch level=debug msg="Address does not match first address, not sending TLS ServerName" first=[2001:1458:d00:19::19f] address=monit-grafana.cern.ch
2022-10-25 14:45:26.404	ts=2022-10-25T12:45:26.404Z caller=main.go:180 module=http_ipv6 target=http://monit-grafana.cern.ch level=debug msg="Received redirect" location=https://monit-grafana.cern.ch/login
2022-10-25 14:45:26.402	ts=2022-10-25T12:45:26.402Z caller=main.go:180 module=http_ipv6 target=http://monit-grafana.cern.ch level=debug msg="Making HTTP request" url=http://[2001:1458:d00:19::19f] host=monit-grafana.cern.ch
2022-10-25 14:45:26.379	ts=2022-10-25T12:45:26.379Z caller=main.go:180 module=http_ipv6 target=http://monit-grafana.cern.ch level=debug msg="Resolved target address" ip=2001:1458:d00:19::19f
2022-10-25 14:45:26.372	ts=2022-10-25T12:45:26.372Z caller=main.go:180 module=http_ipv6 target=http://monit-grafana.cern.ch level=debug msg="Beginning probe" probe=http timeout_seconds=5
2022-10-25 14:45:26.372	ts=2022-10-25T12:45:26.372Z caller=main.go:180 module=http_ipv6 target=http://monit-grafana.cern.ch level=debug msg="Resolving target address" ip_protocol=ip6
2022-10-25 14:45:14.278	ts=2022-10-25T12:45:14.278Z caller=main.go:180 module=http_ipv4 target=http://monit-grafana.cern.ch level=debug msg="Response timings for roundtrip" roundtrip=0 start=2022-10-25T12:45:13.336662639Z dnsDone=2022-10-25T12:45:13.336662639Z connectDone=2022-10-25T12:45:13.337166734Z gotConn=2022-10-25T12:45:13.337201684Z responseStart=2022-10-25T12:45:13.3378599...
2022-10-25 14:45:14.278	ts=2022-10-25T12:45:14.278Z caller=main.go:180 module=http_ipv4 target=http://monit-grafana.cern.ch level=debug msg="Response timings for roundtrip" roundtrip=1 start=2022-10-25T12:45:13.33918198Z dnsDone=2022-10-25T12:45:13.348170336Z connectDone=2022-10-25T12:45:13.348710878Z gotConn=2022-10-25T12:45:13.352489943Z responseStart=2022-10-25T12:45:14.27782874...
2022-10-25 14:45:14.278	ts=2022-10-25T12:45:14.277Z caller=main.go:180 module=http_ipv4 target=http://monit-grafana.cern.ch level=debug msg="Received HTTP response" status_code=200
2022-10-25 14:45:14.278	ts=2022-10-25T12:45:14.278Z caller=main.go:180 module=http_ipv4 target=http://monit-grafana.cern.ch level=debug msg="Probe succeeded" duration_seconds=0.971100154



MONIT WLCG Site Monitoring

- **Data producers:**
 - ETF providing probe metrics (*IPv6* metrics already available)
- **Data processing:**
 - Sites Availability/Reliability status calculated using ETF metrics
 - Apache Spark job running regularly to compute the availability/reliability metrics
 - Different "profiles" supported per VO (ATLAS is testing *IPv6* profile)
- **Data storage:**
 - Elasticsearch
 - 1 hour resolution for Availability/Reliability statistics
 - 10 minutes resolution for the site/endpoint status with shorter retention (1 year)
 - HDFS for long term archive




VO cms Profile CMS_CRITICAL Tier 1 Country All Federation All Site All Flavour All Hostname All Recomputations status

More Dashboards Website

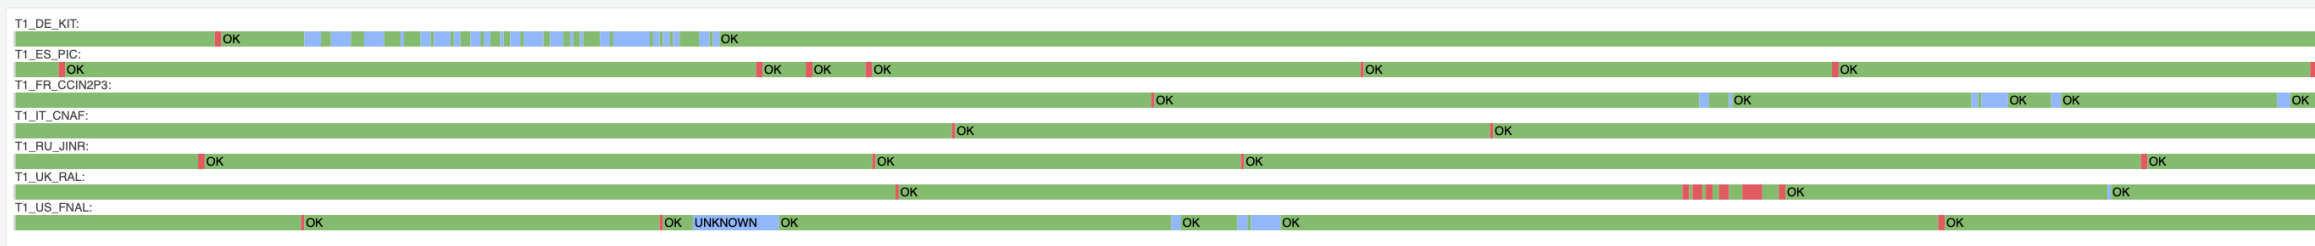
This dashboard shows the availability, reliability, and status computation from ETF tests. To use it:

1. Apply the desired selection using the filters on top (dashboard will also load faster).
2. Click on "Availability & Reliability", "Site Status", "Endpoint Status", or "Test Status" to see results. Scroll up/down for more results.


Availability & Reliability



Site Status



Endpoint Status



Test Status (1 panel)



VO All Tier All Country All Federation All Site T1_UK_RAL Flavour All Hostname arc-ce05.gridpp.rl.ac.uk Metric All Status All Filters +

More Dashboards Website

This dashboard shows all ETF test results. To use it:


1. Apply the desired selection using the filters on top (dashboard will also load faster).
2. Click on "Table" or "Timeline" to see results. Scroll up/down for more results.

Table

Timestamp	Status	Site	Flavour	Hostname	Metric	Summary	In Report
2022-10-25 12:21:53.000	OK	T1_UK_RAL	ARC-CE	arc-ce05.gridpp.rl.ac.uk	org.cms.WN-xrootd-access-/cms/Role=lcgadmin	cmssgm-473074.0-lcg2240.gridpp.rl.ac.uk: OK	true
2022-10-25 12:21:53.000	OK	T1_UK_RAL	ARC-CE	arc-ce05.gridpp.rl.ac.uk	org.cms.WN-analysis-/cms/Role=lcgadmin	cmssgm-473074.0-lcg2240.gridpp.rl.ac.uk: OK	true
2022-10-25 12:21:53.000	OK	T1_UK_RAL	ARC-CE	arc-ce05.gridpp.rl.ac.uk	org.cms.WN-cvmfs-/cms/Role=lcgadmin	cmssgm-473074.0-lcg2240.gridpp.rl.ac.uk: OK. cvmfs vers 2.9.4 (probe 1.3-pr...	true
2022-10-25 12:21:53.000	OK	T1_UK_RAL	ARC-CE	arc-ce05.gridpp.rl.ac.uk	org.cms.WN-xrootd-fallback-/cms/Role=lcgadmin	cmssgm-473074.0-lcg2240.gridpp.rl.ac.uk: OK	true
2022-10-25 12:21:53.000	OK	T1_UK_RAL	ARC-CE	arc-ce05.gridpp.rl.ac.uk	org.cms.WN-frontier-/cms/Role=lcgadmin	cmssgm-473074.0-lcg2240.gridpp.rl.ac.uk: OK	true
2022-10-25 12:21:53.000	OK	T1_UK_RAL	ARC-CE	arc-ce05.gridpp.rl.ac.uk	org.cms.WN-mc-/cms/Role=lcgadmin	cmssgm-473074.0-lcg2240.gridpp.rl.ac.uk: OK	true
2022-10-25 12:21:53.000	OK	T1_UK_RAL	ARC-CE	arc-ce05.gridpp.rl.ac.uk	org.cms.WN-squid-/cms/Role=lcgadmin	cmssgm-473074.0-lcg2240.gridpp.rl.ac.uk: OK	true
2022-10-25 12:21:53.000	OK	T1_UK_RAL	ARC-CE	arc-ce05.gridpp.rl.ac.uk	org.cms.WN-basic-/cms/Role=lcgadmin	cmssgm-473074.0-lcg2240.gridpp.rl.ac.uk: OK	true
2022-10-25 12:21:53.000	OK	T1_UK_RAL	ARC-CE	arc-ce05.gridpp.rl.ac.uk	org.sam.CONDOR-JobSubmit-/cms/Role=lcgadmin	OK - Job successfully completed	true
2022-10-25 12:21:53.000	OK	T1_UK_RAL	ARC-CE	arc-ce05.gridpp.rl.ac.uk	org.cms.WN-isolation-/cms/Role=lcgadmin	cmssgm-473074.0-lcg2240.gridpp.rl.ac.uk: OK	true

Downtimes (1 panel)

Timeline

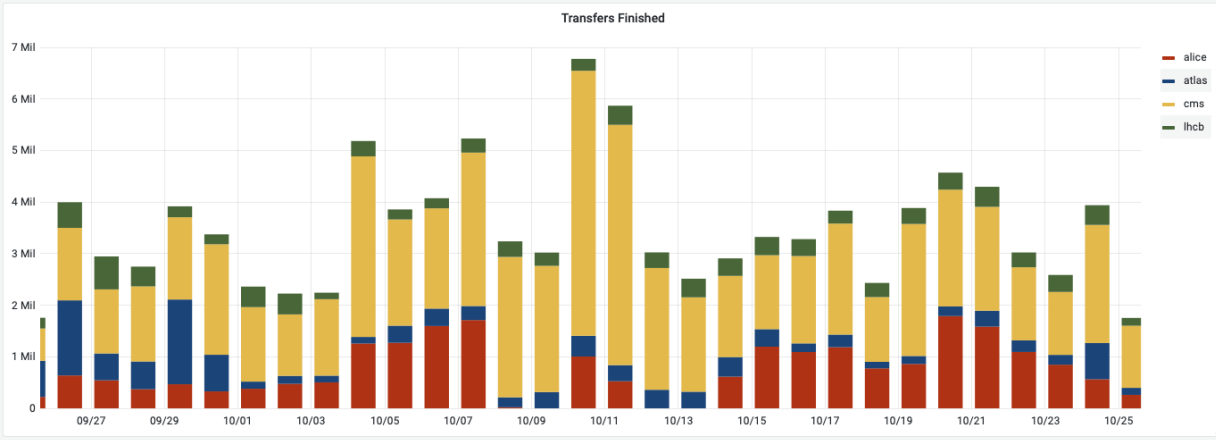
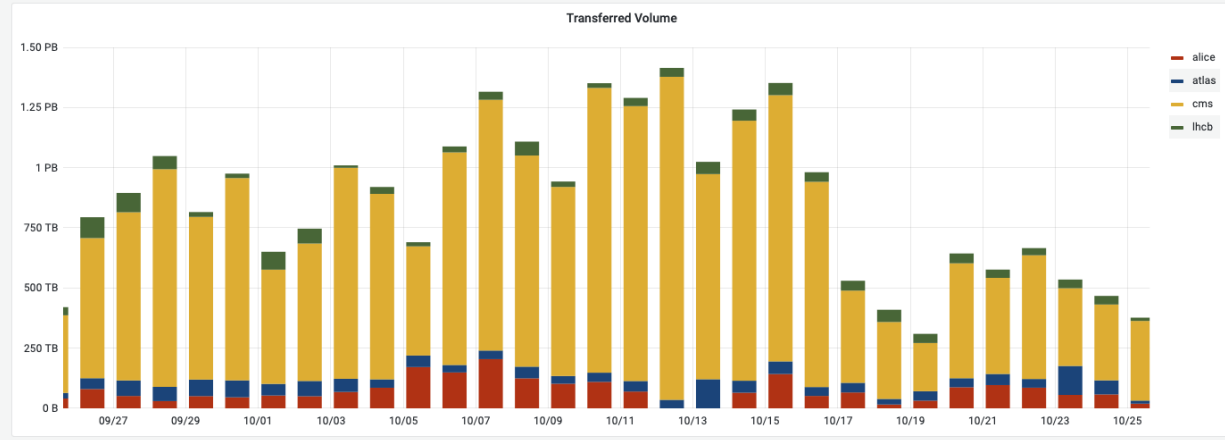









MONIT WLCG Transfers Monitoring

- **Data producers:**
 - FTS and XRootD transfers data
 - Both providing *IPv6* flags information
 - For XRootD the *IPv6* flag is acquired based on the server/client domain name
- **Data processing:**
 - Data enrichment with topology information
 - Data aggregation for long term storage
 - Apache Spark streaming and batch jobs creating the new data documents
- **Data storage:**
 - Elasticsearch with different retention for enriched/aggregated data
 - HDFS for long term archive



Group By vo Bin 1d VO All Source Country All Dest Country All Source Site All Dest Site All Source Exp. Site All Dest Exp. Site All Source Tier All Dest Tier All Remote Access true IPv6 All Filters +





Summary

- **MONIT provides infrastructure for processing and storing metrics and logs**
 - Base CERN DC Monitoring
 - Service metrics and logs
 - Remote probes
 - WLCG monitoring
- **MONIT source agents are running in dual stack**
 - Receiving data from IPv6 hosts
- **IPv6 related monitoring data could be adopted if provided by the producer**

Thank you!

IT Monitoring Service

- [Main site](#)
- [Users documentation](#)
- [WLCG Sitemon](#)
- [MONIT Grafana](#)

Questions?



home.cern