
CERN

IT Monitoring Service

Status and Progress

Alberto AIMAR CERN-IT
for the MONIT Team

Outline

- Scope and Mandate
- Architecture and Data Flow
- Current Status and Progress
- Strategy and Plans

Scope and Mandate

Monitoring - Scope

Data Centres Monitoring

- Monitoring of DCs at CERN and Wigner
- Hardware, operating system, and services
- Data Centres equipment, PDUs, temperature sensors, etc.
- Metrics and logs

Experiment Dashboards

- WLCG Monitoring
- Sites availability, data transfers, job information, reports
- Used by WLCG, experiments, sites and users

Data Centres Monitoring (meter)

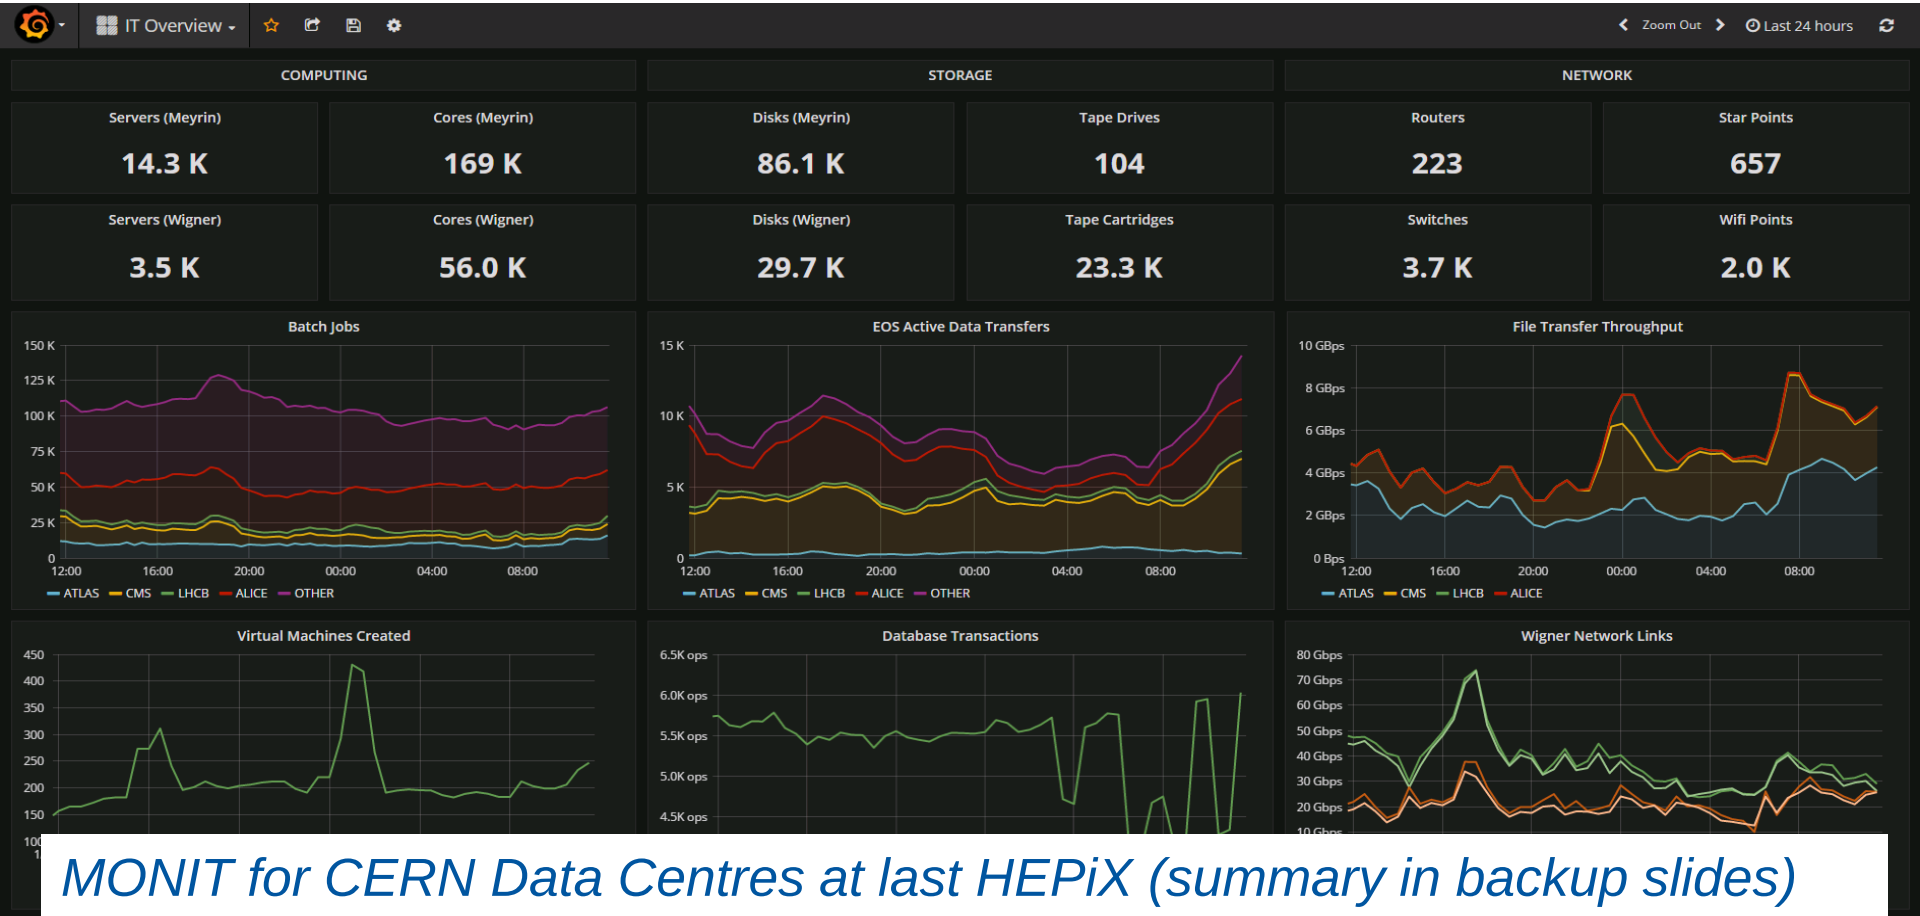

MONIT for CERN Data Centres at last HEPiX (summary in backup slides)

Experiment Dashboards

- Job monitoring, sites availability, data management and transfers
- Used by experiments operation teams, sites, users, WLCG

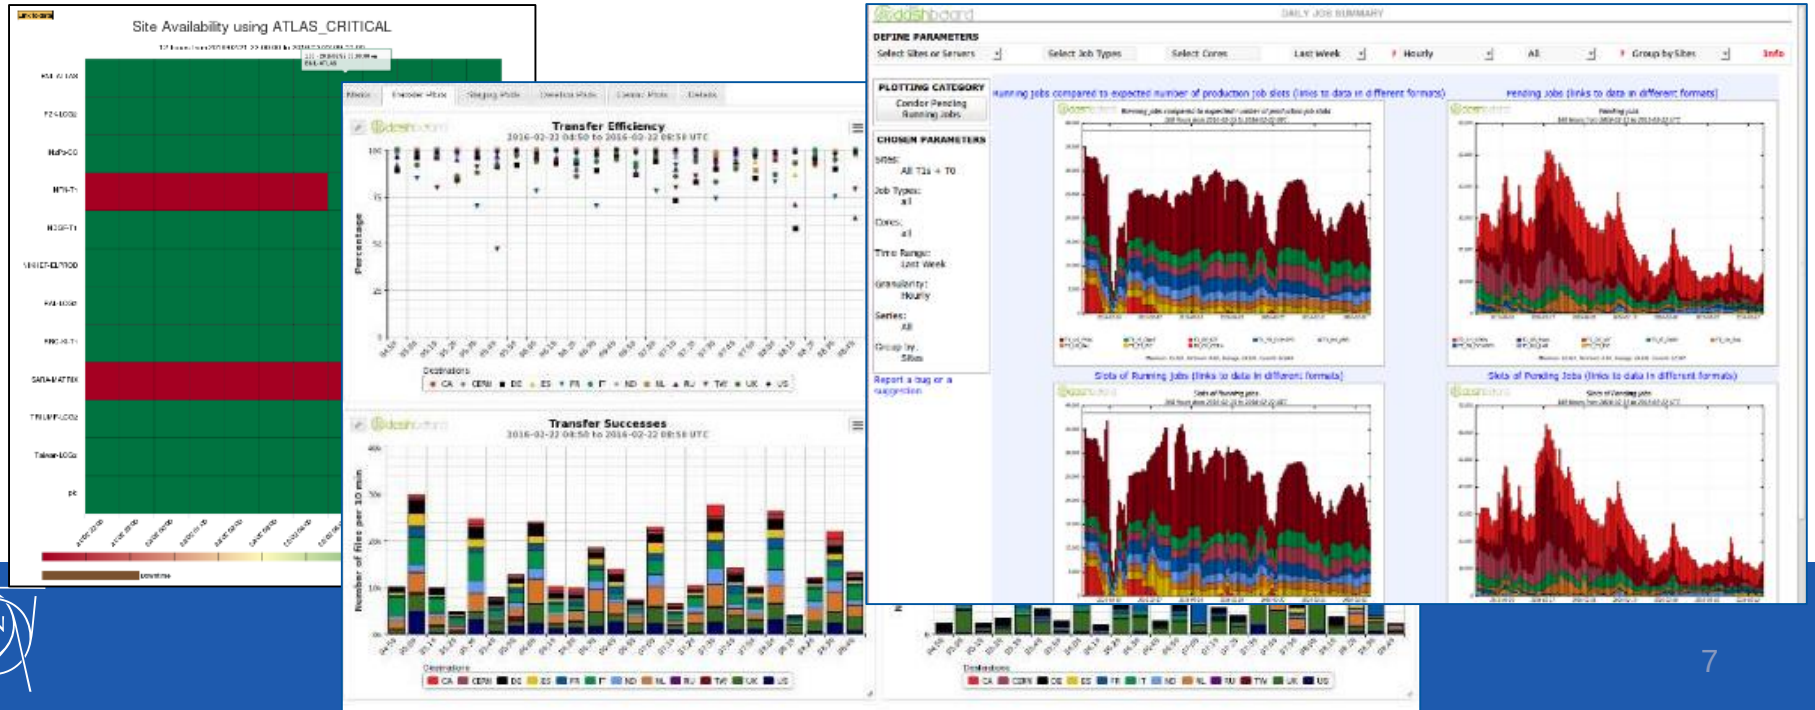

WLCG Monitoring - Mandate

- Regroup monitoring activities hosted by CERN IT
  - Monitoring of Data Centres, WLCG and Experiment Dashboards
  - ETF, HammerCloud testing frameworks
- Uniform with standard CERN IT practices
- Management of services, communication, tools
- Review existing monitoring usage and needs (IT, WLCG, etc.)
- Investigate, implement established open source technologies
- Reduce dependencies on in-house software and on few experts
- Continue support, while preparing the new services

Architecture and Data Flow

Previous Monitoring

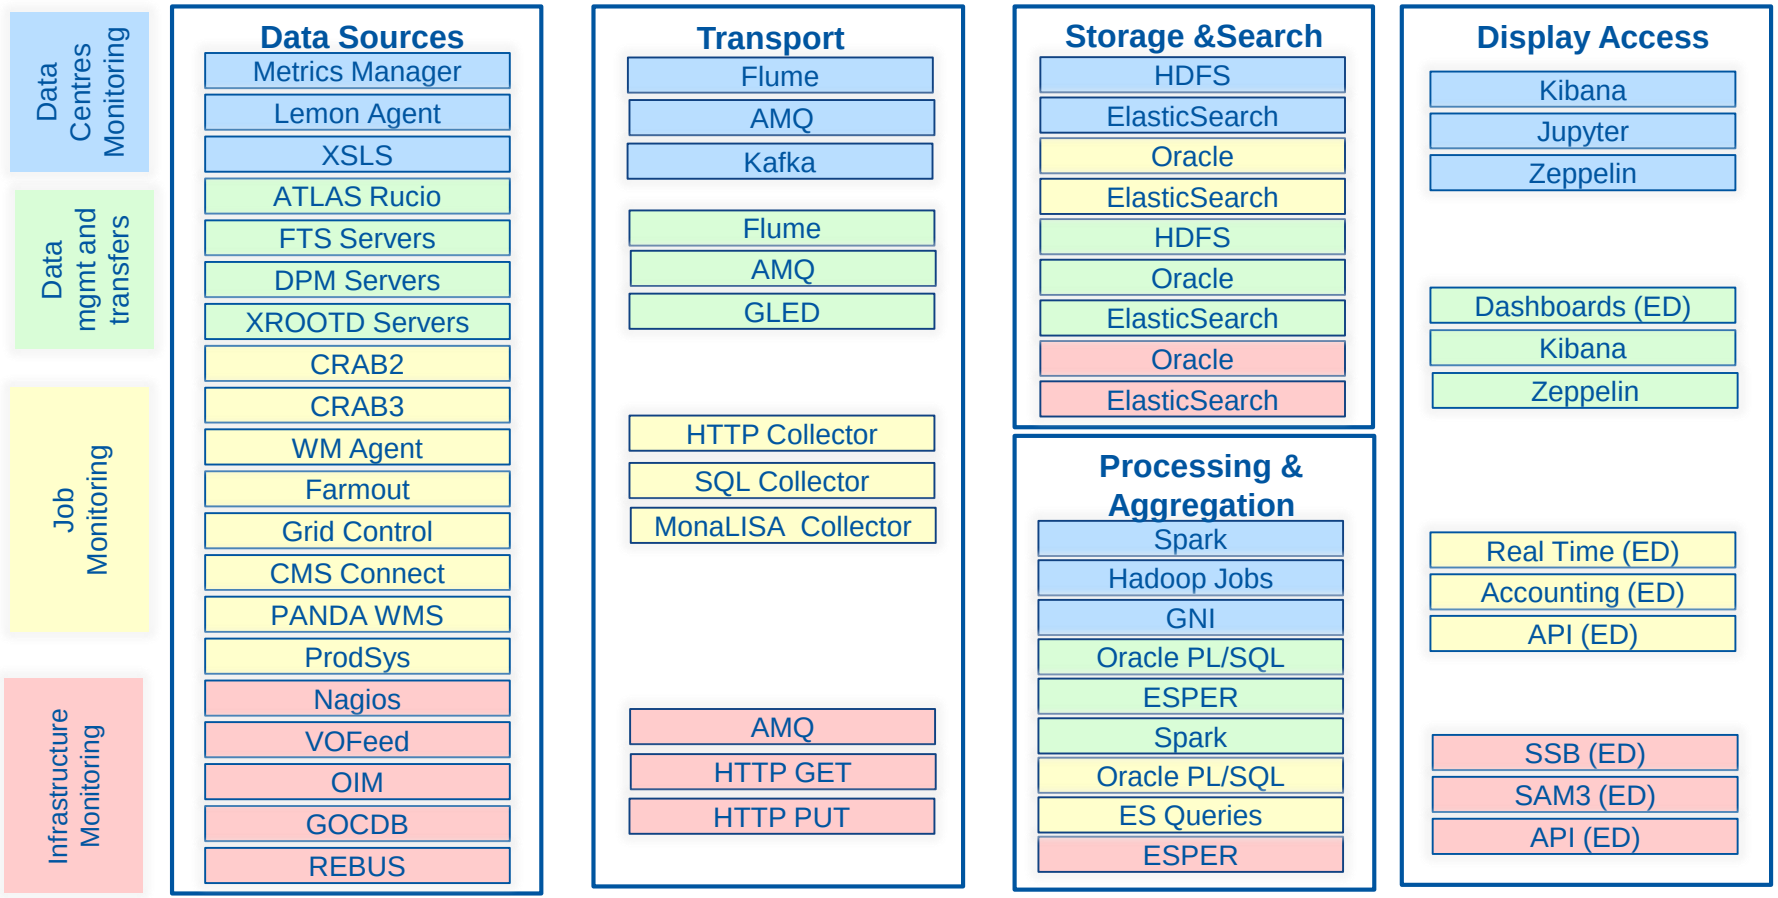

Unified Monitoring

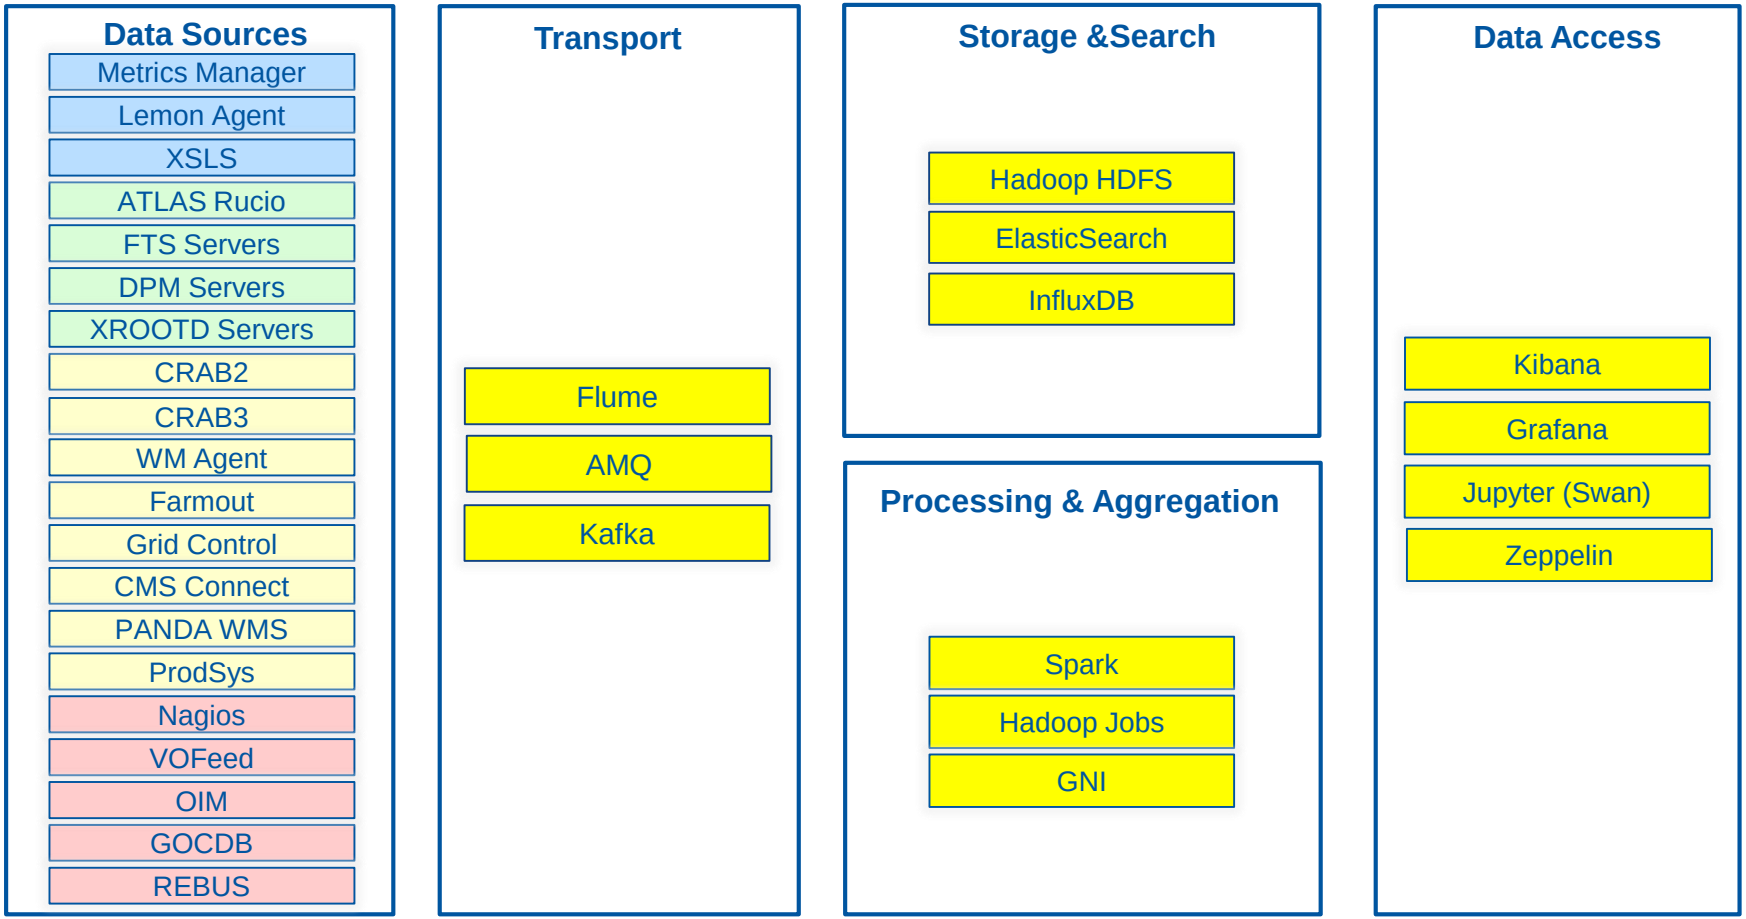

Unified Monitoring Architecture

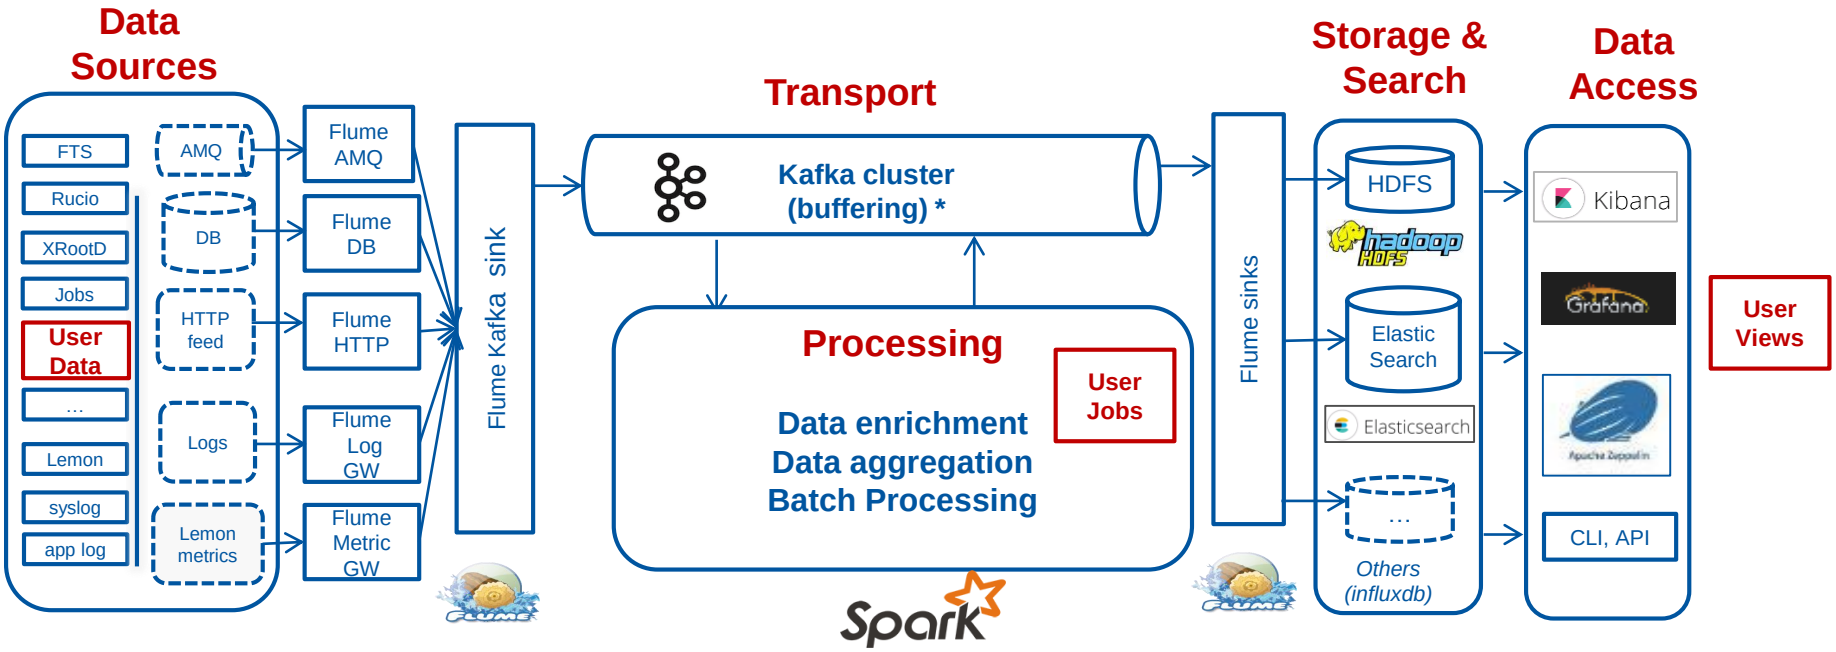

Today: > 500 GB/day, 72h buffering

Unified Data Sources

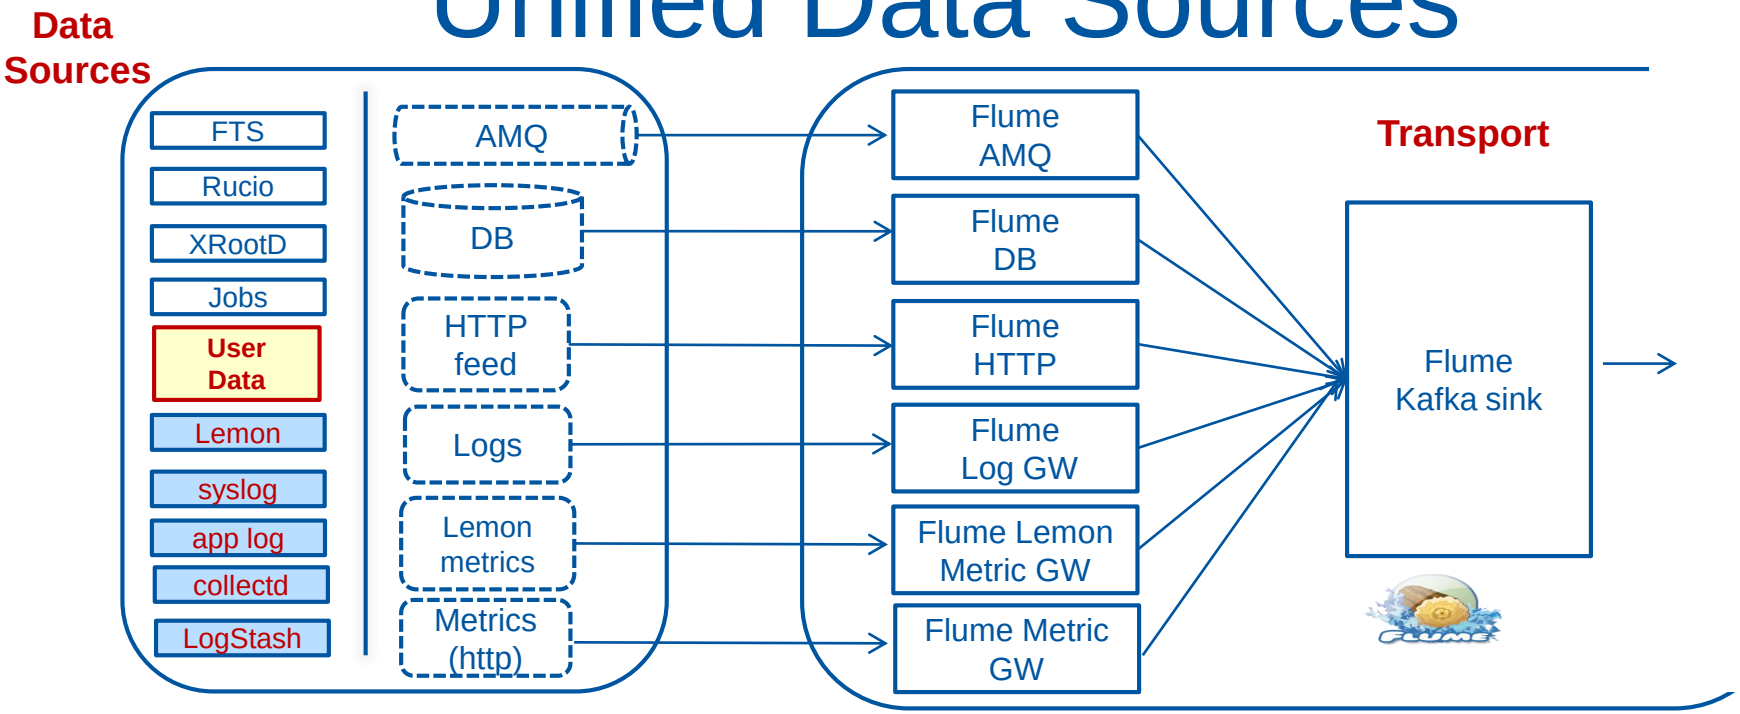

- Data is all channeled via Flume via gateways
  - Validated and normalized if necessary (e.g. standard names, date formats)
- Adding new Data Sources is documented and fairly simple (User Data)
  - Available both for Metrics (IT, WLCG, etc.) and Logs (hw logs, OS logs, syslogs, app logs)

Unified Processing

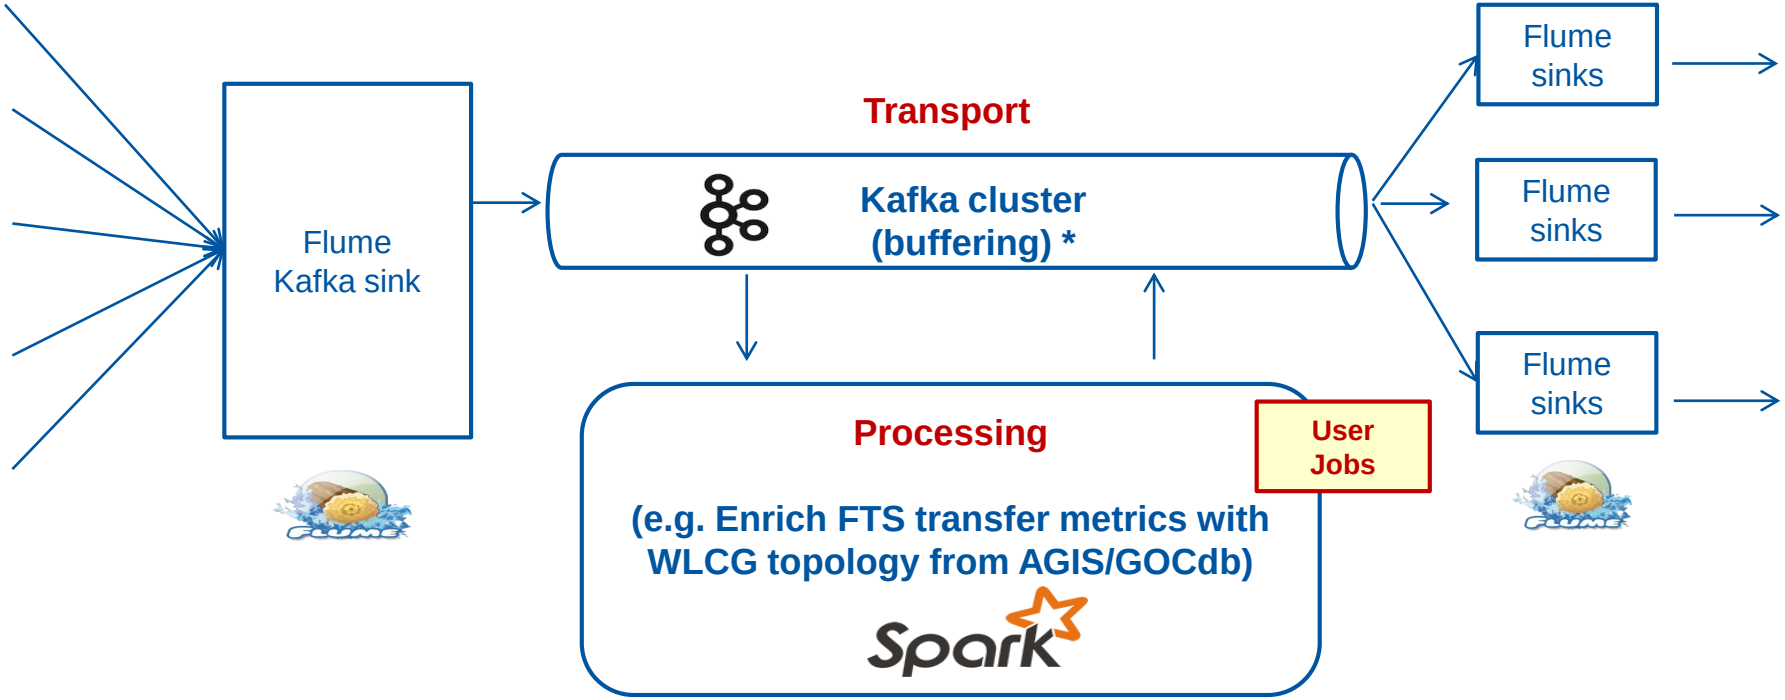

Proven useful many times

Data Processing

Stream processing

Data enrichment

- Join information from several sources (e.g. WLCG topology)

Data aggregation

- Over time (e.g. summary statistics for a time bin)
- Over other dimensions (e.g. compute a cumulative metric for a set of machines hosting the same service)

Data correlation

- Advanced Alarming: detect anomalies and failures correlating data from multiple sources (e.g. data centre topology-aware alarms)

Batch processing

- Reprocessing, data compression, historical data, periodic reports

Unified Access

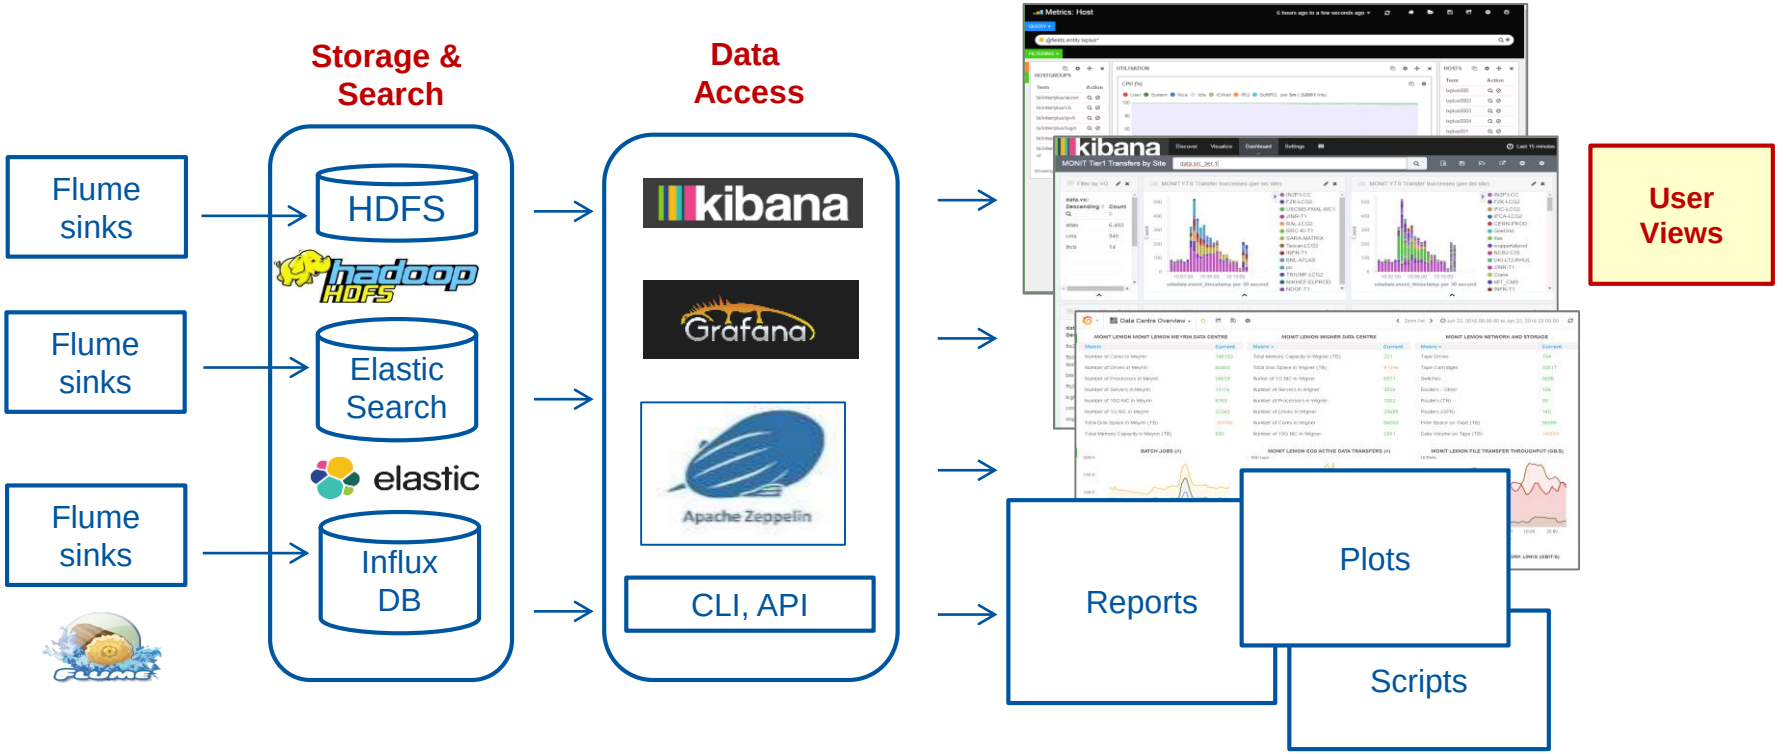

- Default dashboards, and can be customized and extended fairly easily
- Multiple data access methods (dashboards, notebooks, CLI)
- Dashboards, reports and access via scripts to create new User Views or reports

Current Status and Progress

Services Proposed

Monitor, collect, visualize, process, aggregate, alarm

- Metrics and Logs

Infrastructure operations and scale

Helping and supporting

- Interfacing new data sources
- Developing custom processing, aggregations, alarms
- Building dashboards and reports

Infrastructure(s) - Current Numbers

- Designed, developed and deployed a new monitoring infrastructure capable of handling CERN IT and WLCG data
  - ~ 150 VMs, ~ 500 GB/day, ~ 1 billion docs /day

- Maintenance legacy WLCG infrastructure and tools
  - ~ 90 VMs
- Maintenance legacy ITMON infrastructure and tools (i.e. meter, timber)
  - ~ 150 VMs

Infrastructure(s) - Operations

- Building and tuning the complete infrastructure
- Supporting existing services
- Depending on many external services
  - ES, InfluxDB, HDFS
  - Some also new and being set up
- Securing infrastructure
  - Flume/Kafka/Spark/ES/HDFS
- Configuring infrastructure (Puppet 4)
- User Documentation & Training

Current Data / WLCG

WLCG Data Sources	Additional Data Sources
FTS	ASAP ATLAS
XROOTD	
XROOTD ALICE	CRAB OPS
DDM RUCIO	CMS SPACEMON
DDM TRACES	GLIDEINWMS
DDM ACCOUNTING	
PHEDEX	
ATLAS JM – PANDA/PRODSYS	LHCOPN
CMS JM - HT CONDOR	BOINC - LHCATHOME
SAM3 ETF	PROTODUNE DAQ
AGIS	WMAGENTS
VOFEED	WM ARCHIVE
REBUS	
GOCDB	WLCG SPACE ACCOUNTING
OIM	

Current Processing / WLCG

Validation and Transformation	
	Fields Verifications (e.g. check timestamp in milliseconds in all doc)
	Fields Extractions (e.g. extract FTS log link, transfer ID)
	Fields Computations (e.g. create unique document ID based on other fields Field Normalization: apply common names (e.g. dst_site, dst_country, lowercase, etc.)
Enrichment	
	Topology Resolution
	Join raw data with AGIS, VOFeed, REBUS, etc.
Aggregation	
	Binning over time
	Summary data (e.g. for a given interval)
Specific Processing	Specific Spark Jobs (e.g. efficiency = success vs. failures)
	DDM Site avail
	Job monitoring and accounting
	FTS, XRootD transfers, rates
	Sites Availability, profiles (prototype)

MONIT Portal

- Users are lost among different technologies, dashboards, web pages, notebooks

- Single entry point for the MONIT dashboards and reports
- Direct links to dashboards and reports

Current Views: MONIT Portal

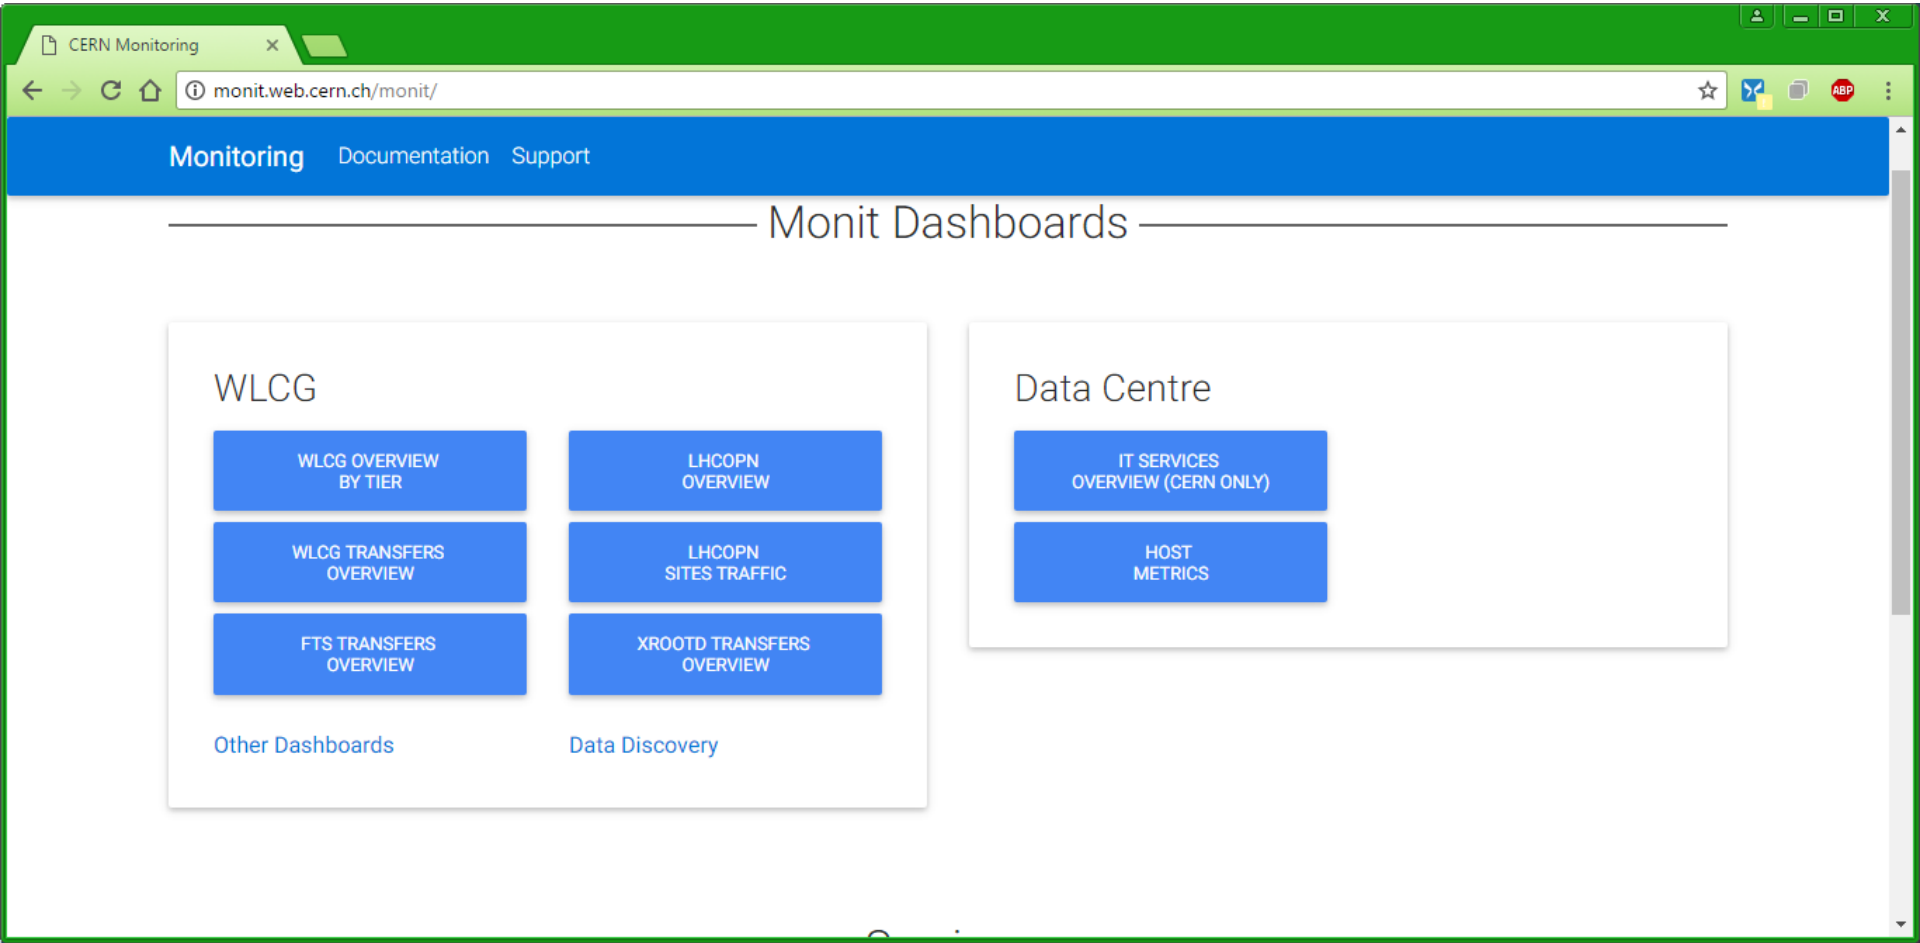

WLCG Transfers

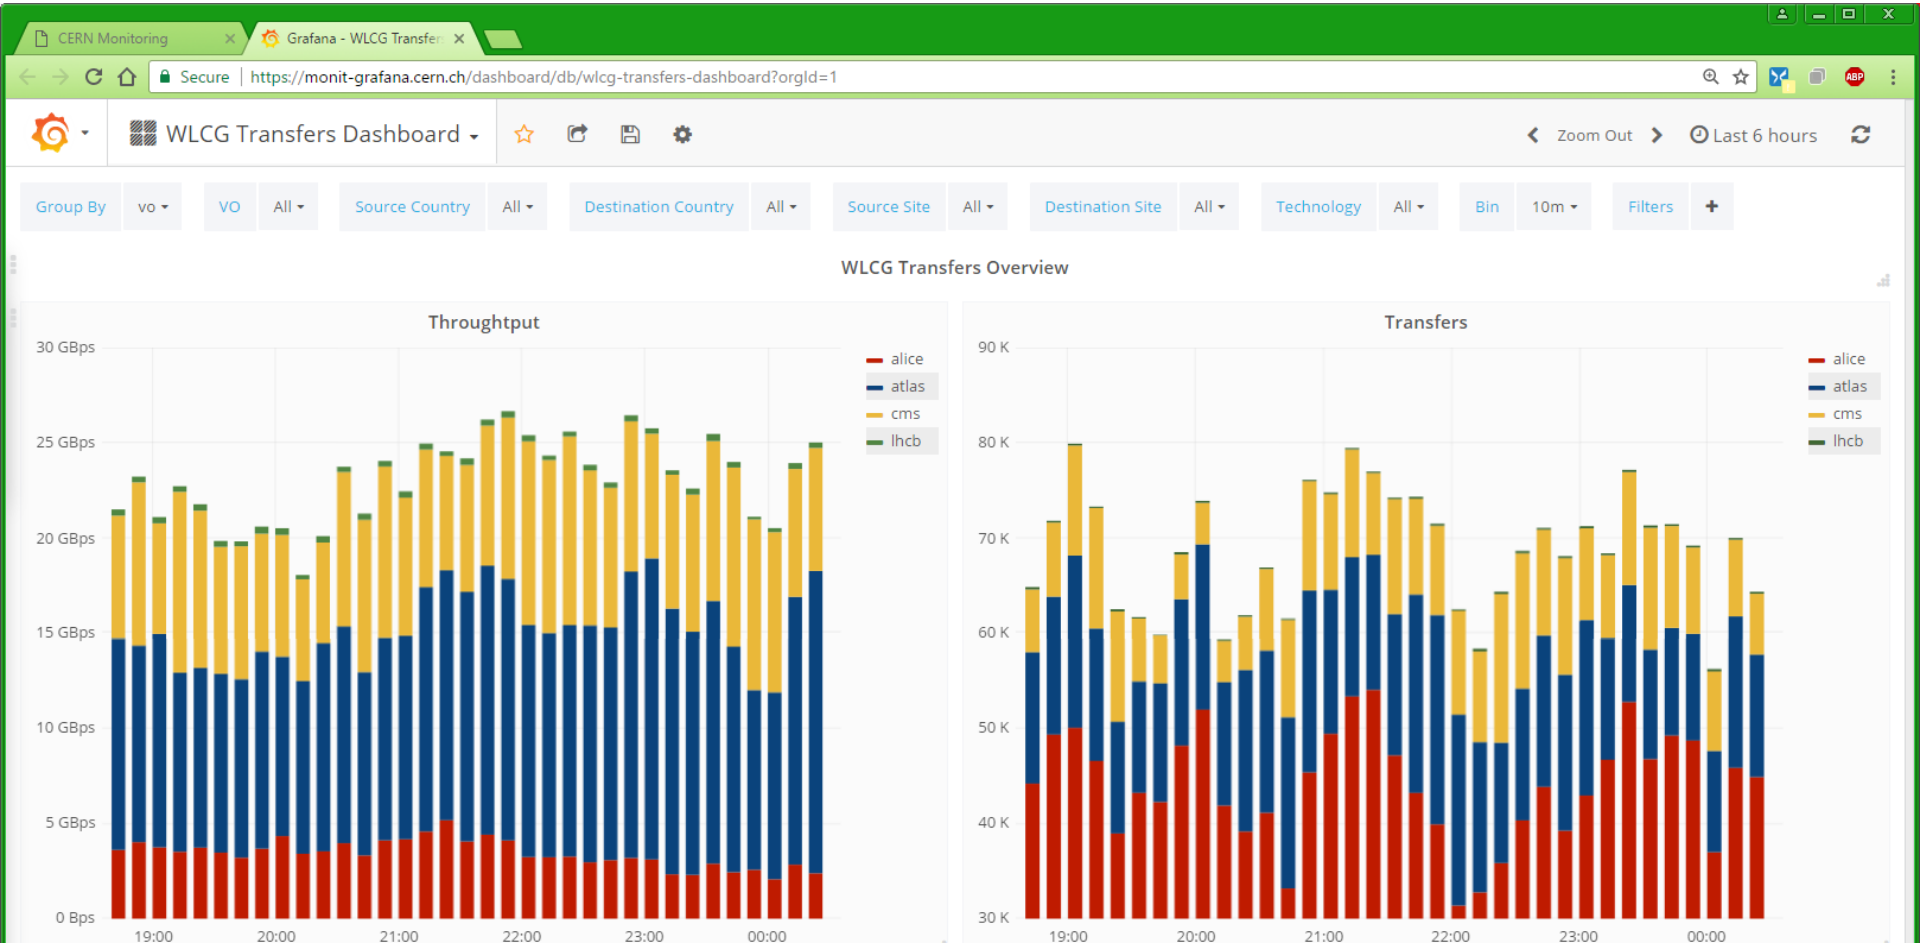

FTS Transfers

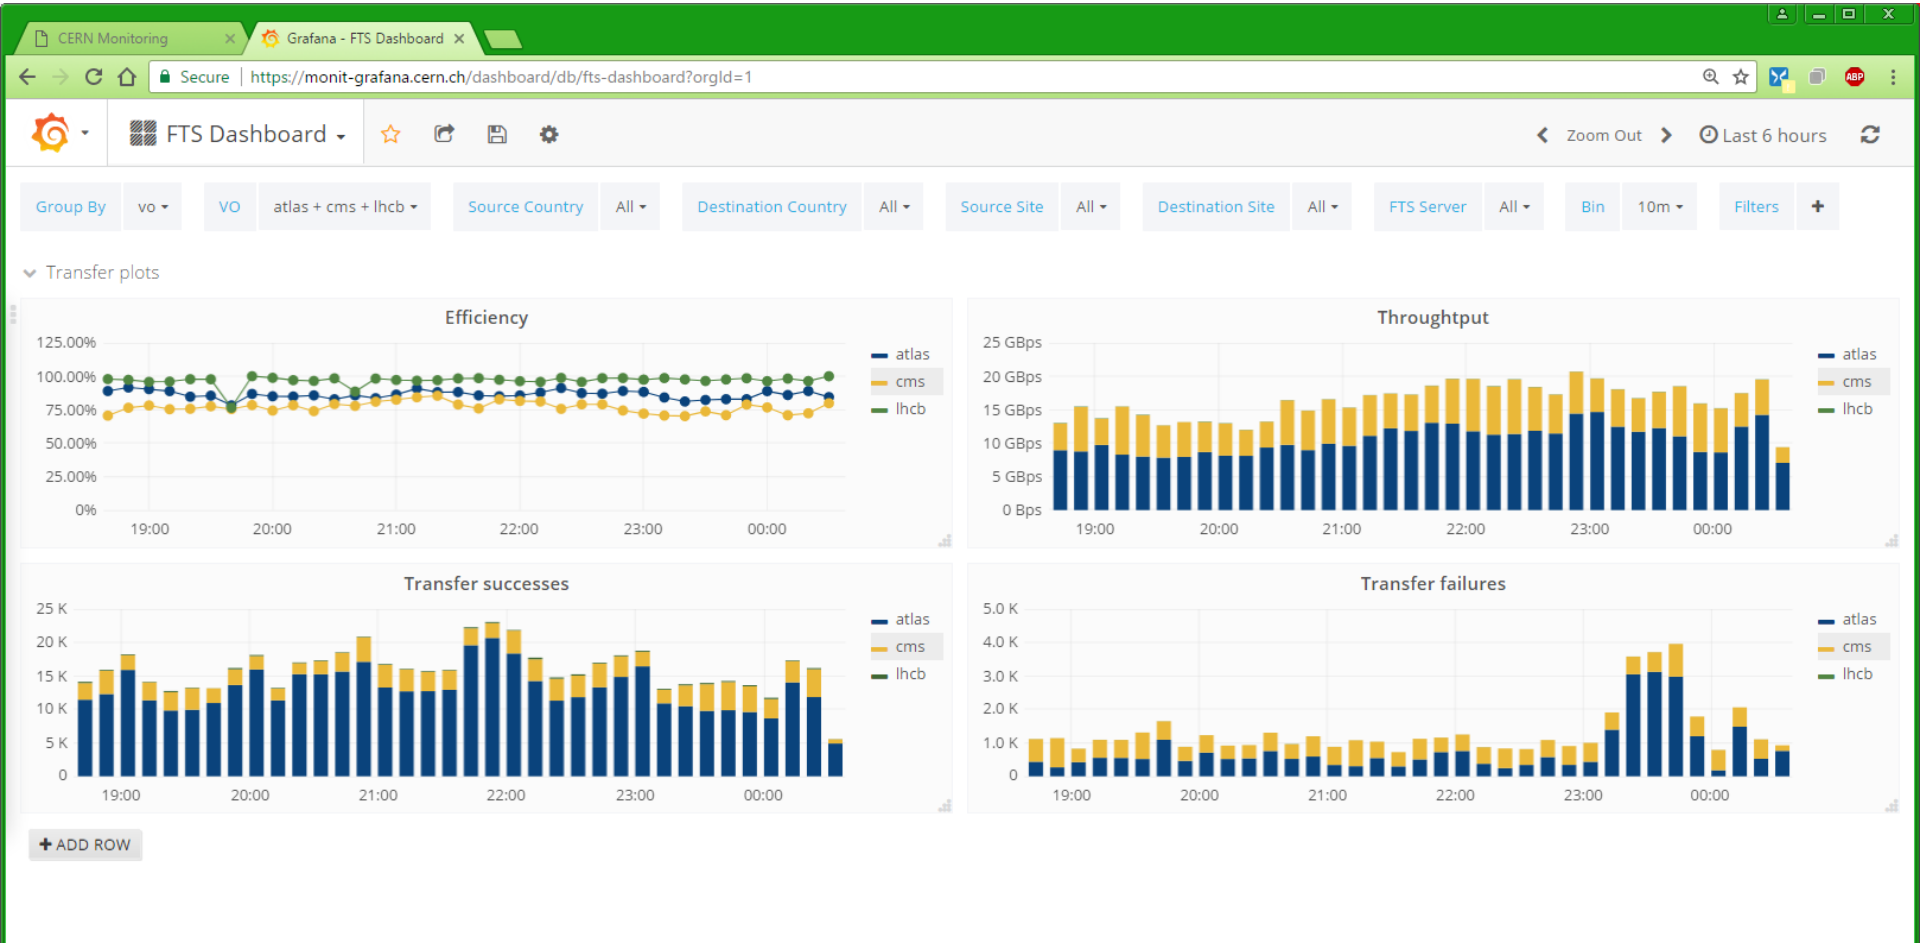

Transfers Overviews (T0, T1, T2)

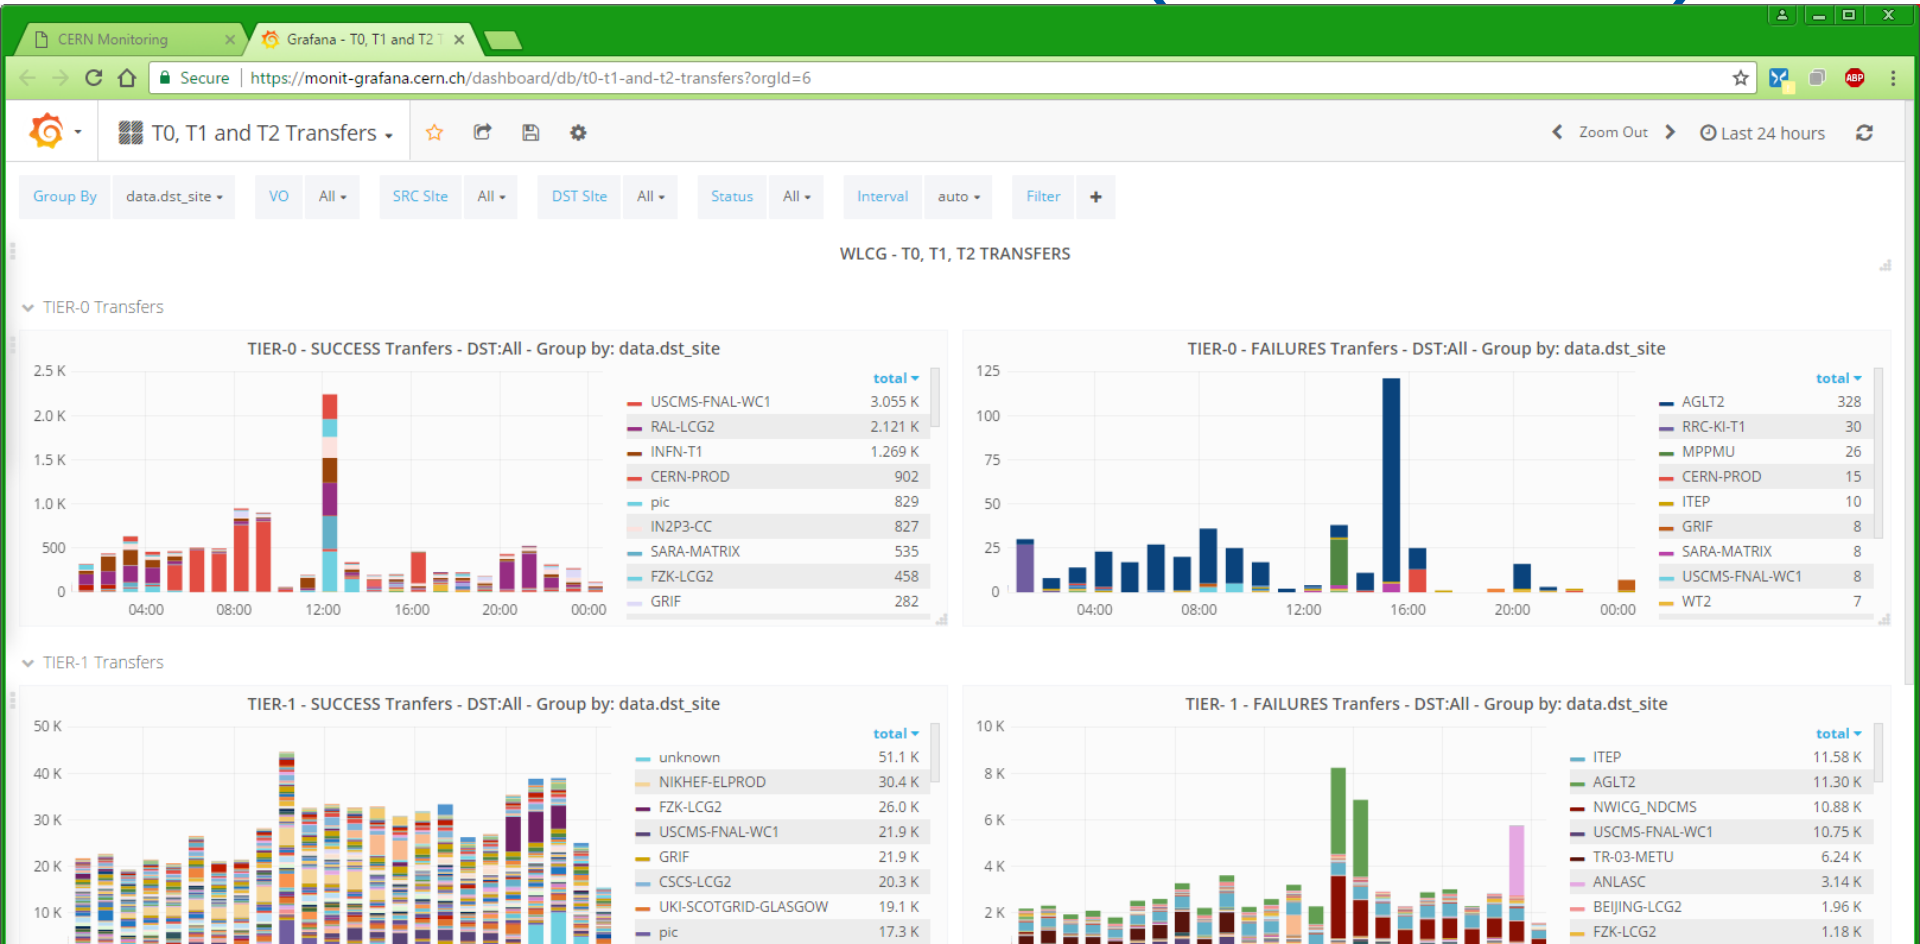

Transfers Efficiency (matrix, prototype)

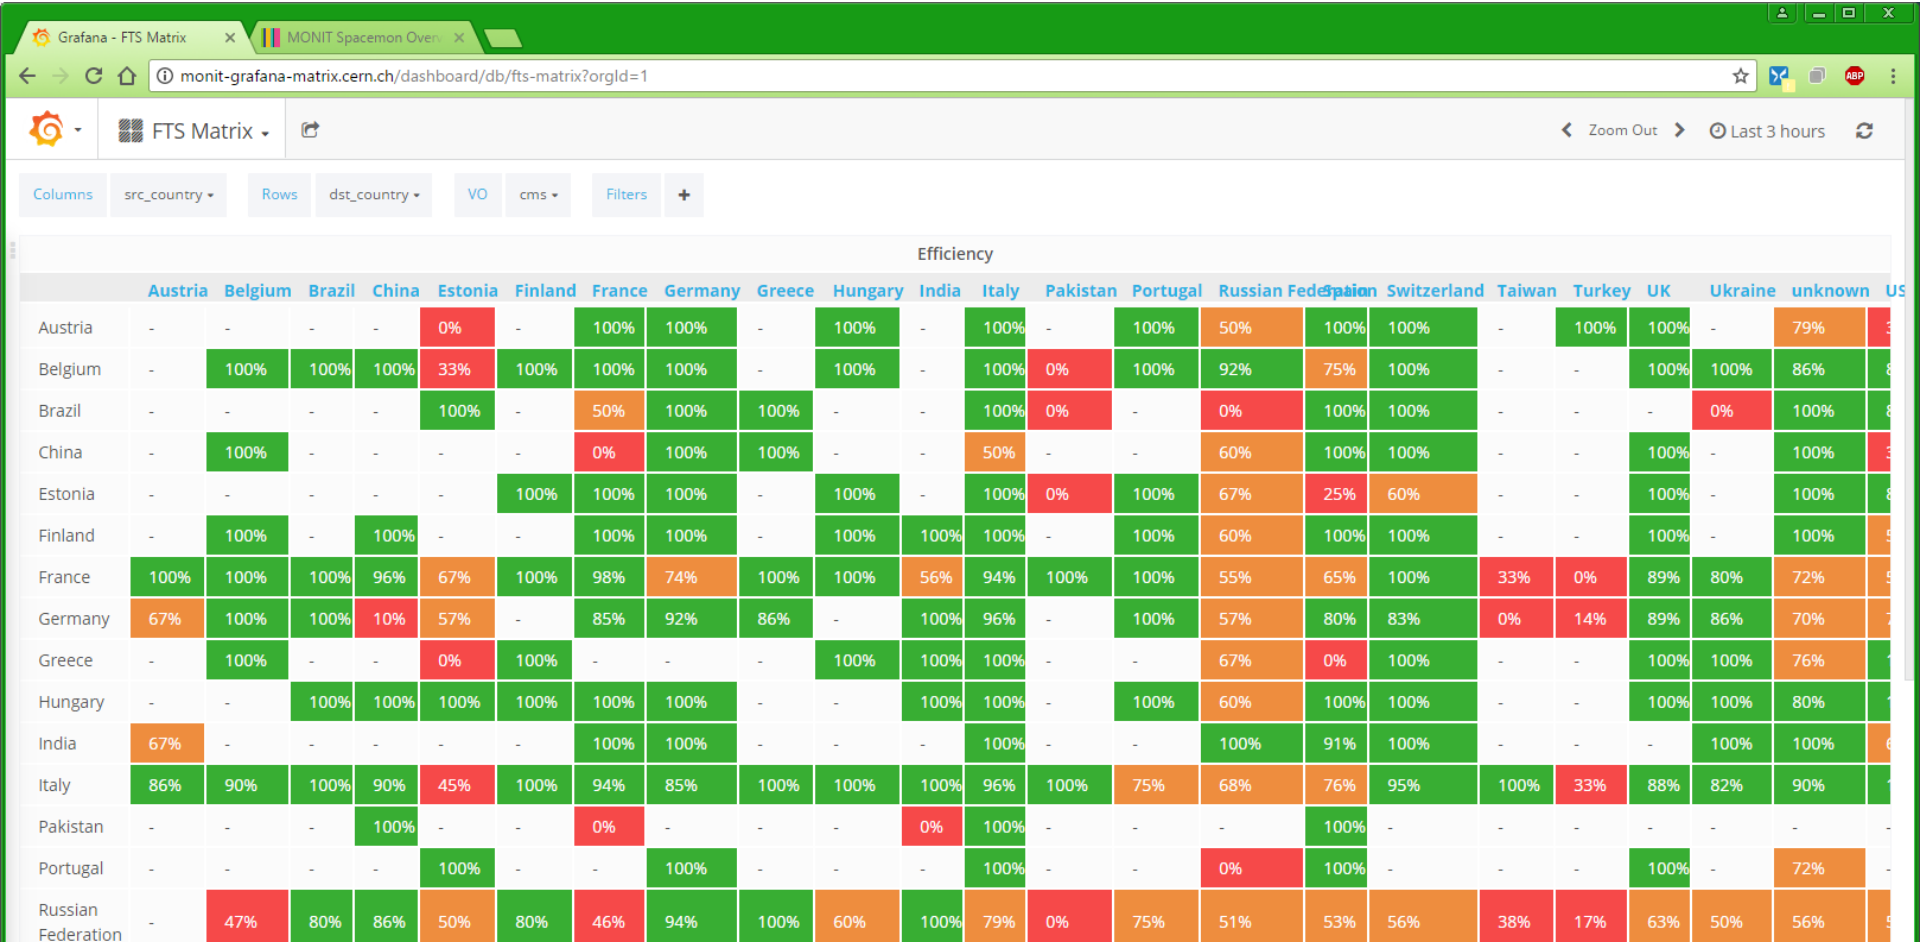

	Austria	Belgium	Brazil	China	Estonia	Finland	France	Germany	Greece	Hungary	India	Italy	Pakistan	Portugal	Russian Federation	Spain	Switzerland	Taiwan	Turkey	UK	Ukraine	unknown	US
Austria	-	-	-	-	0%	-	100%	100%	-	100%	-	100%	-	100%	50%	100%	100%	-	100%	100%	-	79%	
Belgium	-	100%	100%	100%	33%	100%	100%	100%	-	100%	-	100%	0%	100%	92%	75%	100%	-	-	100%	100%	86%	
Brazil	-	-	-	-	100%	-	50%	100%	100%	-	-	100%	0%	-	0%	100%	100%	-	-	-	0%	100%	
China	-	100%	-	-	-	-	0%	100%	100%	-	-	50%	-	-	60%	100%	100%	-	-	100%	-	100%	
Estonia	-	-	-	-	-	100%	100%	100%	-	100%	-	100%	0%	100%	67%	25%	60%	-	-	100%	-	100%	
Finland	-	100%	-	100%	-	-	100%	100%	-	100%	100%	100%	-	100%	60%	100%	100%	-	-	100%	-	100%	
France	100%	100%	100%	96%	67%	100%	98%	74%	100%	100%	56%	94%	100%	100%	55%	65%	100%	33%	0%	89%	80%	72%	
Germany	67%	100%	100%	10%	57%	-	85%	92%	86%	-	100%	96%	-	100%	57%	80%	83%	0%	14%	89%	86%	70%	
Greece	-	100%	-	-	0%	100%	-	-	-	100%	100%	100%	-	-	67%	0%	100%	-	-	100%	100%	76%	
Hungary	-	-	100%	100%	100%	100%	100%	100%	-	-	100%	100%	-	100%	60%	100%	100%	-	-	100%	100%	80%	
India	67%	-	-	-	-	-	100%	100%	-	-	-	100%	-	-	100%	91%	100%	-	-	-	100%	100%	
Italy	86%	90%	100%	90%	45%	100%	94%	85%	100%	100%	100%	96%	100%	75%	68%	76%	95%	100%	33%	88%	82%	90%	
Pakistan	-	-	-	100%	-	-	0%	-	-	-	0%	100%	-	-	-	100%	-	-	-	-	-	-	-
Portugal	-	-	-	-	100%	-	-	100%	-	-	-	100%	-	-	0%	100%	-	-	-	100%	-	72%	-
Russian Federation	-	47%	80%	86%	50%	80%	46%	94%	100%	60%	100%	79%	0%	75%	51%	53%	56%	38%	17%	63%	50%	56%	

Portal: Experiment Dashboards

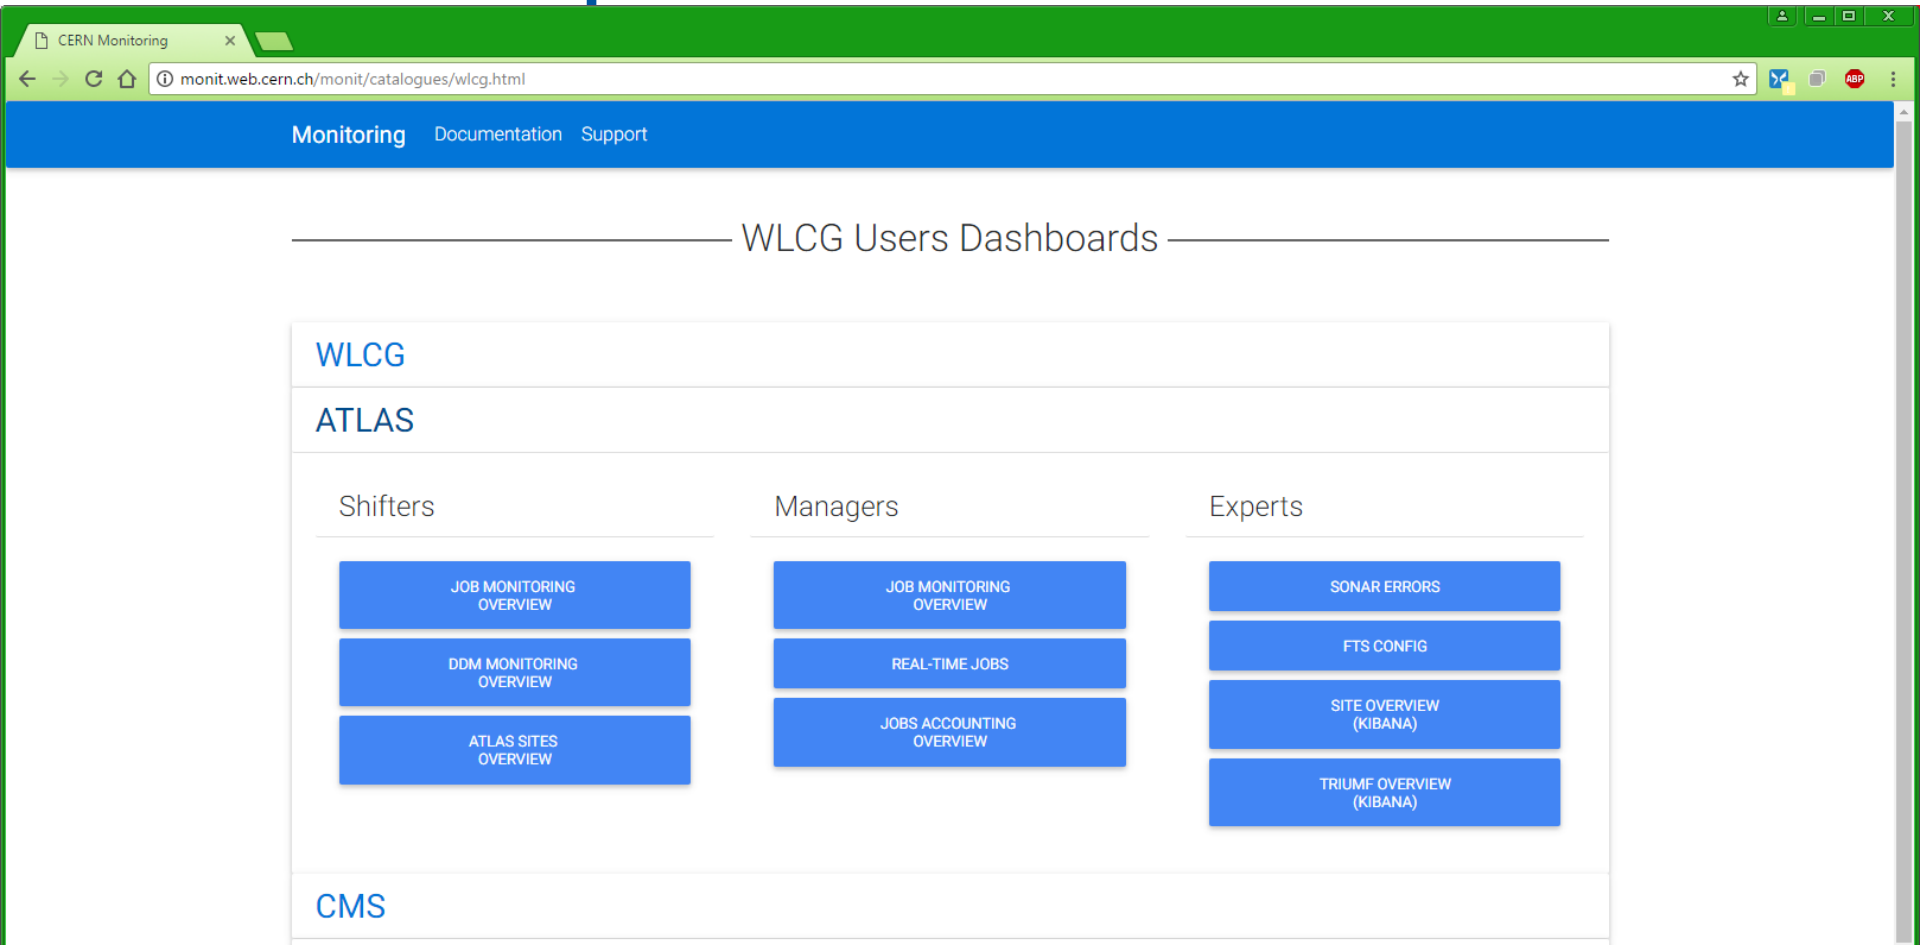

ATLAS Running Jobs

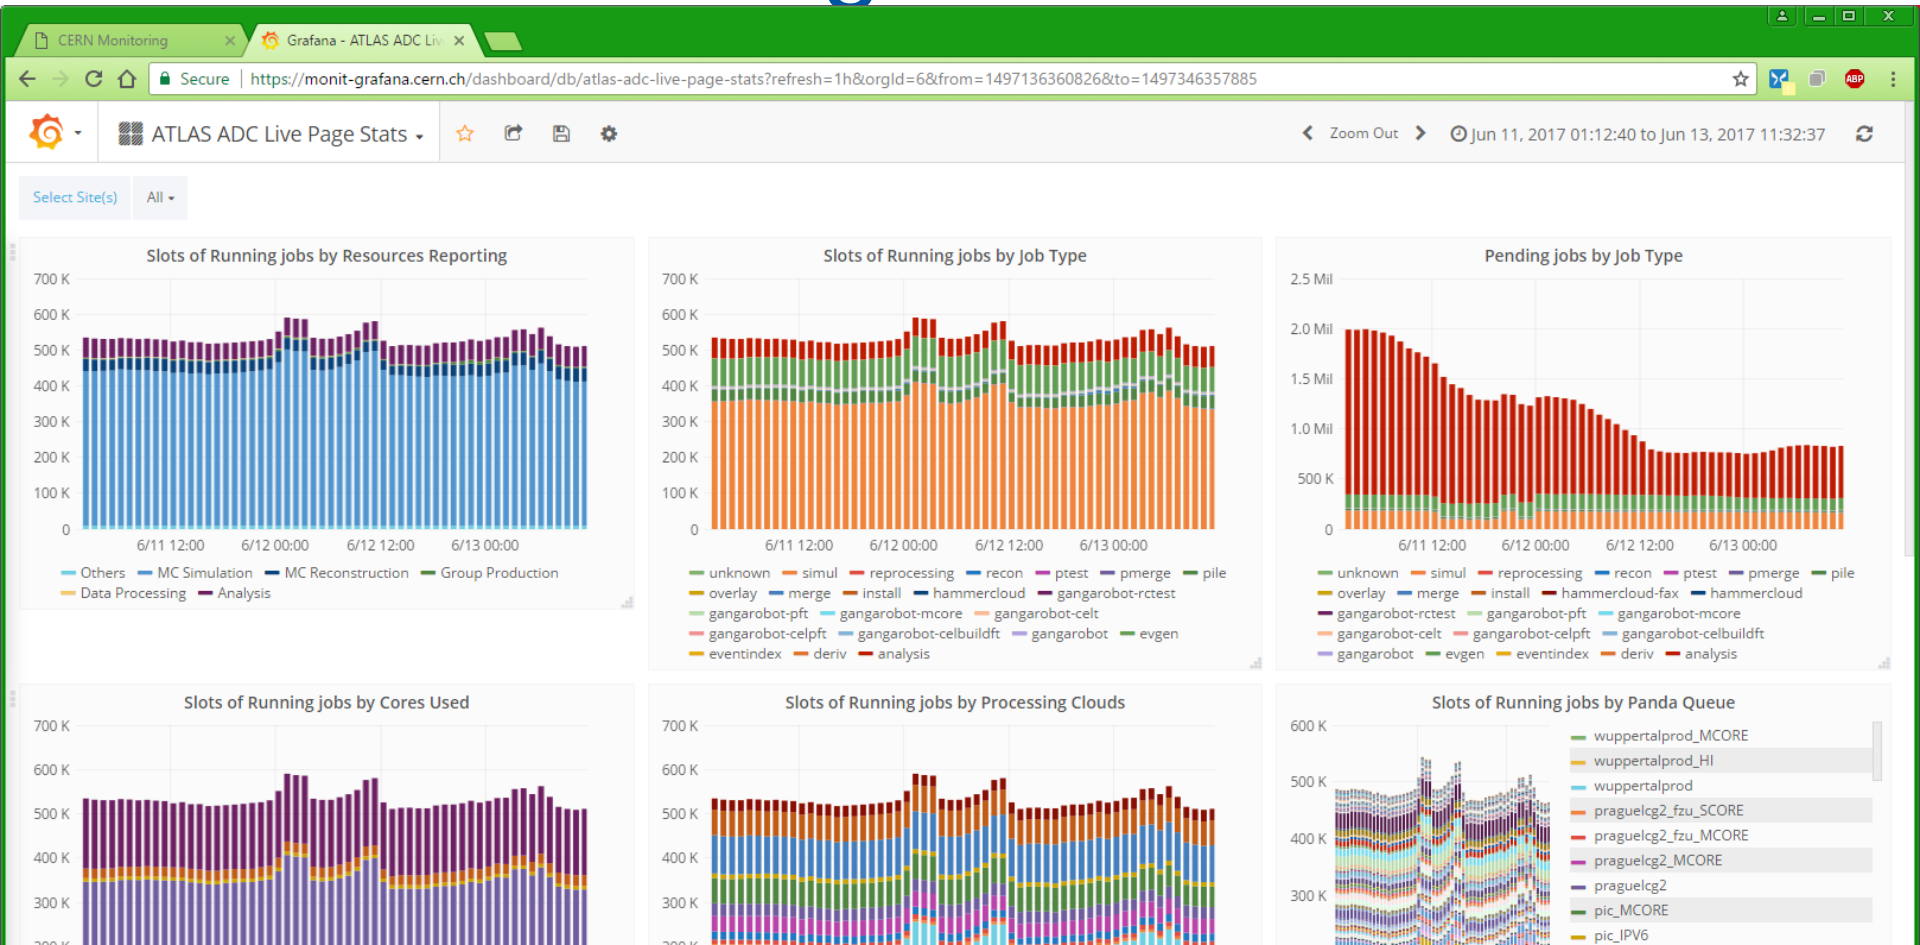

ATLAS Real Time

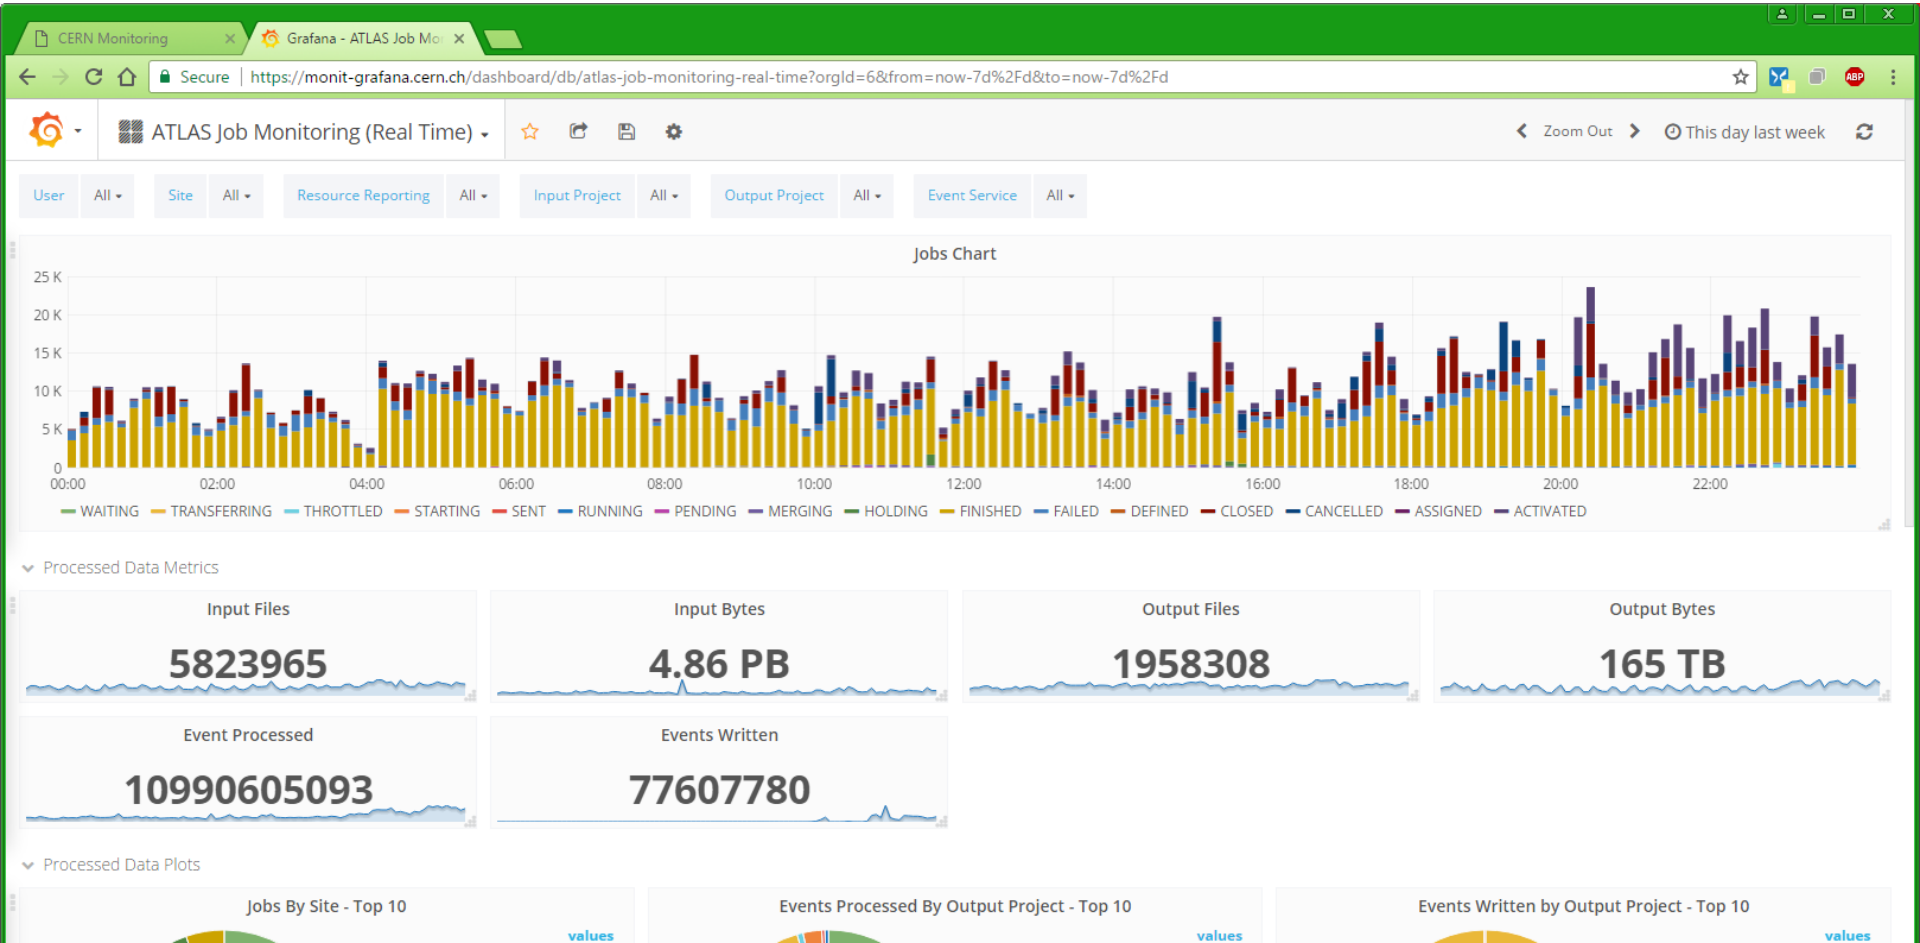

ATLAS DDM Dashboard

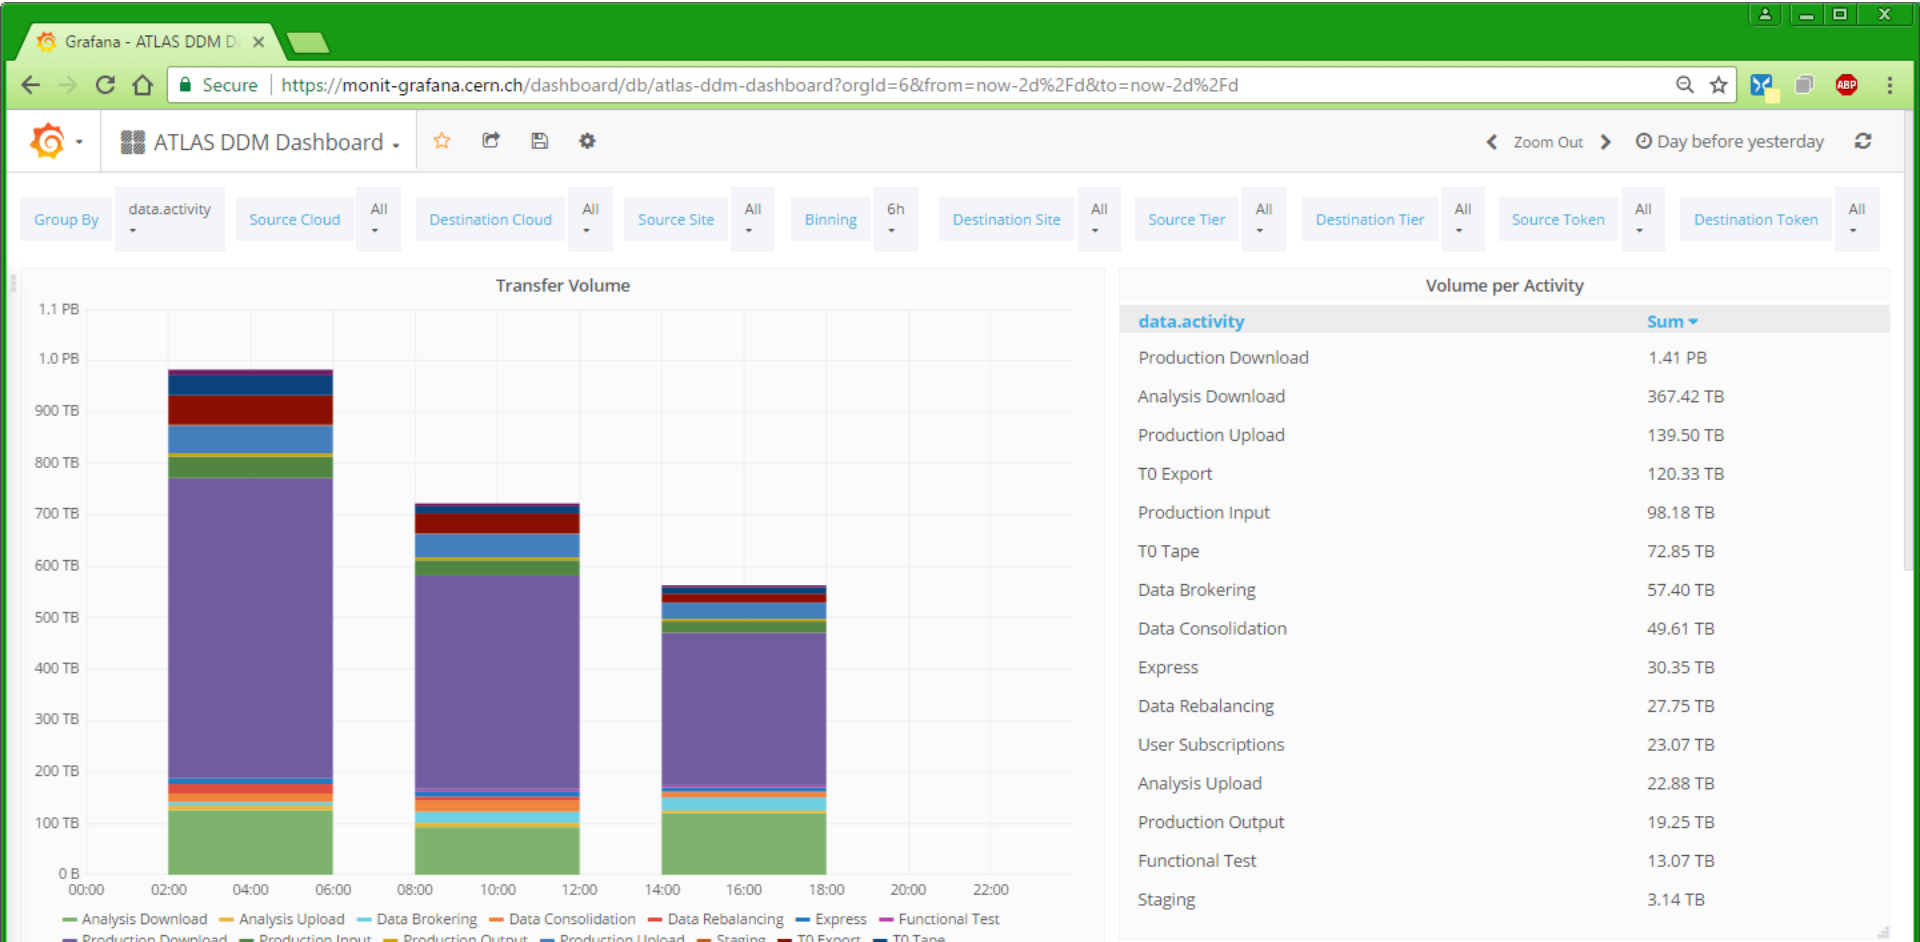

CMS Job Monitoring (new via HTCondor)

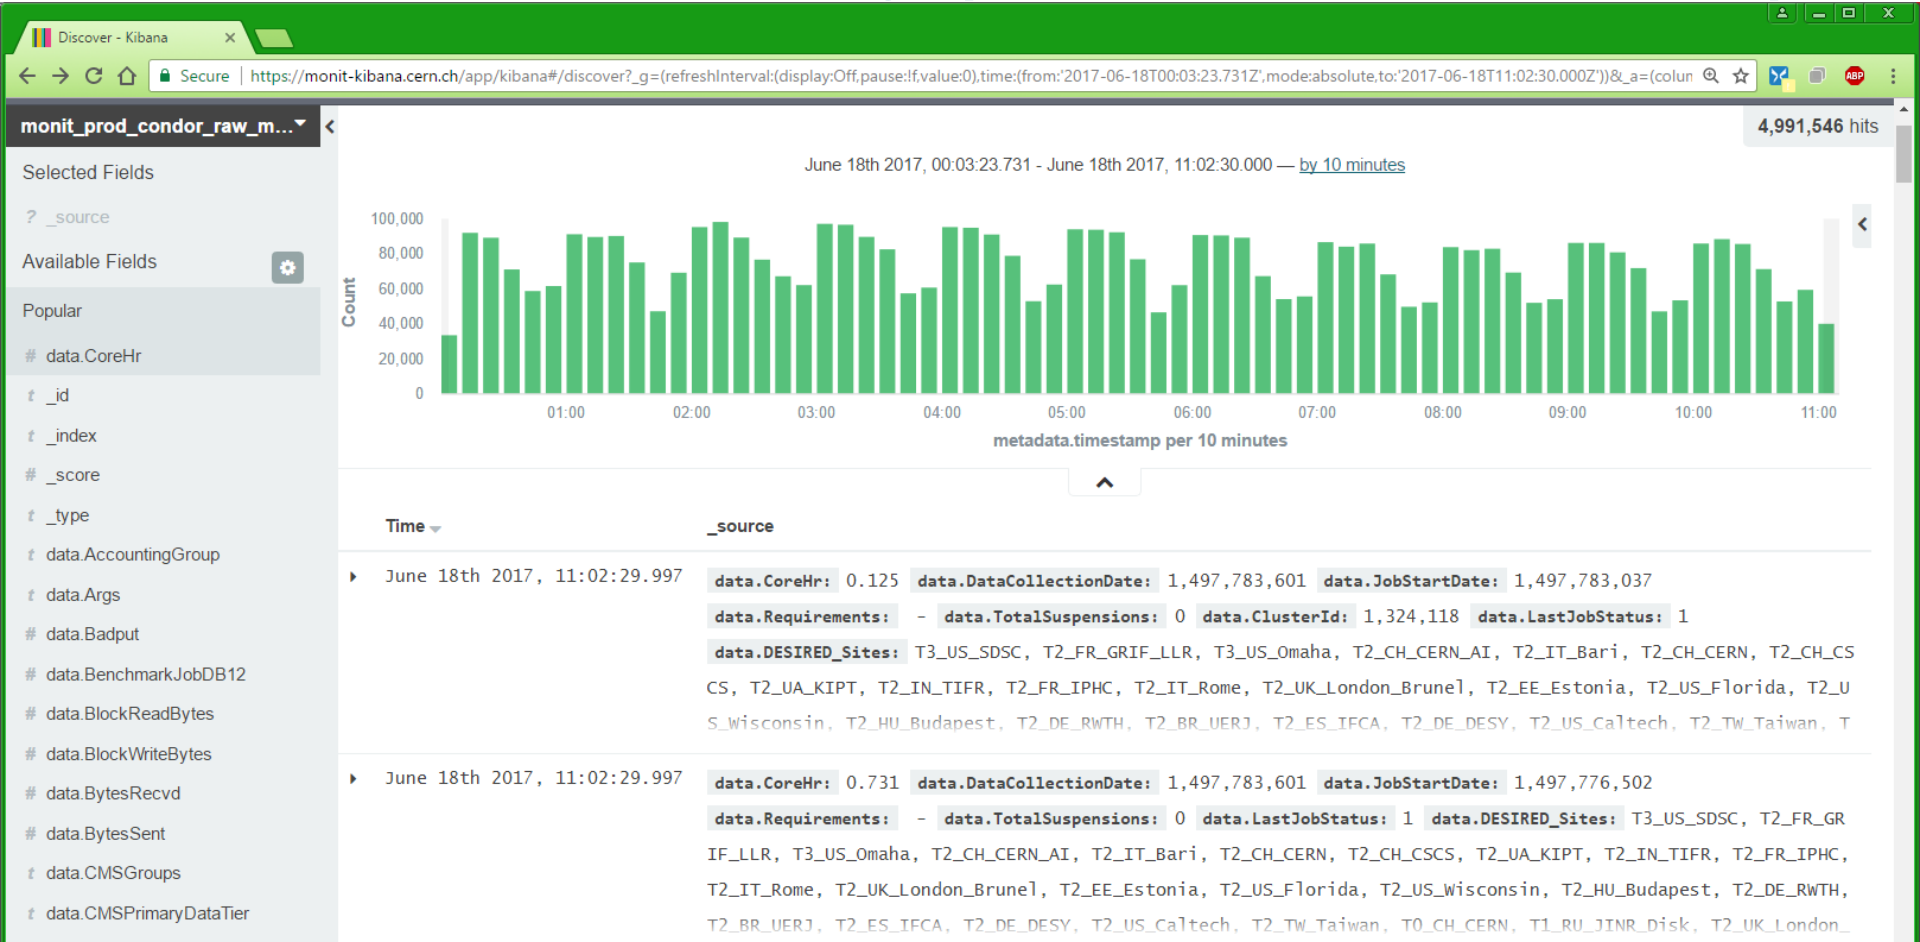

CRAB and ASO Operations

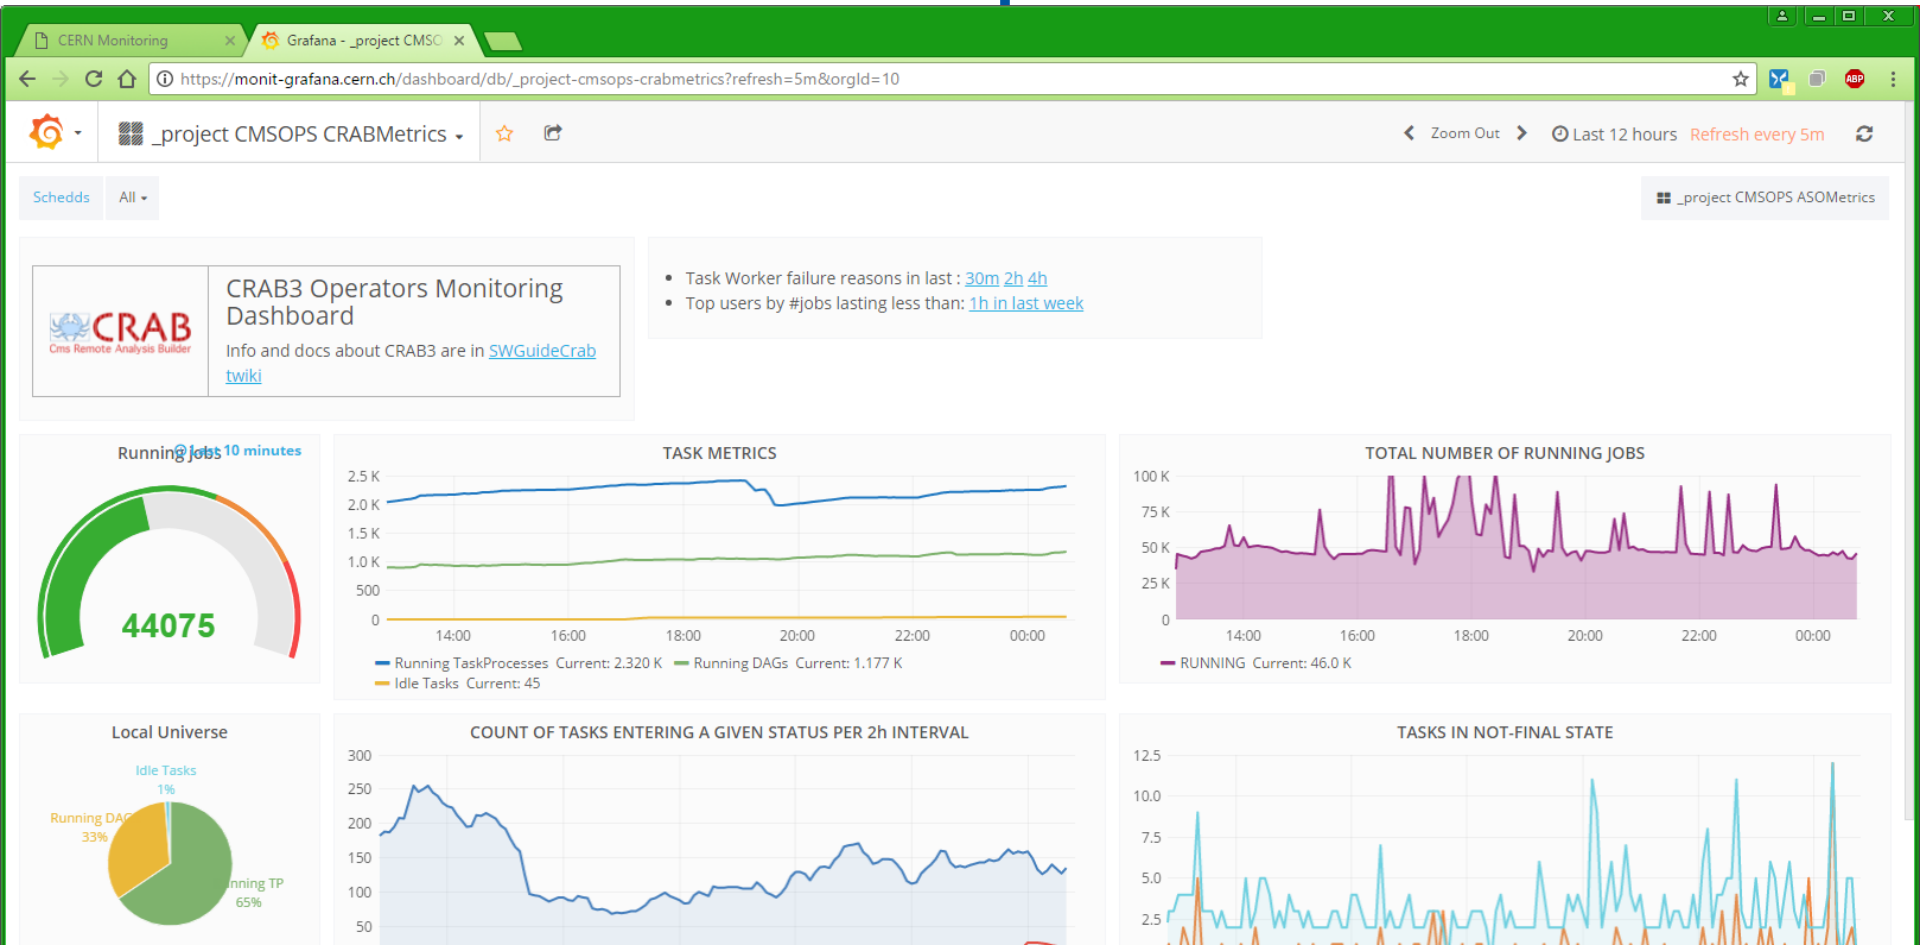

Portal: Other Dashboards

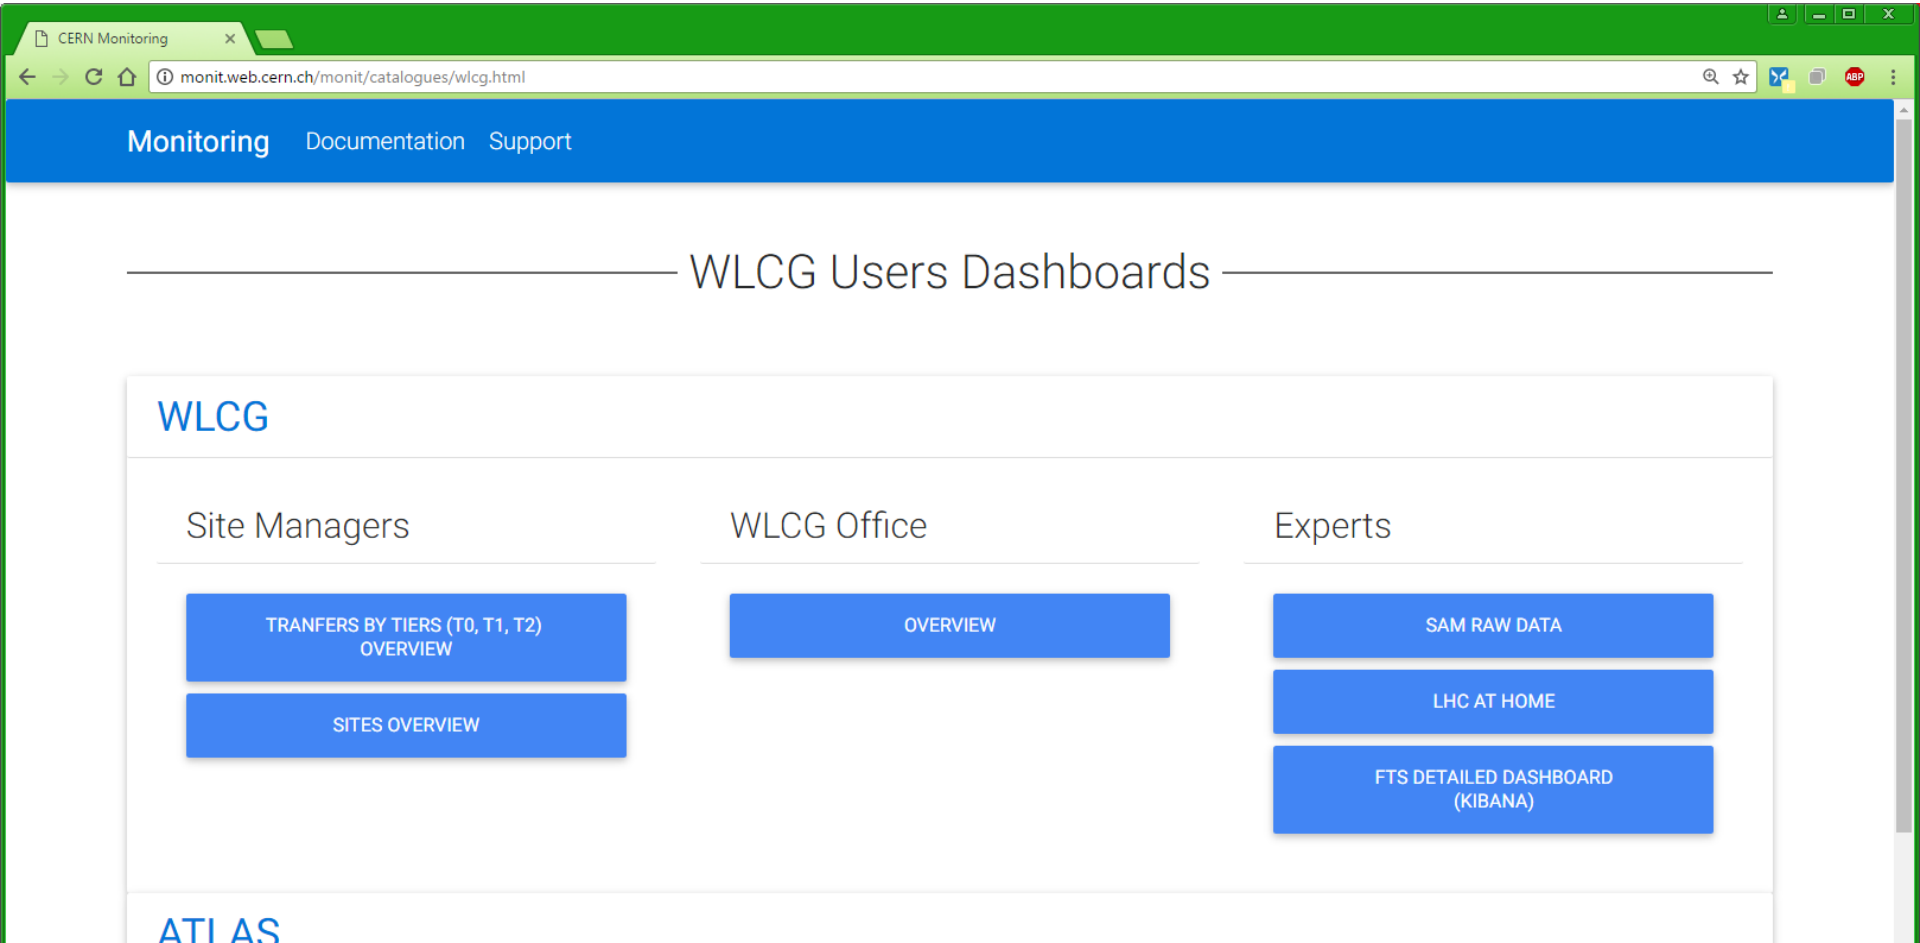

SAM Data

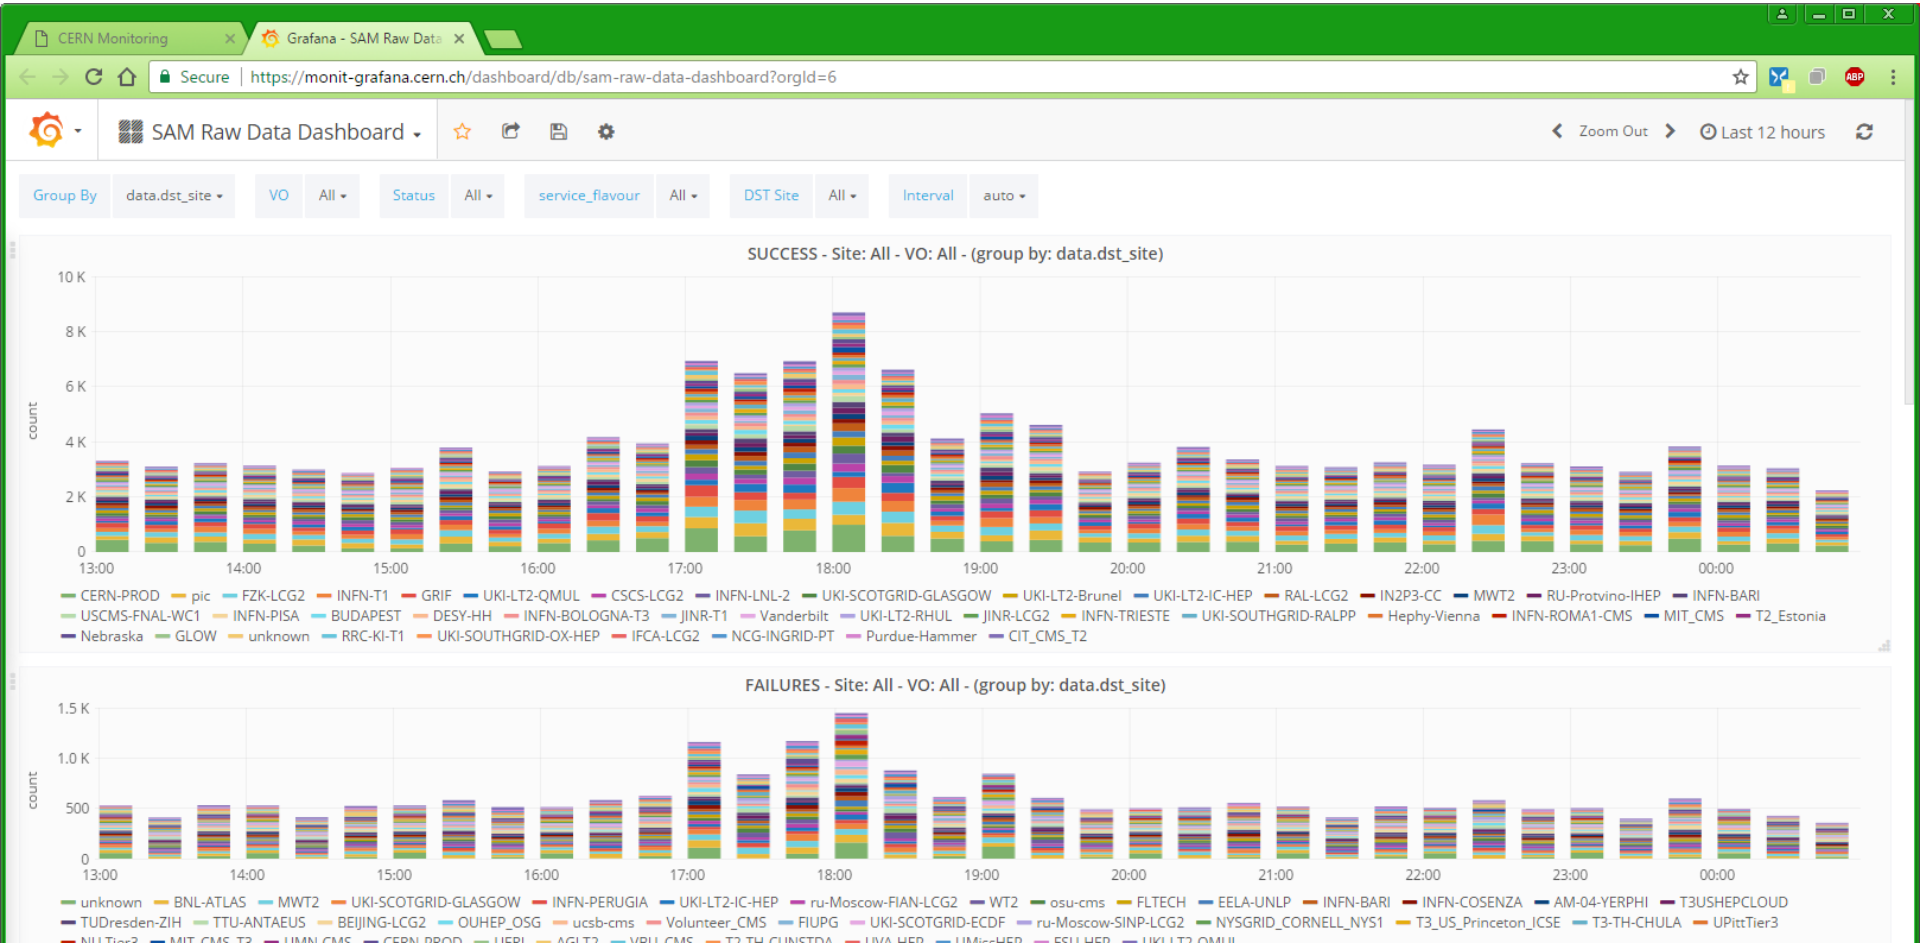
CERN Monitoring
Grafana - SAM Raw Data
Secure | https://monit-grafana.cern.ch/dashboard/db/sam-raw-data-dashboard?orgId=6
SAM Raw Data Dashboard
Zoom Out
Last 12 hours
Group By
data.dst_site
VO
All
Status
All
service_flavour
All
DST Site
All
Interval
auto
SUCCESS - Site: All - VO: All - (group by: data.dst_site)
count
10 K
8 K
6 K
4 K
2 K
0
13:00
14:00
15:00
16:00
17:00
18:00
19:00
20:00
21:00
22:00
23:00
00:00
CERN-PROD
pic
FZK-LCG2
INFN-T1
GRIF
UKI-LT2-QMUL
CSCS-LCG2
INFN-LNL-2
UKI-SCOTGRID-GLASGOW
UKI-LT2-Brunel
UKI-LT2-IC-HEP
RAL-LCG2
IN2P3-CC
MWT2
RU-Protvino-IHEP
INFN-BARI
USCMS-FNAL-WC1
INFN-PISA
BUDAPEST
DESY-HH
INFN-BOLOGNA-T3
JINR-T1
Vanderbilt
UKI-LT2-RHUL
JINR-LCG2
INFN-TRIESTE
UKI-SOUTHGRID-RALPP
Hephy-Vienna
INFN-ROMA1-CMS
MIT_CMS
T2_Estonia
Nebraska
GLOW
unknown
RRC-KI-T1
UKI-SOUTHGRID-OX-HEP
IFCA-LCG2
NCG-INGRID-PT
Purdue-Hammer
CIT_CMS_T2
FAILURES - Site: All - VO: All - (group by: data.dst_site)
count
1.5 K
1.0 K
500
0
13:00
14:00
15:00
16:00
17:00
18:00
19:00
20:00
21:00
22:00
23:00
00:00
unknown
BNL-ATLAS
MWT2
UKI-SCOTGRID-GLASGOW
INFN-PERUGIA
UKI-LT2-IC-HEP
ru-Moscow-FIAN-LCG2
WT2
osu-cms
FLTECH
EELA-UNLP
INFN-BARI
INFN-COSENZA
AM-04-YERPHI
T3USHEPCLOUD
TUDresden-ZIH
TTU-ANTAEUS
BEIJING-LCG2
OUHEP_OSG
ucsb-cms
Volunteer_CMS
FIUPG
UKI-SCOTGRID-ECDF
ru-Moscow-SINP-LCG2
NYSGRID_CORNELL_NYS1
T3_US_Princeton_ICSE
T3-TH-CHULA
UPittTier3

SAM - Availability Profiles

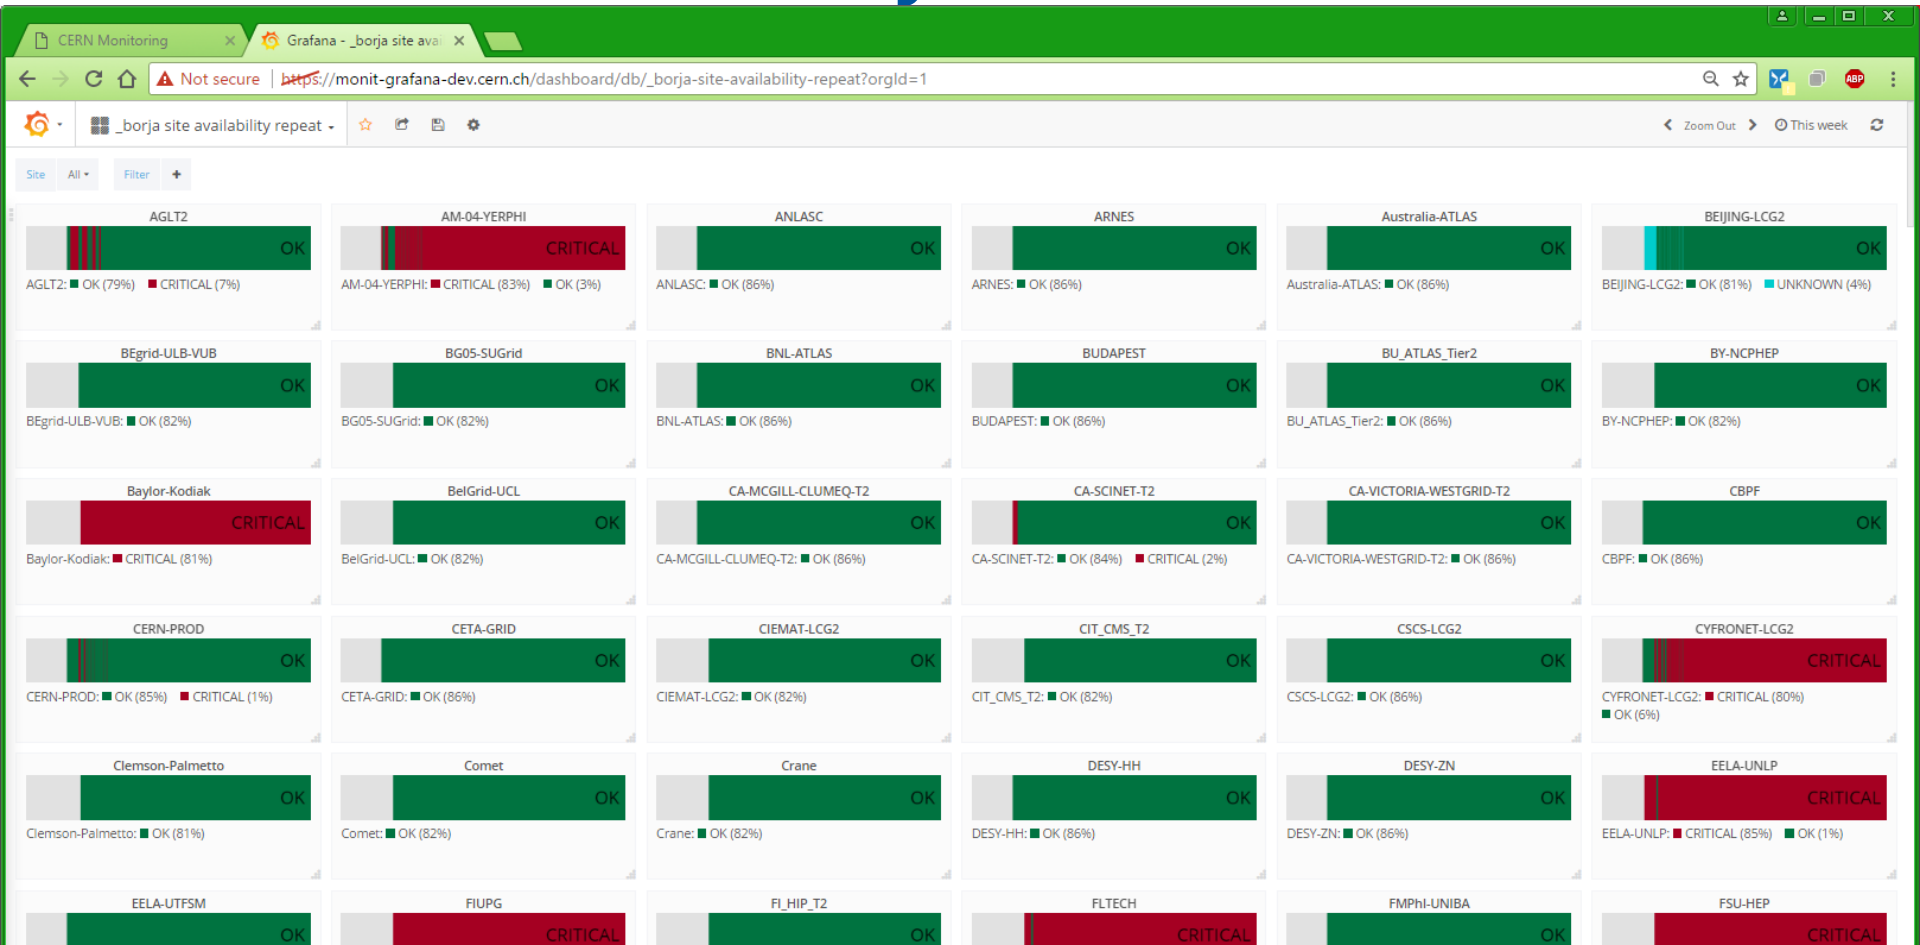

SSB Functionalities in MONIT

- Many WLCG metrics are already in
- Sites/Experiments can already inject SSB metrics in MONIT
- Instead of XML send the same data in JSON
- Use standard MONIT features, processing, dashboards, visualization, extraction
  - e.g. GlideInWMS by S.Lammel

Other Data in MONIT for WLCG

- Several other sources
  - LHCOPN network traffic
  - LHC@HOME (BOINC)
  - WLCG Space Accounting
  - Several CMS data sources
  - ATLAS ASAP metrics
  - OpenStack at CERN

LHCOPN

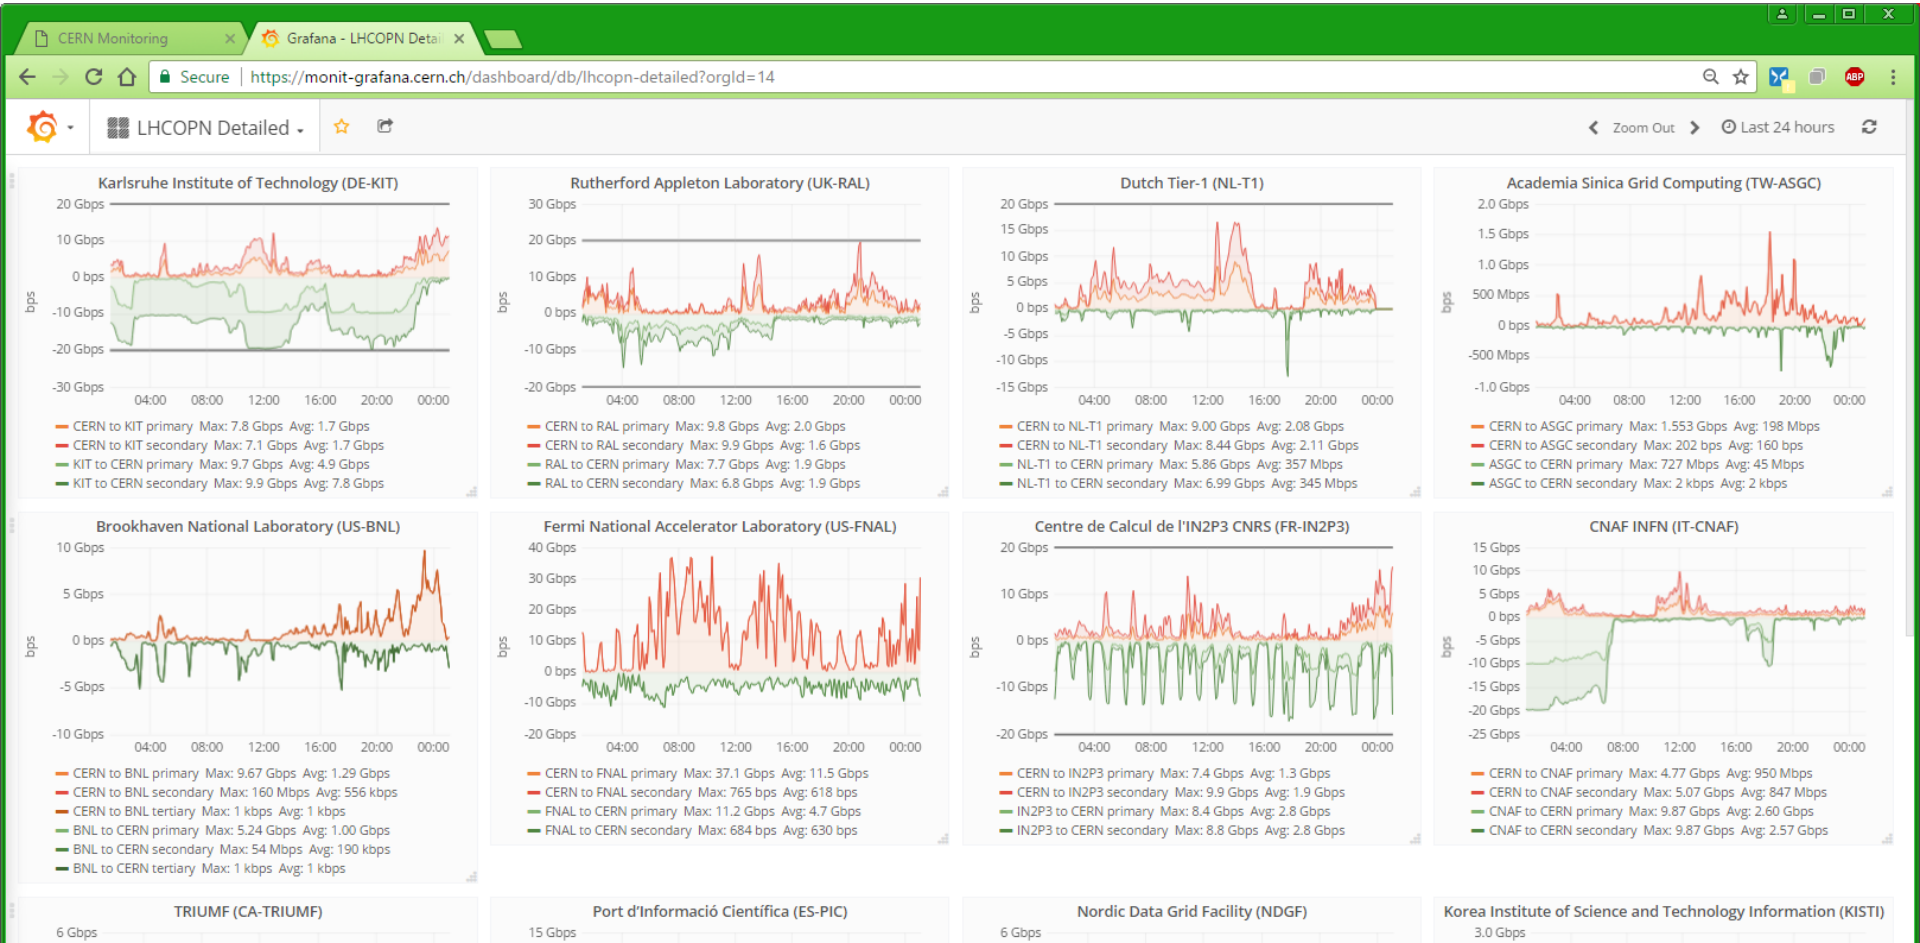
CERN Monitoring
Grafana - LHCOPN Detai
Secure | https://monit-grafana.cern.ch/dashboard/db/lhcopn-detailed?orgId=14
LHCOPN Detailed
Zoom Out
Last 24 hours
Karlsruhe Institute of Technology (DE-KIT)
CERN to KIT primary Max: 7.8 Gbps Avg: 1.7 Gbps
CERN to KIT secondary Max: 7.1 Gbps Avg: 1.7 Gbps
KIT to CERN primary Max: 9.7 Gbps Avg: 4.9 Gbps
KIT to CERN secondary Max: 9.9 Gbps Avg: 7.8 Gbps
Rutherford Appleton Laboratory (UK-RAL)
CERN to RAL primary Max: 9.8 Gbps Avg: 2.0 Gbps
CERN to RAL secondary Max: 9.9 Gbps Avg: 1.6 Gbps
RAL to CERN primary Max: 7.7 Gbps Avg: 1.9 Gbps
RAL to CERN secondary Max: 6.8 Gbps Avg: 1.9 Gbps
Dutch Tier-1 (NL-T1)
CERN to NL-T1 primary Max: 9.00 Gbps Avg: 2.08 Gbps
CERN to NL-T1 secondary Max: 8.44 Gbps Avg: 2.11 Gbps
NL-T1 to CERN primary Max: 5.86 Gbps Avg: 357 Mbps
NL-T1 to CERN secondary Max: 6.99 Gbps Avg: 345 Mbps
Academia Sinica Grid Computing (TW-ASGC)
CERN to ASGC primary Max: 1.553 Gbps Avg: 198 Mbps
CERN to ASGC secondary Max: 202 bps Avg: 160 bps
ASGC to CERN primary Max: 727 Mbps Avg: 45 Mbps
ASGC to CERN secondary Max: 2 kbps Avg: 2 kbps
Brookhaven National Laboratory (US-BNL)
CERN to BNL primary Max: 9.67 Gbps Avg: 1.29 Gbps
CERN to BNL secondary Max: 160 Mbps Avg: 556 kbps
CERN to BNL tertiary Max: 1 kbps Avg: 1 kbps
BNL to CERN primary Max: 5.24 Gbps Avg: 1.00 Gbps
BNL to CERN secondary Max: 54 Mbps Avg: 190 kbps
BNL to CERN tertiary Max: 1 kbps Avg: 1 kbps
Fermi National Accelerator Laboratory (US-FNAL)
CERN to FNAL primary Max: 37.1 Gbps Avg: 11.5 Gbps
CERN to FNAL secondary Max: 765 bps Avg: 618 bps
FNAL to CERN primary Max: 11.2 Gbps Avg: 4.7 Gbps
FNAL to CERN secondary Max: 684 bps Avg: 630 bps
Centre de Calcul de l'IN2P3 CNRS (FR-IN2P3)
CERN to IN2P3 primary Max: 7.4 Gbps Avg: 1.3 Gbps
CERN to IN2P3 secondary Max: 9.9 Gbps Avg: 1.9 Gbps
IN2P3 to CERN primary Max: 8.4 Gbps Avg: 2.8 Gbps
IN2P3 to CERN secondary Max: 8.8 Gbps Avg: 2.8 Gbps
CNAF INFN (IT-CNAF)
CERN to CNAF primary Max: 4.77 Gbps Avg: 950 Mbps
CERN to CNAF secondary Max: 5.07 Gbps Avg: 847 Mbps
CNAF to CERN primary Max: 9.87 Gbps Avg: 2.60 Gbps
CNAF to CERN secondary Max: 9.87 Gbps Avg: 2.57 Gbps
TRIUMF (CA-TRIUMF)
Port d'Informació Científica (ES-PIC)
Nordic Data Grid Facility (NDGF)
Korea Institute of Science and Technology Information (KISTI)

LHCOPN vs WLCG Transfers

WLCG Space Accounting (prototype)

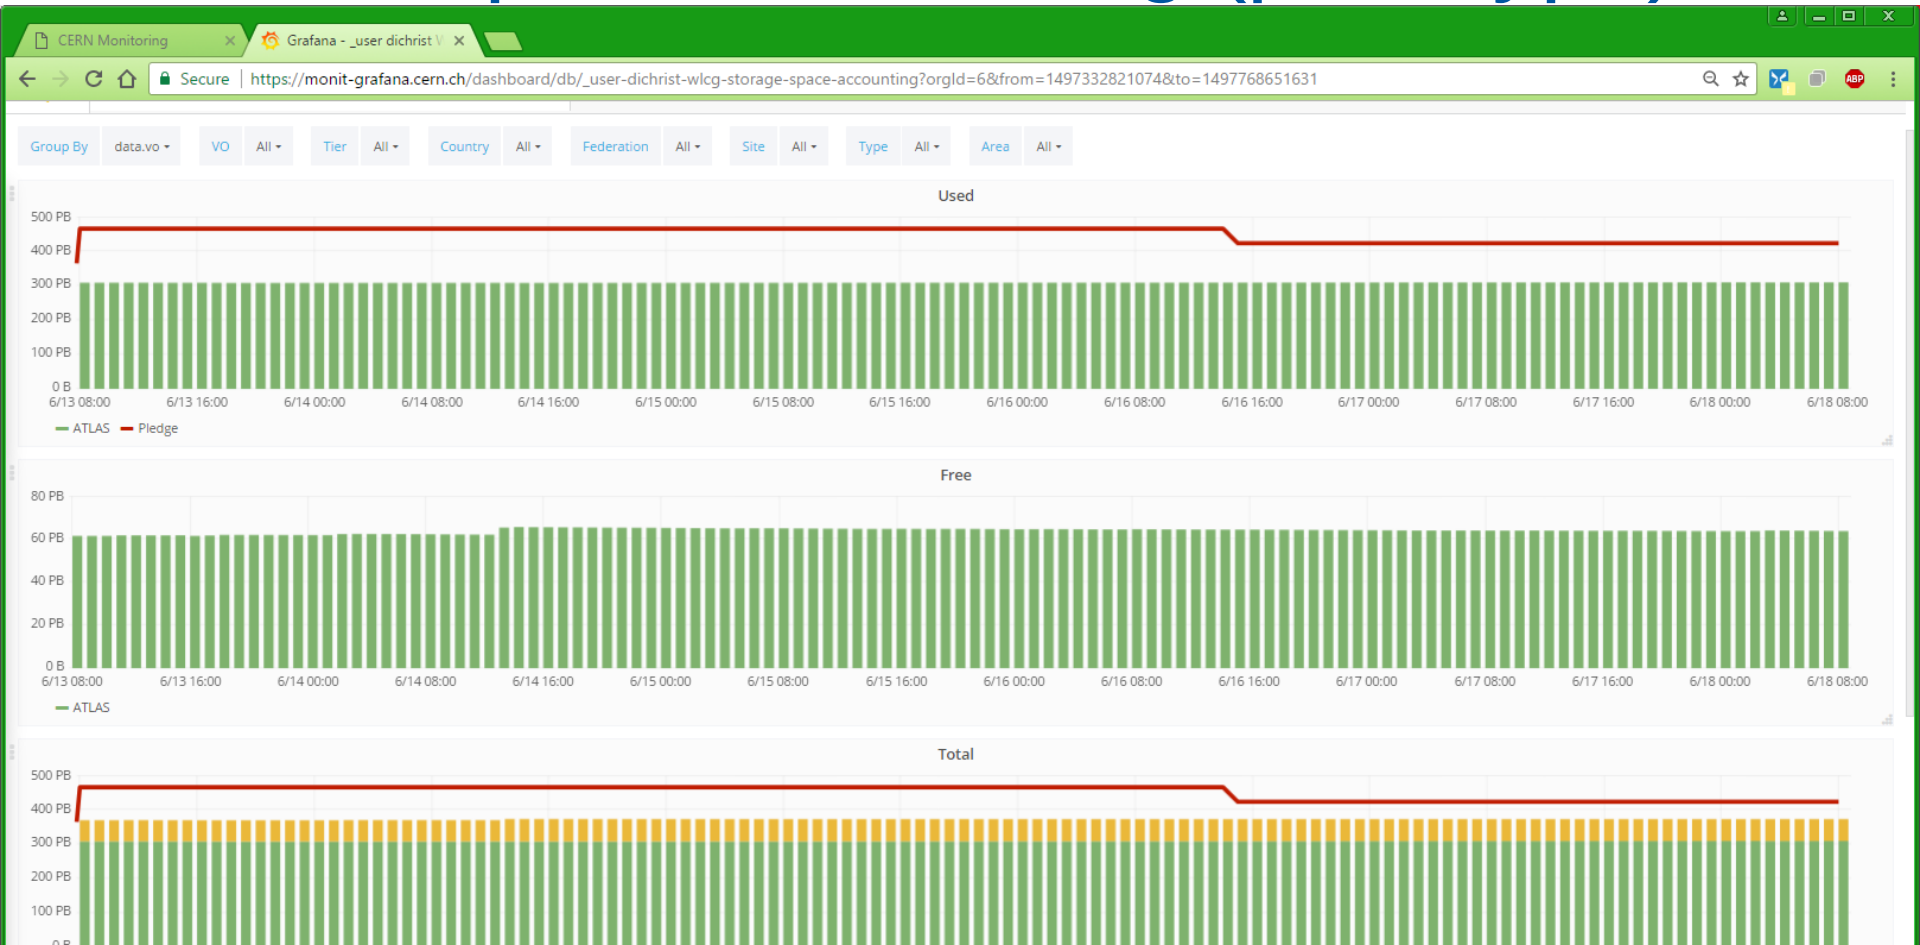

Strategy and Plans

WLCG Workflow in MONIT

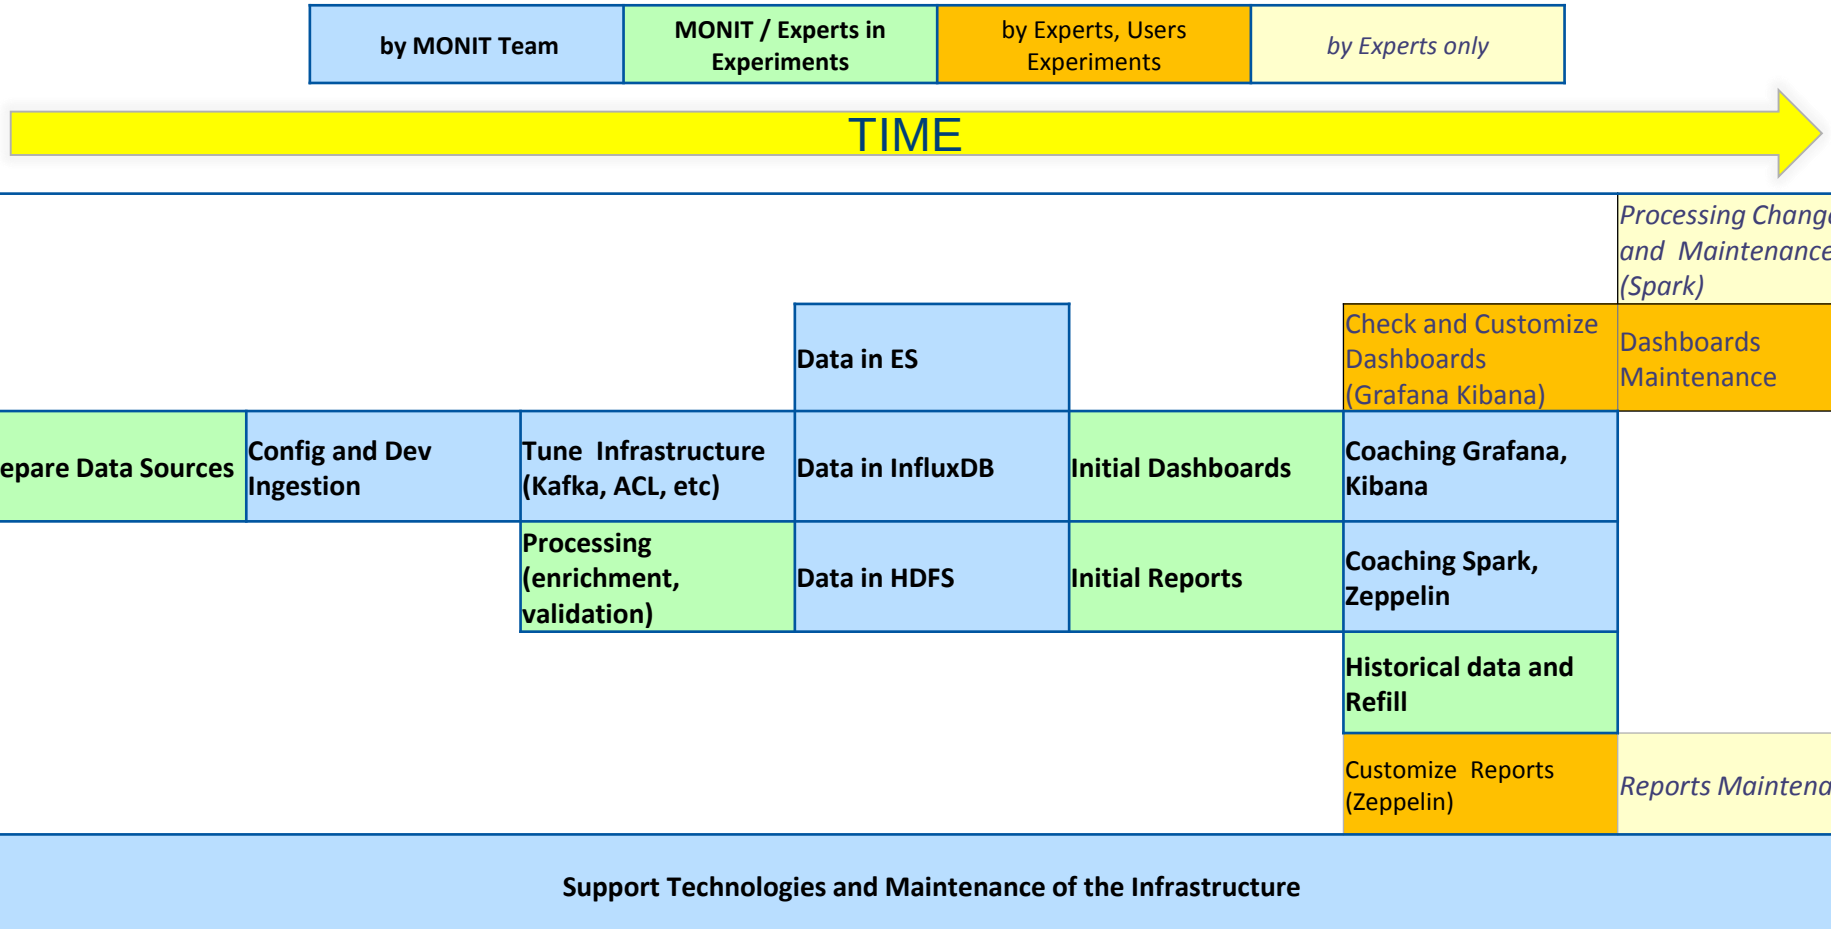

WLCG MONIT: Where we are now

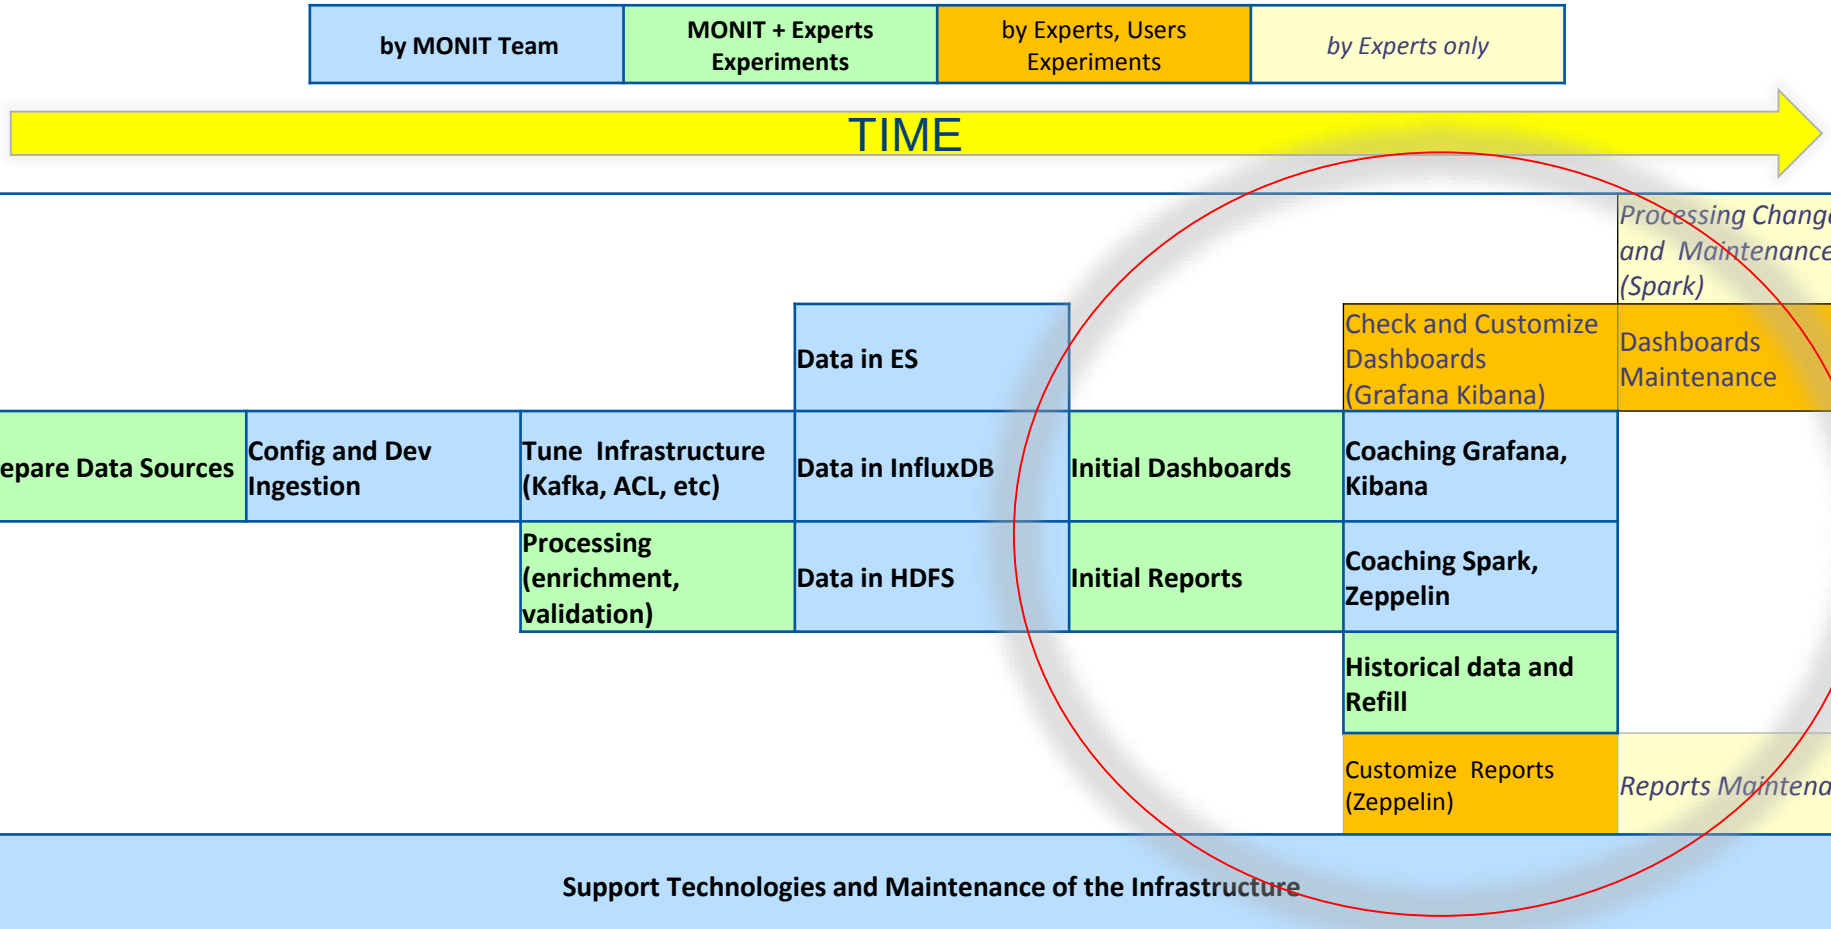

Working with Experiments

- WLCG Dashboards for all Experiments
  - Transfers, Sites Availability, etc.
- ATLAS
  - Fortnightly updates at ADC Analytics and Monitoring meetings
  - ATLAS Working Groups for Transfers, DDM, Jobs, and SAM
- CMS
  - Fortnightly updates at the CMS Monitoring Forum
  - Few advanced users. Several new data sources, dashboards.
- ALICE and LHCb
  - Presented MONIT and in periodic contact
  - Not using MONIT. New data sources easy to add, if useful
- Other experiments
  - ProtoDUNE: Sending logs and metrics (prototype)

Time Scale 2017H2 - 2018

2017H2

- Data already in MONIT. Processing completed. Initial dashboards are there.
- Import all historical data is ongoing
- Ready to provide training, on dashboards and reports
- Running the two WLCG infrastructures in parallel
- Needs help from experts to check and develop together final new dashboards
- Share with final users the new solutions (shifters, sites, managers, experts, etc.)

2018

- Progressively stop supporting old WLCG dashboards and reports
- Migrate final users. First Transfers, then Jobs, then SAM
- MONIT team can focus more on tuning and operating the infrastructure, supporting users and training experts

Next Steps

Focus on tuning the MONIT infrastructure

- performance, special types of plots, reports
- provide training sessions, dedicated coaching

Need help from the experts to complete migration of dashboards and reports

- Develop final real use cases
  - shifters, experts, managers, site managers
- Verify data quality and plots
- Extend and create new dashboards
- Spread the word
  - training, presentation in experiments, etc.

Reference and Contact

Dashboards (CERN SSO login)

monit.cern.ch

Feedback/Requests (SNOW)

cern.ch/monit-support

Documentation

cern.ch/monitdocs


CERN

Additional Info and Backup Slides

Sources, Storage, Visualization and Access

Data Sources	
Collects Metrics	Uses collectd plug-ins pre-installed for CERN host and services
Log Sources	Move logs data via a Flume local agent
HTTP Sources	End point for data from external sources (few, well connected)
Messaging Producers	AMQ Messaging end point to received data from external sources

Storage and Search	
ElasticSearch	Short-term storage and index (1 month, depends on resources available)
InfluxDB	Short-term time series storage (months, years aggregated)
HDFS	Long-term archive (years raw data)

Visualization and Access	
Kibana	Data from ElasticSearch. Full search/filter/discovery of data
Grafana	Data from ElasticSearch, InfluxDB. Dashboards optimized for time series plots
Zeppelin	Data from HDFS, ElasticSearch, InfluxDB. Notebooks for analysis, reports and plots Native support for Spark
API and CLIs	Access from external applications, scripts etc.

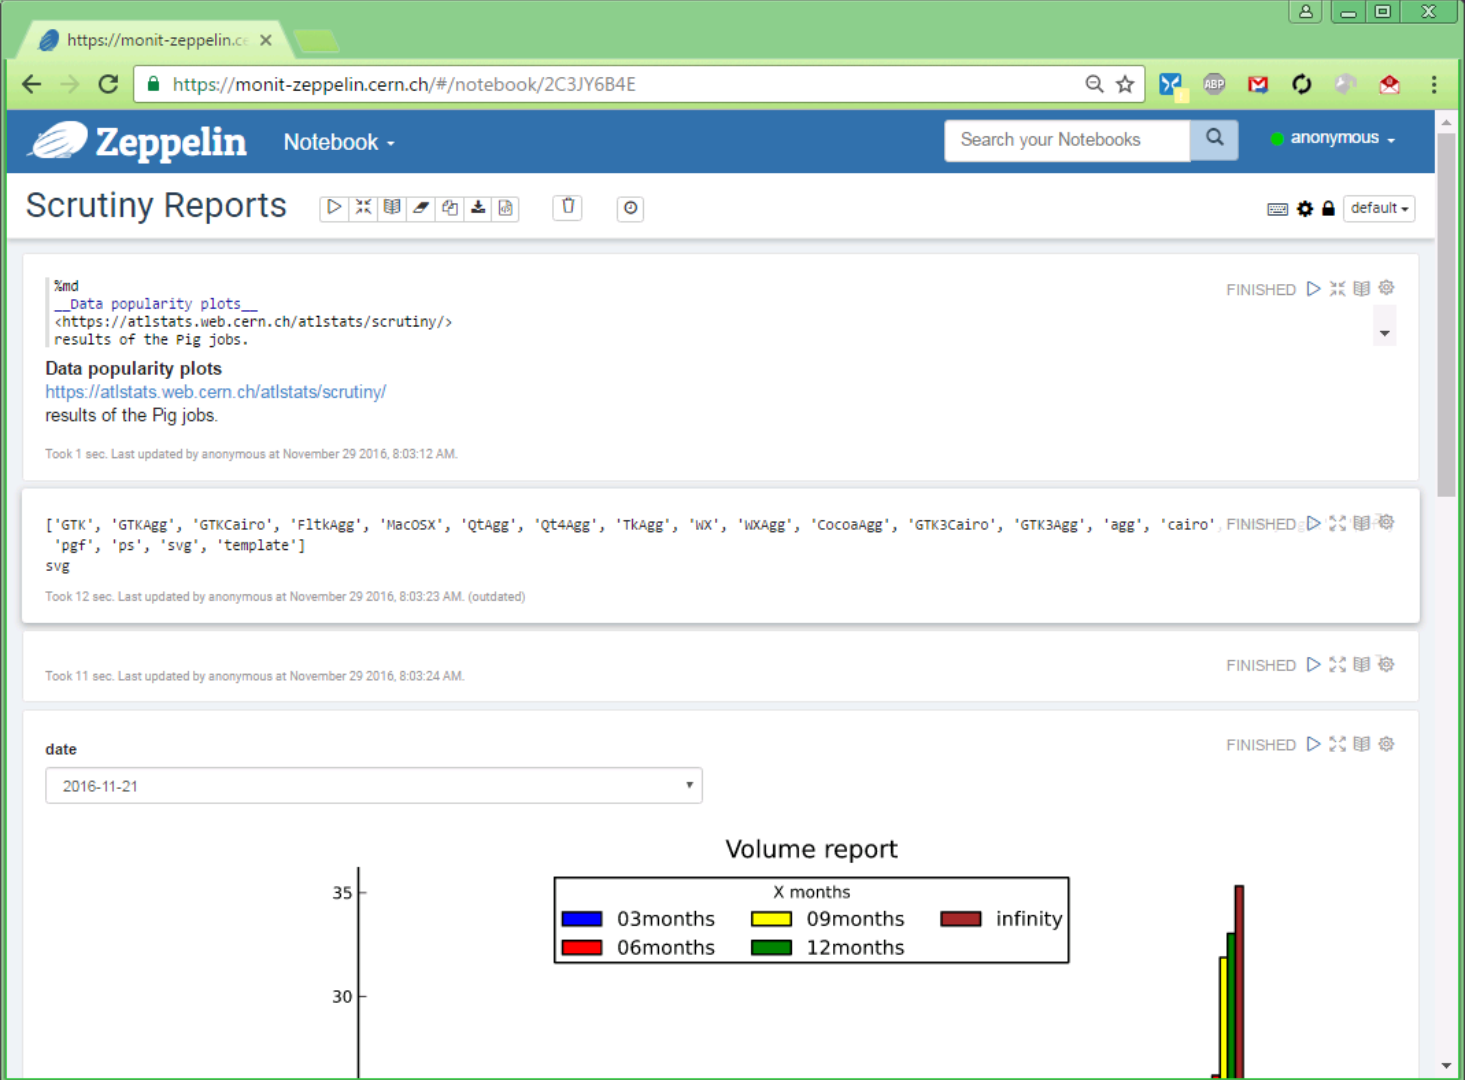

Notebooks with Zeppelin

Extract Data from HDFS or ES

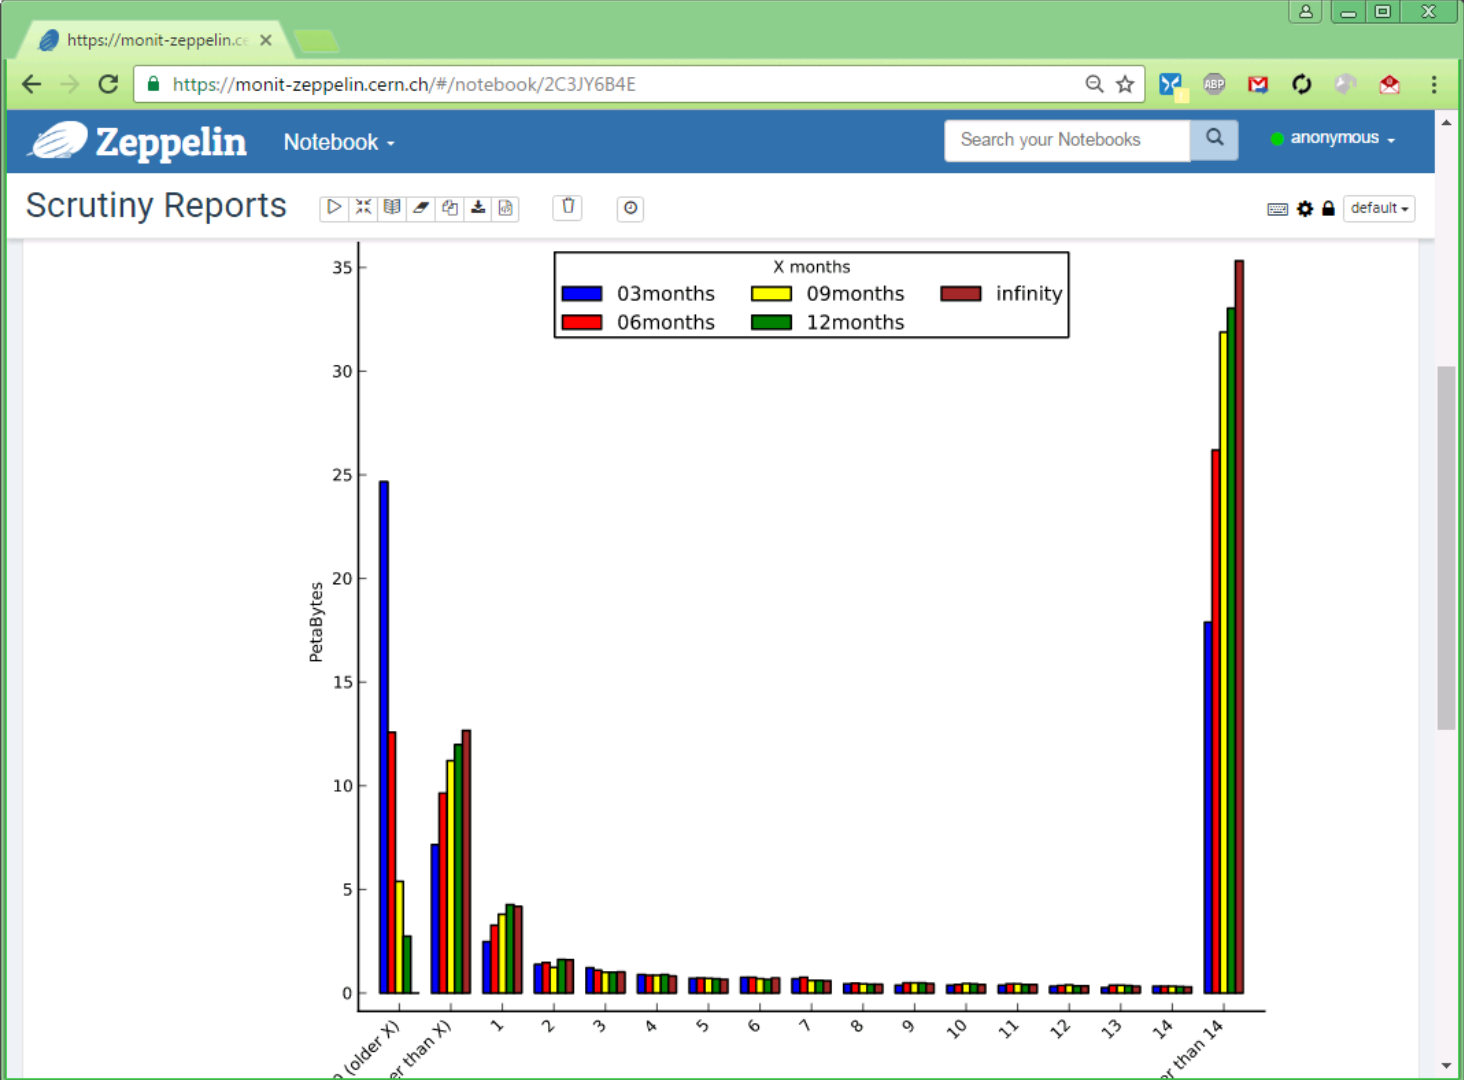

Manipulate the data and plot with common languages and tools

Python

Scala

numpy

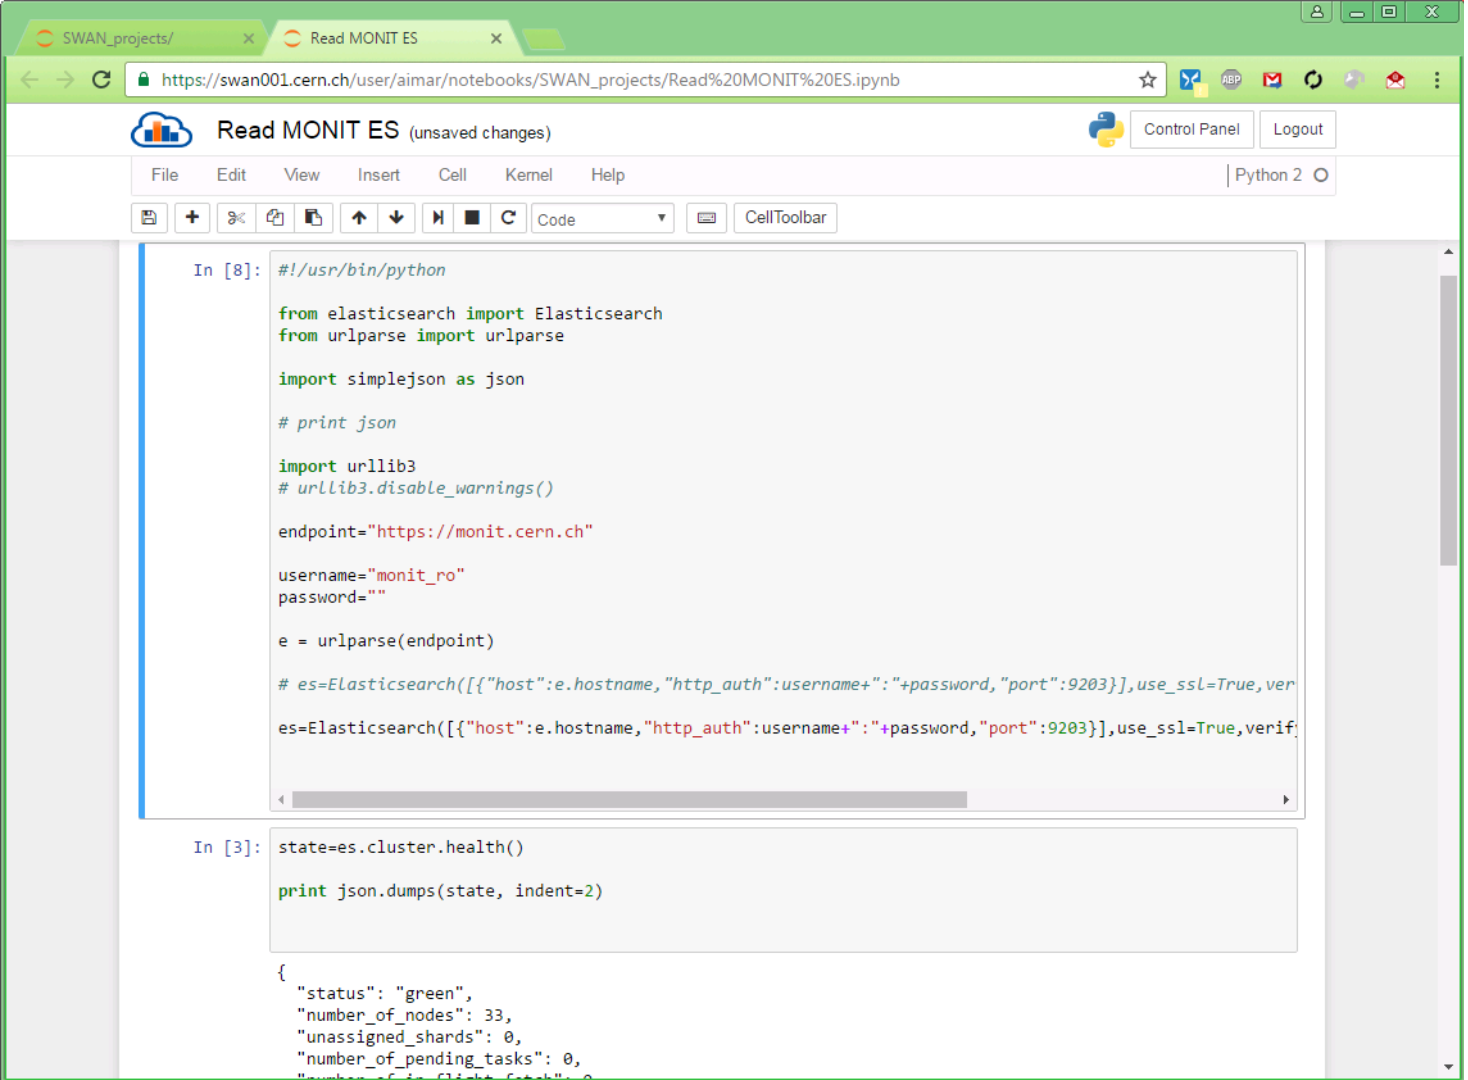

Notebooks with Swan

Extract Data from ES

ROOT

Python

C++

CVMFS

New DC Monitoring using Collectd

- Lemon Agent is the last component in production from the old Lemon/DC Monitoring
- Moving to collectd
  - collect system and service metrics
  - optimized to handle thousands of metrics
  - modular and portable with hundreds of plugins available
  - easy to develop new plugins in Python/Java/C
  - continuously improving and well documented

Moved Data Sources: IT and Logs

IT Data Sources	
	COLLECTD QA (3K HOSTS)
(tbd)	COLLECTD PROD (30K HOSTS)

Logs Data Sources	
	SYSLOG
	CASTOR
	EOS
	OPENSTACK
	HAMMERCLOUD
	FTS SERVERS
	DNS LOAD BALANCER
	BATCH
	PUNCH - PUPPET
	OPENSHIFT
	SQUID
	INSPIRE

Collect Data Source

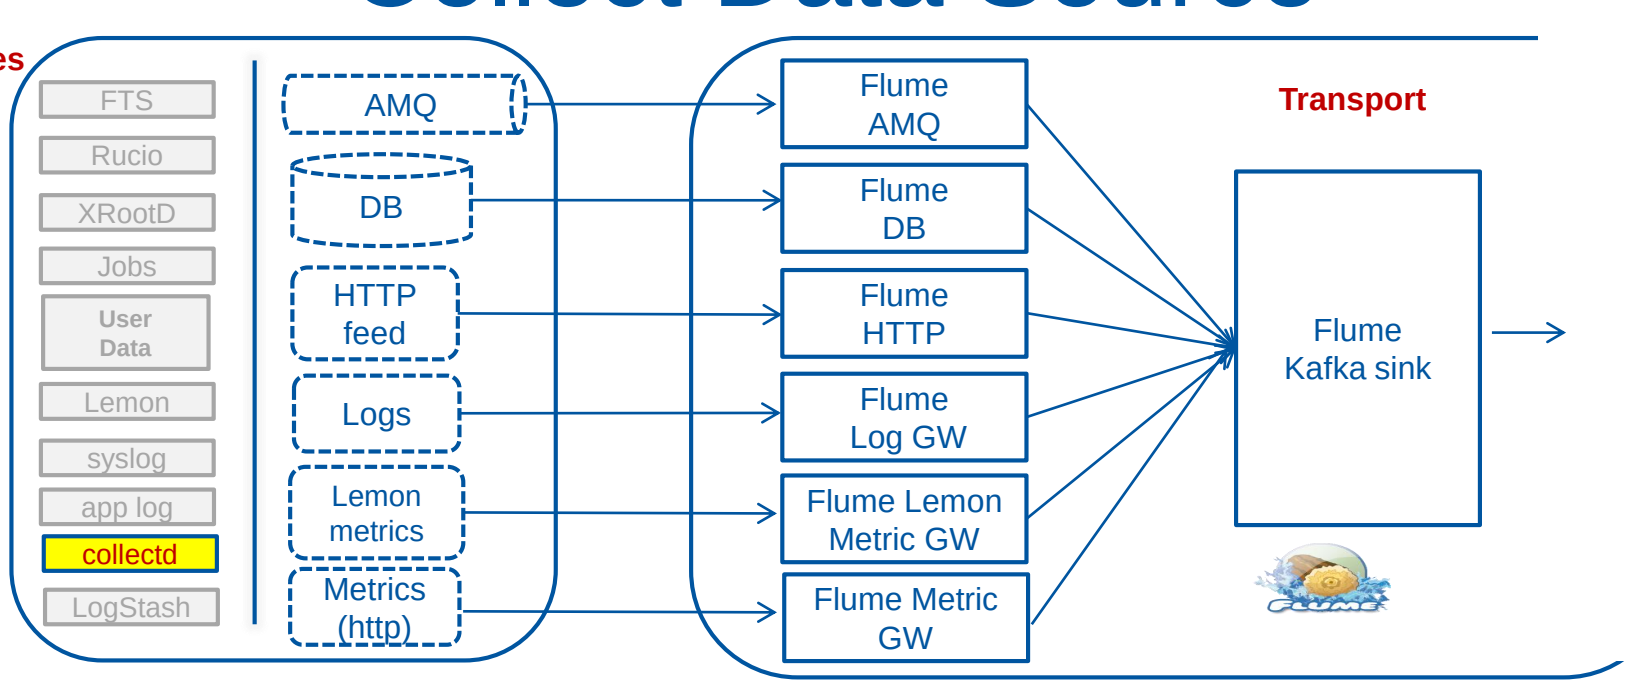

- Only component to add and use the full MONIT infrastructure
- All existing IT monitoring will be replaced (meter, notifications, dashbords)

Collectd Data Source

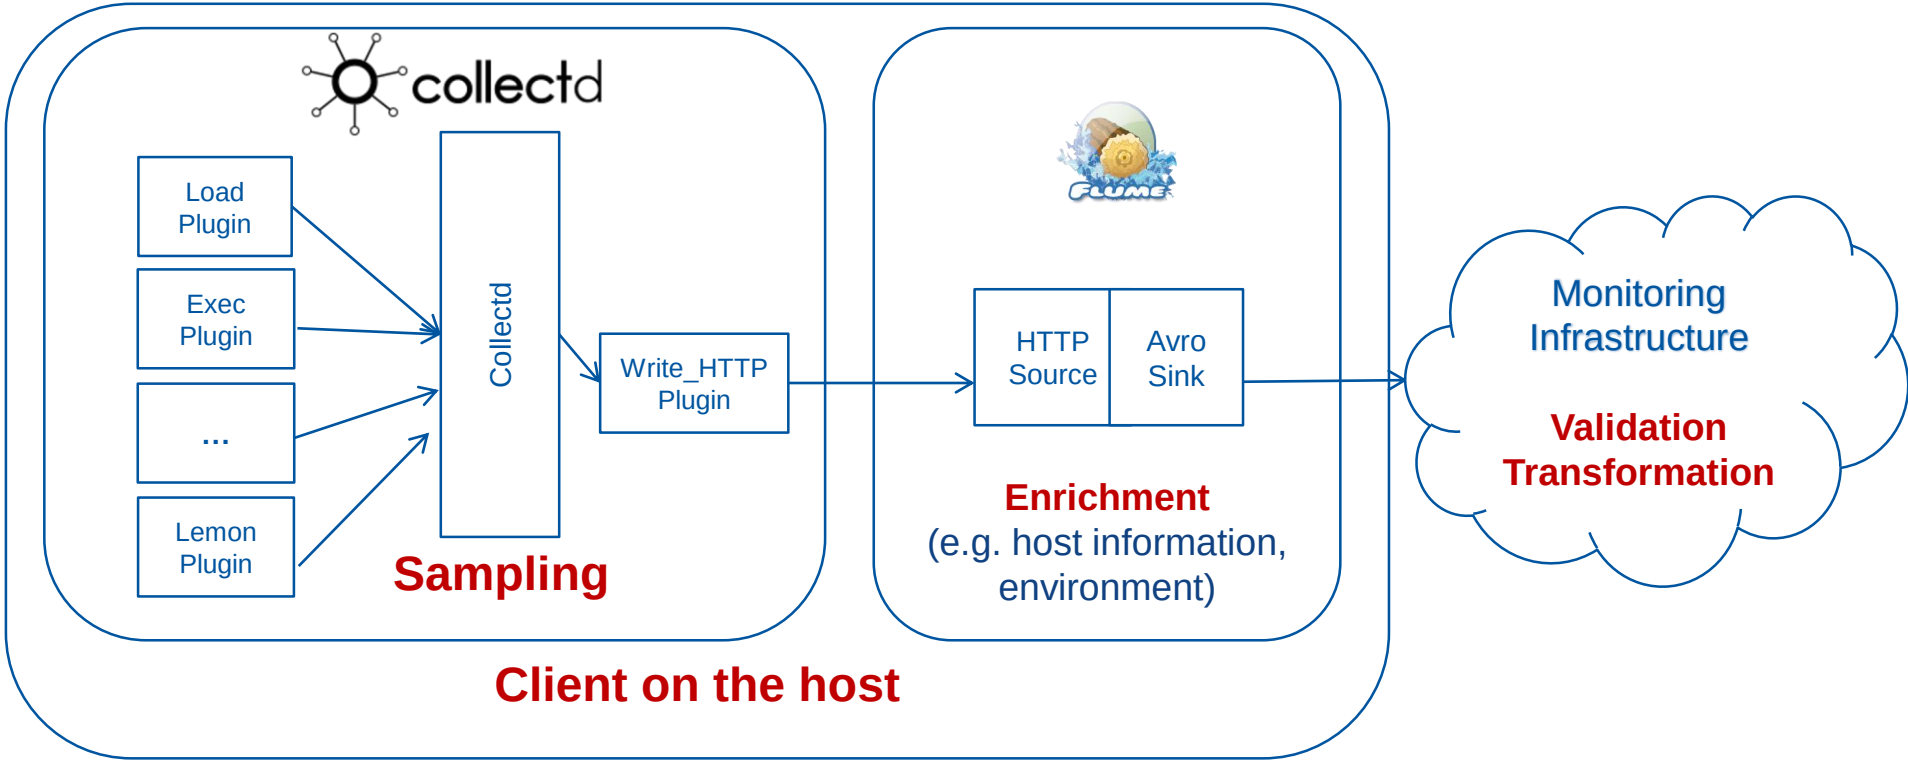

Collectd Metrics and Plugins

Lemon Metric Classes	Lemon Sensors	Collectd Support	Collectd Plugin
file.filecount	file	Official	https://collectd.org/wiki/index.php/Plugin:FileCount
file.size	file	Official	https://collectd.org/wiki/index.php/Plugin:FileCount
file.spaceUsed	file	Official	https://collectd.org/wiki/index.php/Plugin:FileCount
file.sslmtime	file	?	
log.Parse	parseLog	Official	https://collectd.org/wiki/index.php/Plugin:Tail
log.Parse	parseLog	Official	https://collectd.org/wiki/index.php/Plugin:Tail
cmd.ParseCmd	parse-cmd	Official	https://collectd.org/wiki/index.php/Plugin:Exec
system.bootTime	linux	Official	https://collectd.org/wiki/index.php/Plugin:Uptime
system.contextSwitches	linux	Official	https://collectd.org/wiki/index.php/Plugin:ContextSwitch
system.CPUCount	linux	Official	https://collectd.org/wiki/index.php/Plugin:CPU
system.CPUInfo	linux	?	
system.CPUutil	linux	Official	https://collectd.org/wiki/index.php/Plugin:CPU
system.CPUutilization	linux	Official	https://collectd.org/wiki/index.php/Plugin:CPU
system.createdProcesses	linux	Official	https://collectd.org/wiki/index.php/Plugin:Processes
system.diskStats	linux	Official	https://collectd.org/wiki/index.php/Plugin:Disk
system.existingProcesses	linux	Official	https://collectd.org/wiki/index.php/Plugin:Processes
system.exitCode	linux	Official	https://collectd.org/wiki/index.php/Plugin:Exec
system.fullLoadAvg	linux	Official	https://collectd.org/wiki/index.php/Plugin:Load
system.interrupts	linux	Official	https://collectd.org/wiki/index.php/Plugin:IRQ
system.loadAvg	linux	Official	https://collectd.org/wiki/index.php/Plugin:Load
system.meminfo	linux	Official	https://collectd.org/wiki/index.php/Plugin:Memory
system.memoryShared	linux	Official	https://collectd.org/wiki/index.php/Plugin:Memory
system.memoryStats	linux	Official	https://collectd.org/wiki/index.php/Plugin:Memory
system.networkInterfaceDropped	linux	Official	https://collectd.org/wiki/index.php/Plugin:Interface
system.networkInterfaceInfo	linux	?	
system.networkInterfaceIO	linux	Official	https://collectd.org/wiki/index.php/Plugin:Interface
system.numberOfSockets	linux	Official	https://collectd.org/wiki/index.php/Plugin:TCPConns
system.numberOfUsers	linux	Official	https://collectd.org/wiki/index.php/Plugin:Users

Replacement Strategy

1. Use an existing collectd plugin **(recommended)**
   - Straightforward: main logic can be reused
   - Many similarities at API level
     - registerMetric() => register_read()
     - storeSample() => dispatch()
2. Extend standard collectd plugin
   - Requires development
3. Run lemon sensor using collectd wrapper