



WLCG infrastructure monitoring proposal

Pablo Saiz
IT/SDC/MI

16th August 2013



Table of contents

- I. Summary of the progress
- II. Desired structure of applications
- III. Proposal for infrastructure monitoring

I. Summary

Motivation

- Reduction on number of people
- Redefining scope of applications
- Combining expertise
- Step out and evaluate other alternatives
- Goal:
 - Offer (at least) same QoS with less resources

Status so far

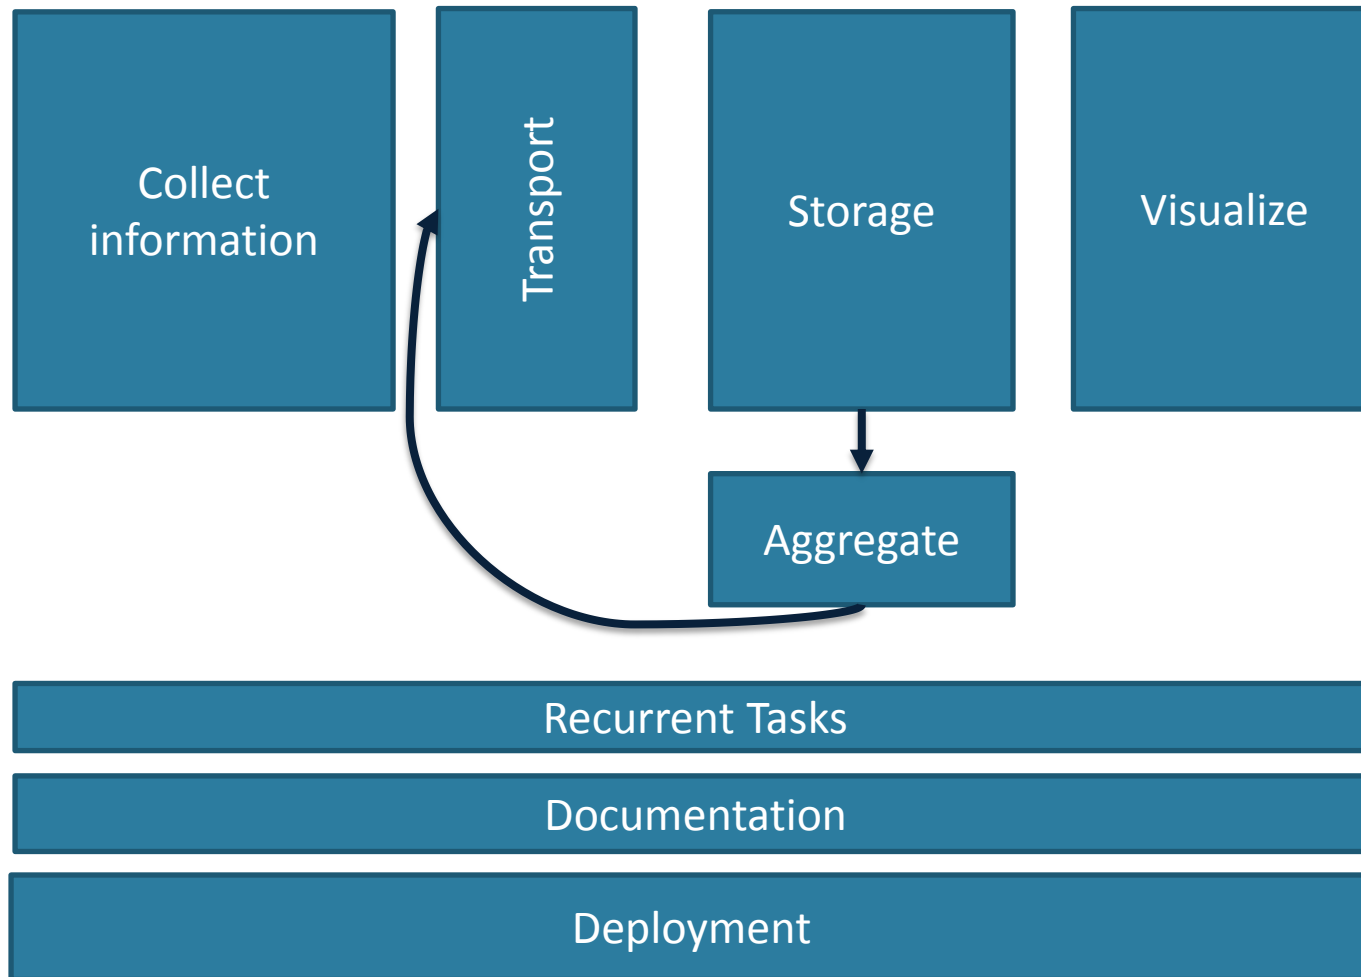
- WLCG monitoring consolidation group created
- Applications supported by the section
- Applications used
- ... so now we know what to provide

How to provide it

- Input from our experience
- Input from other groups
- What is available out there
 - Split in different areas of work
 - Source of Information
 - Visualization
 - Transport
 - Documentation
 - Storage
 - Deployment
 - Aggregation
 - Recurrent tasks
- Review of the areas

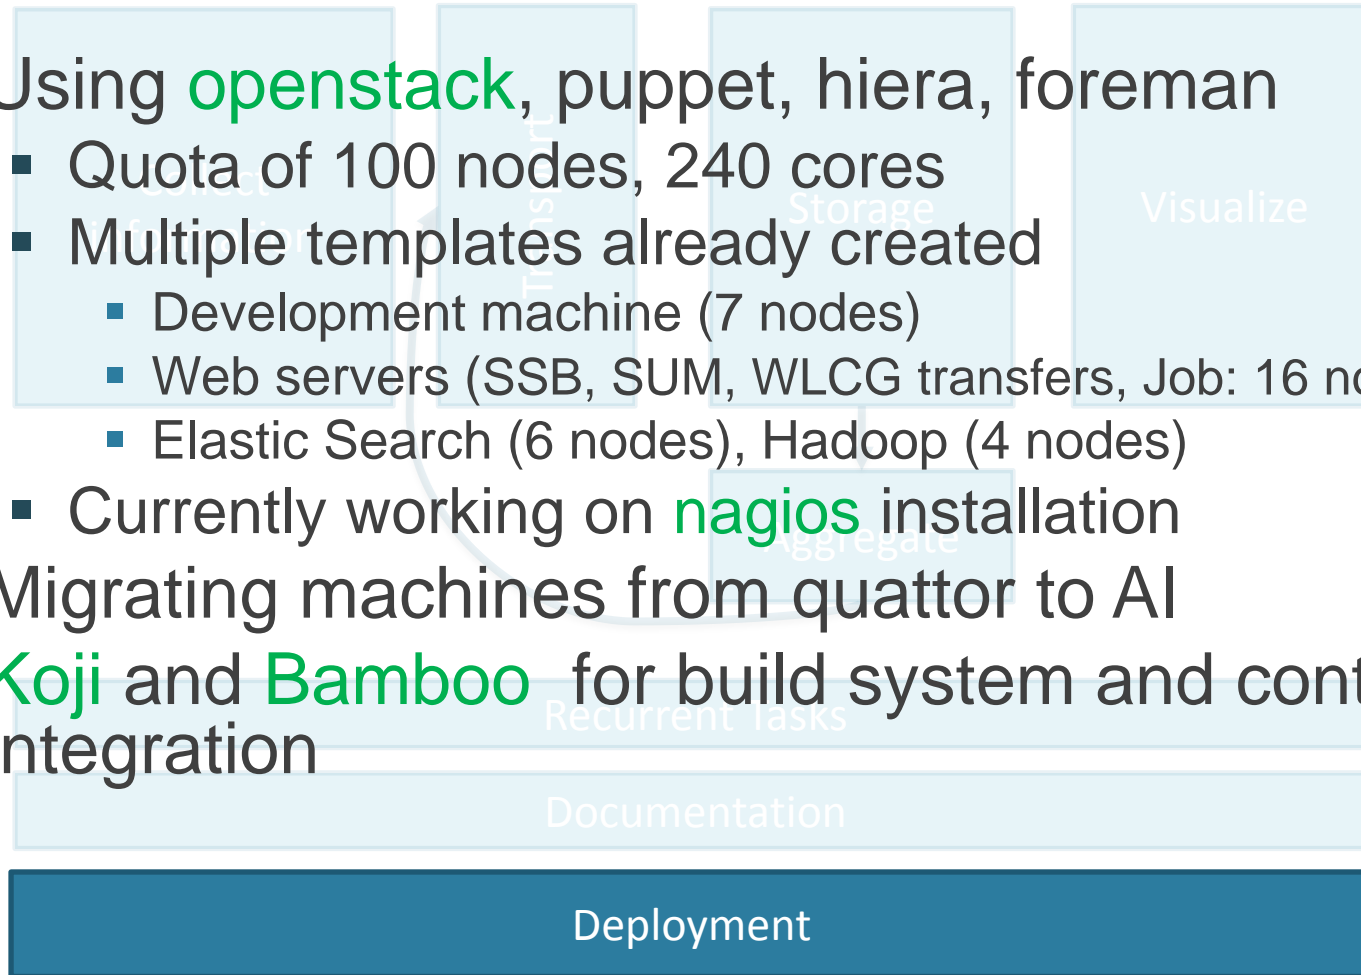
II. Structure of applications

Different layers of applications

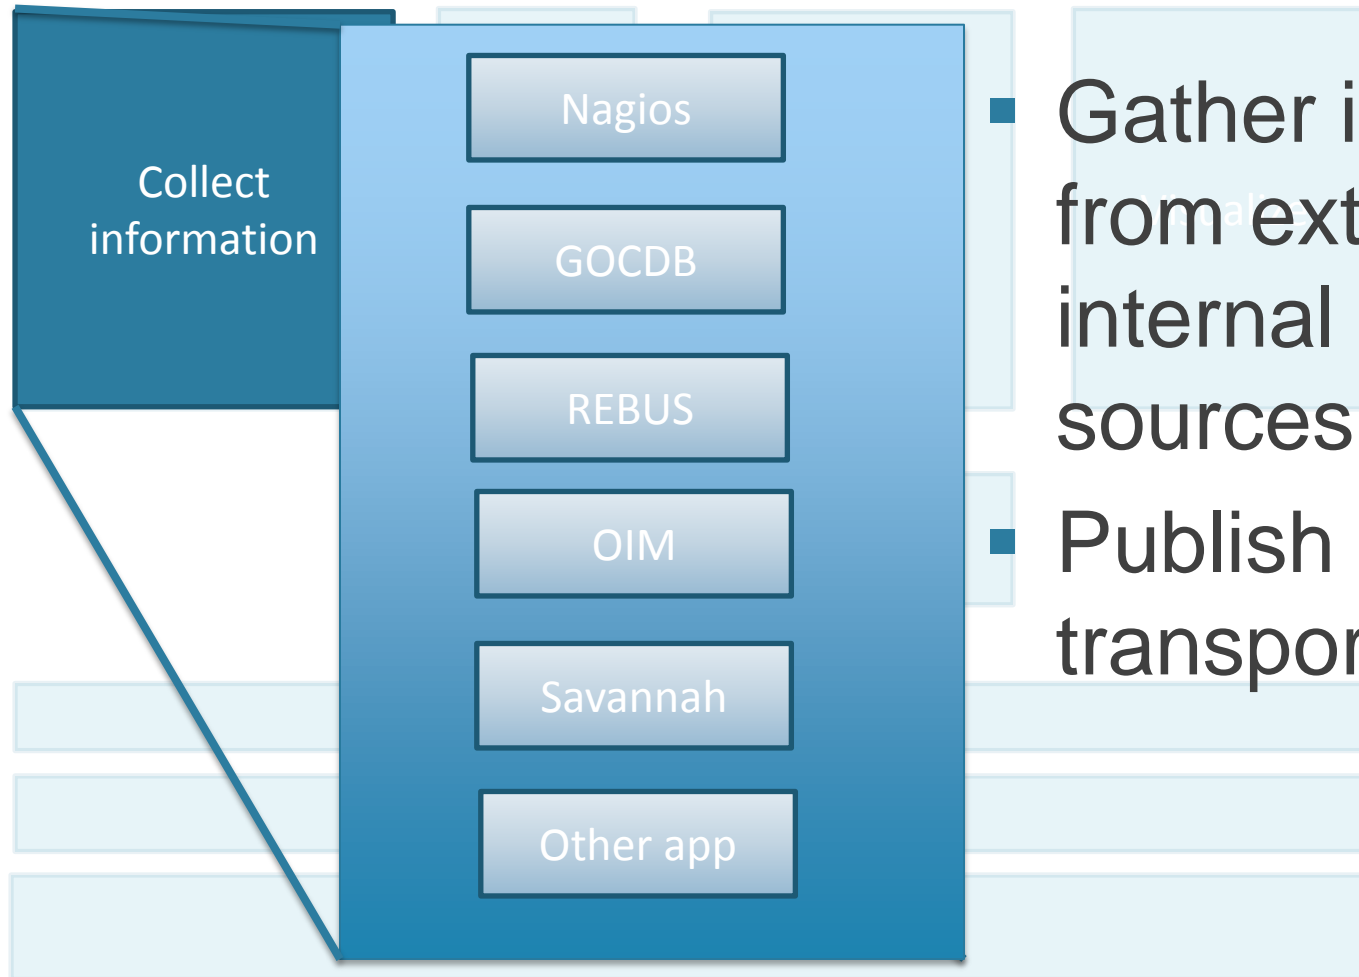


Deployment

- Using **openstack**, puppet, hiera, foreman
 - Quota of 100 nodes, 240 cores
 - Multiple templates already created
 - Development machine (7 nodes)
 - Web servers (SSB, SUM, WLCG transfers, Job: 16 nodes)
 - Elastic Search (6 nodes), Hadoop (4 nodes)
 - Currently working on **nagios** installation
- Migrating machines from quattor to AI
- **Koji** and **Bamboo** for build system and continuous integration

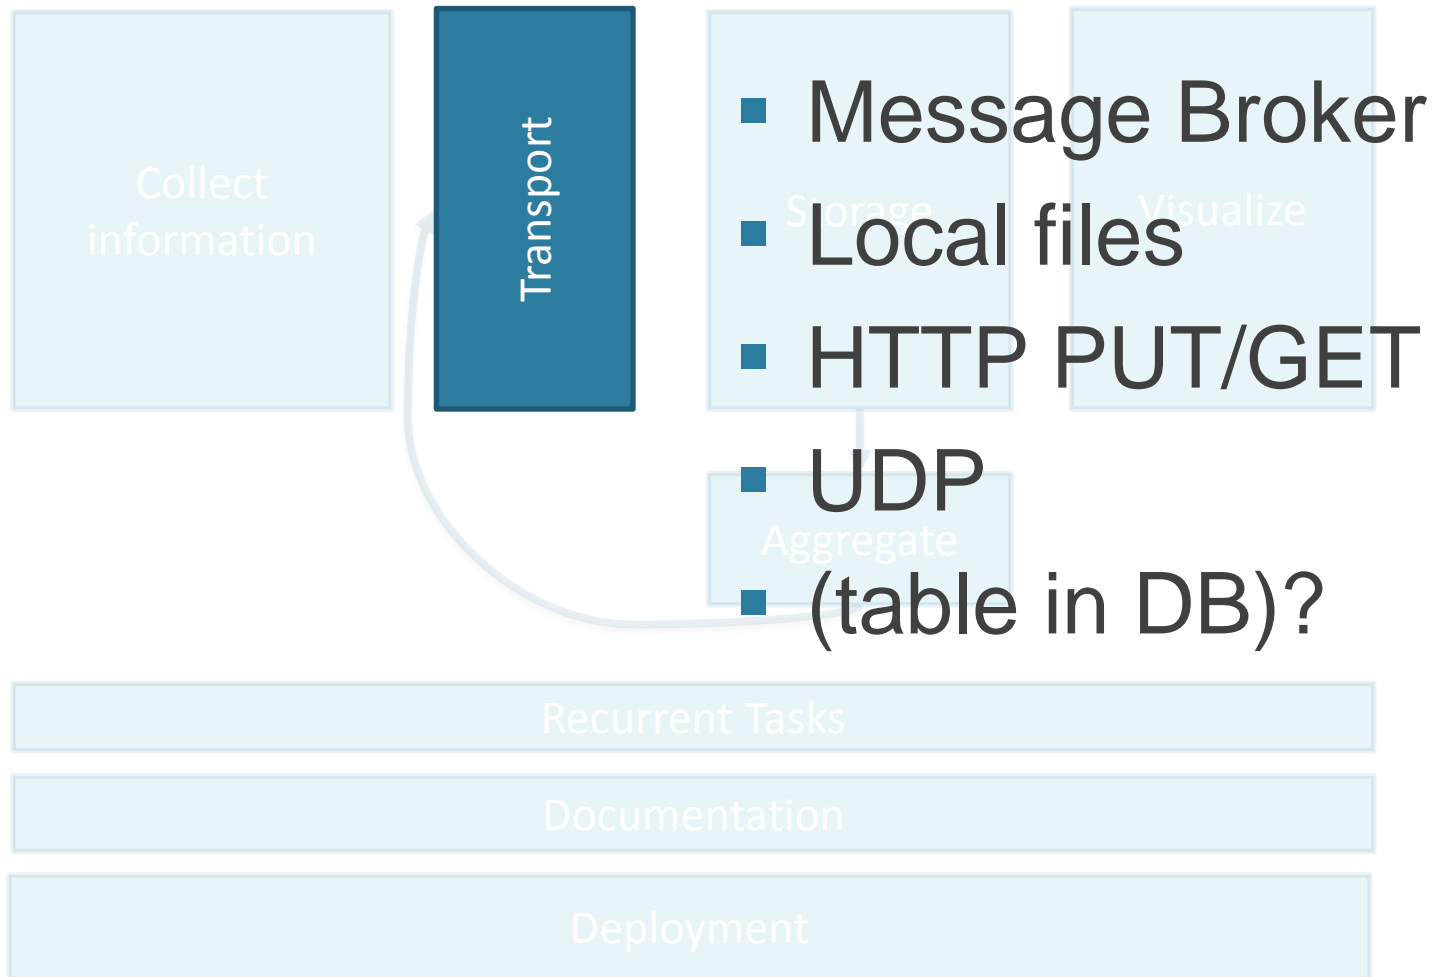


Source of information



- Gather info from external, internal sources.
- Publish it in the transport layer

Transport



Storage

- **Accepts any data**

- #jobs, status of a service, downtime, pledges, channel status
- Metric, Instance, Time Range, Value

- **Archival**

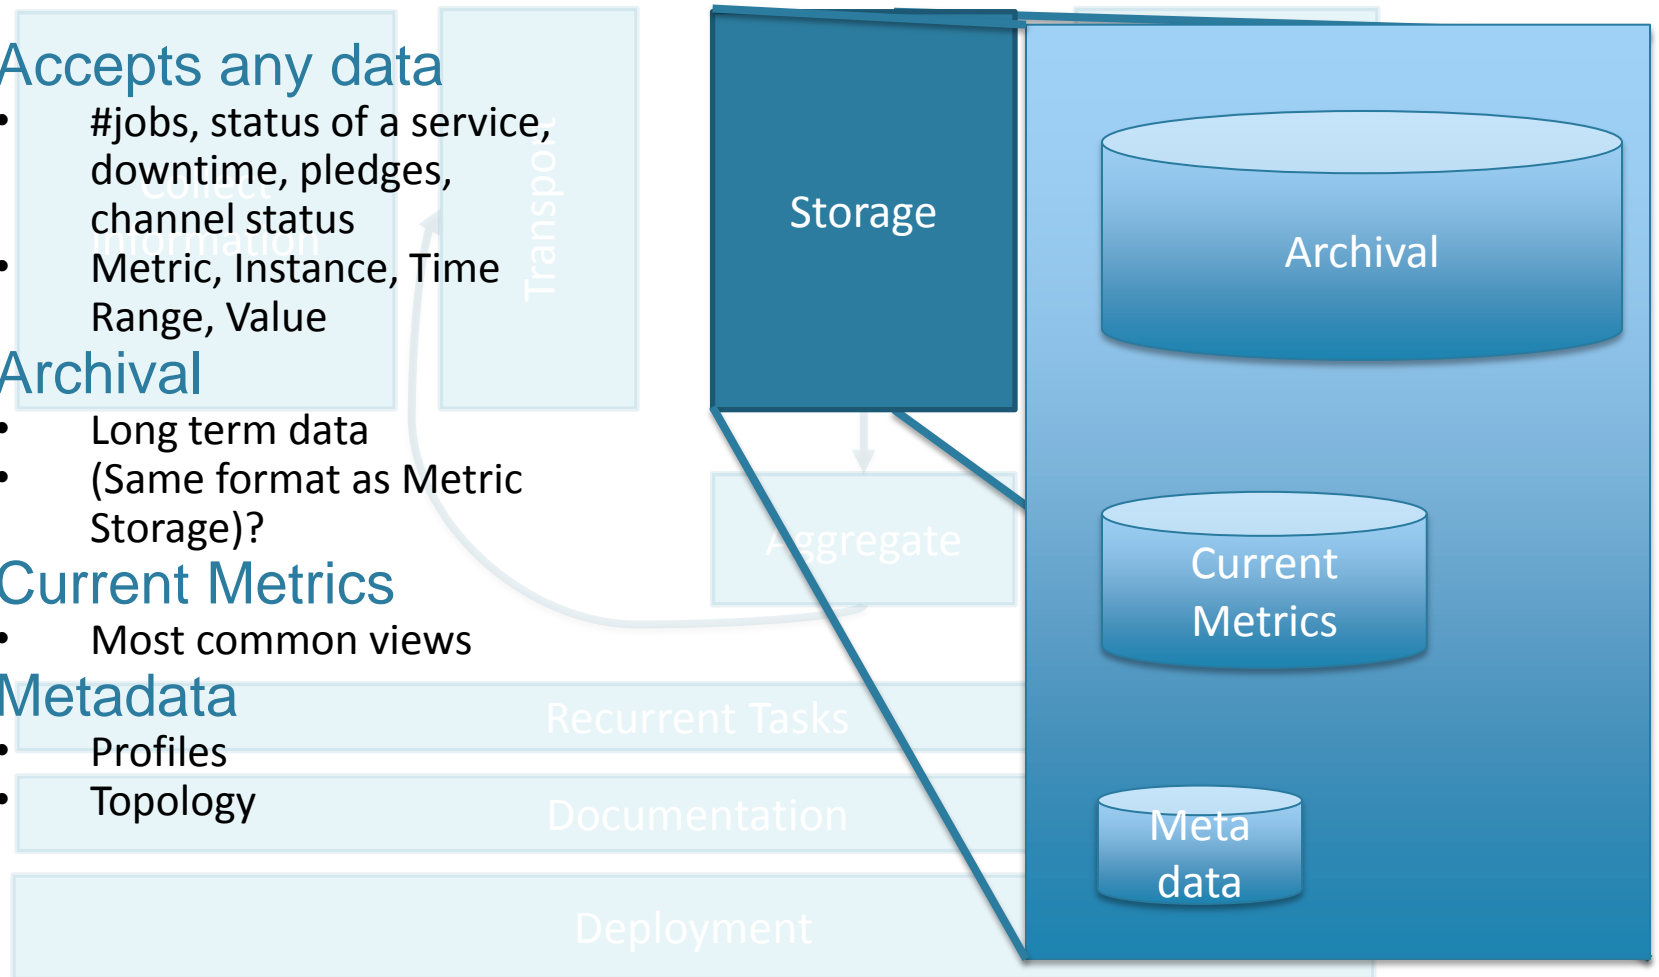
- Long term data
- (Same format as Metric Storage)?

- **Current Metrics**

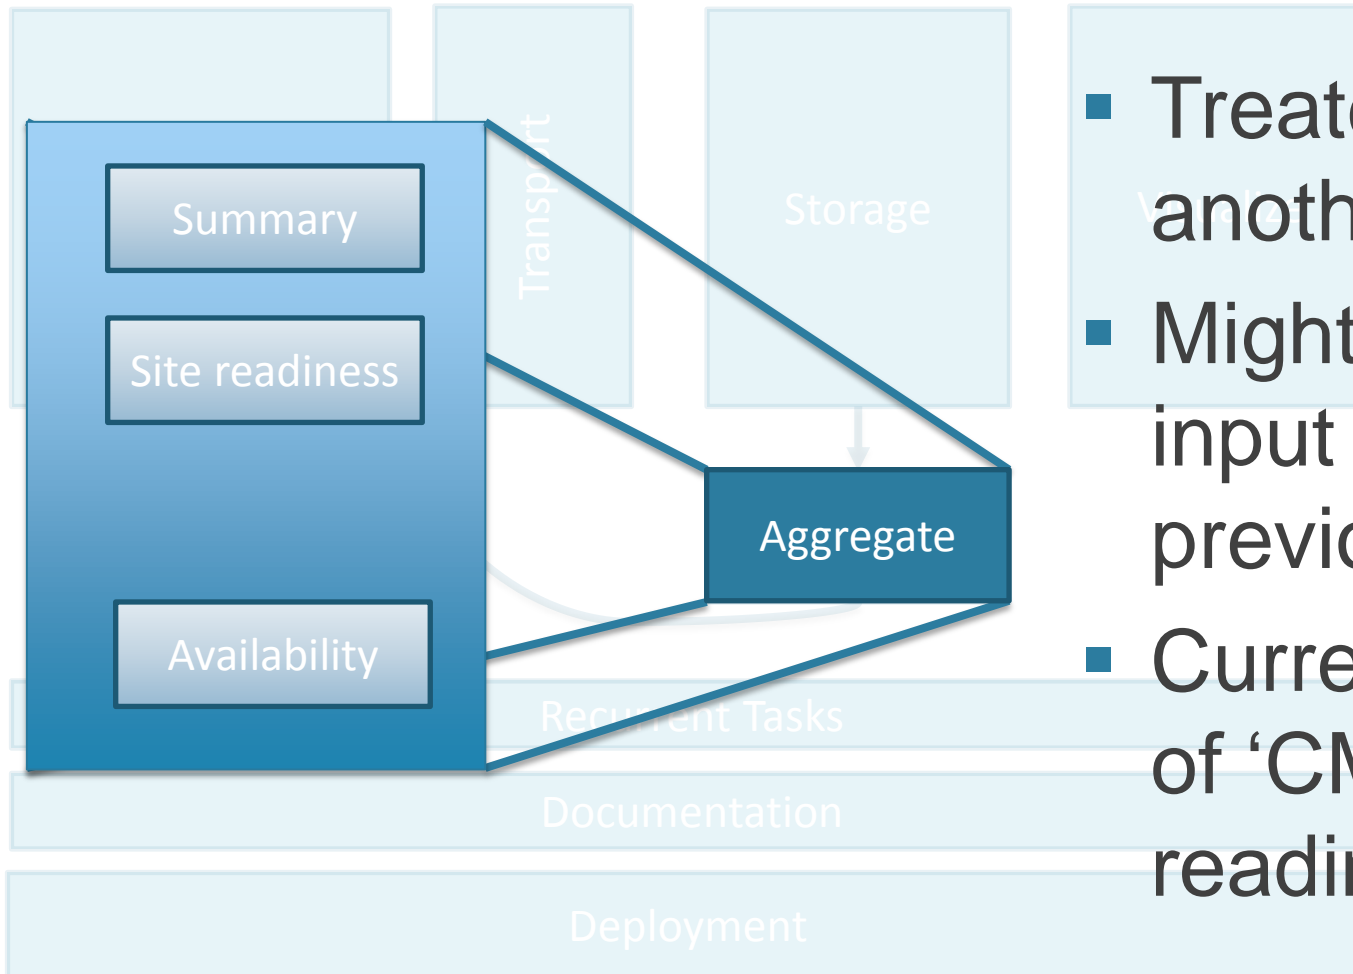
- Most common views

- **Metadata**

- Profiles
- Topology



Aggregation



- Treated as another metric
- Might collect input from previous metrics
- Current schema of 'CMS Site readiness'

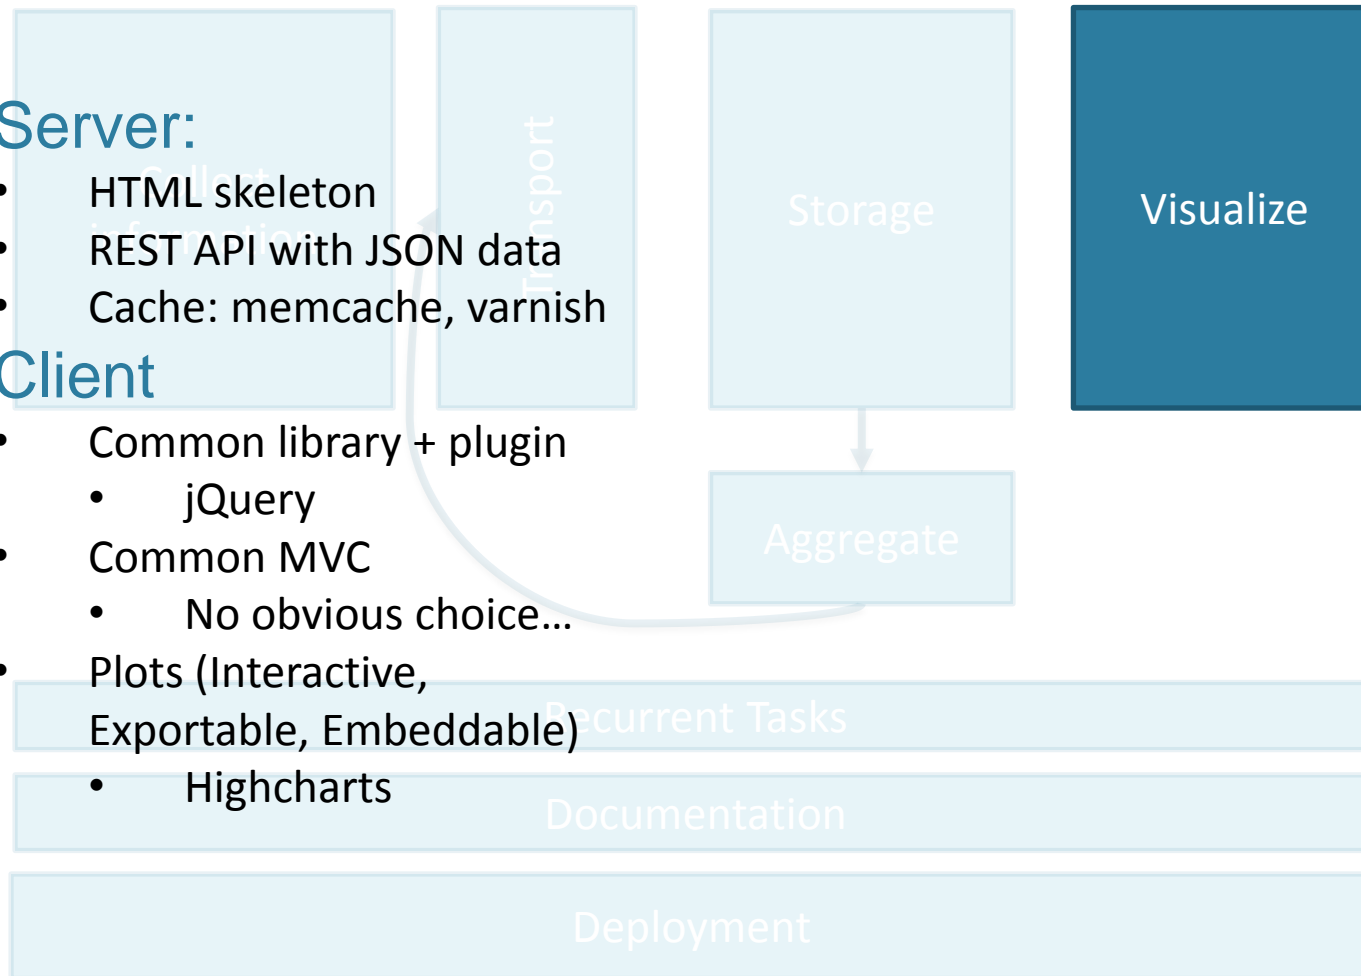
Visualization

- **Server:**
 - HTML skeleton
 - REST API with JSON data
 - Cache: memcache, varnish

- **Client**

- Common library + plugin
 - jQuery
- Common MVC
 - No obvious choice...
- Plots (Interactive, Exportable, Embeddable)

- Highcharts

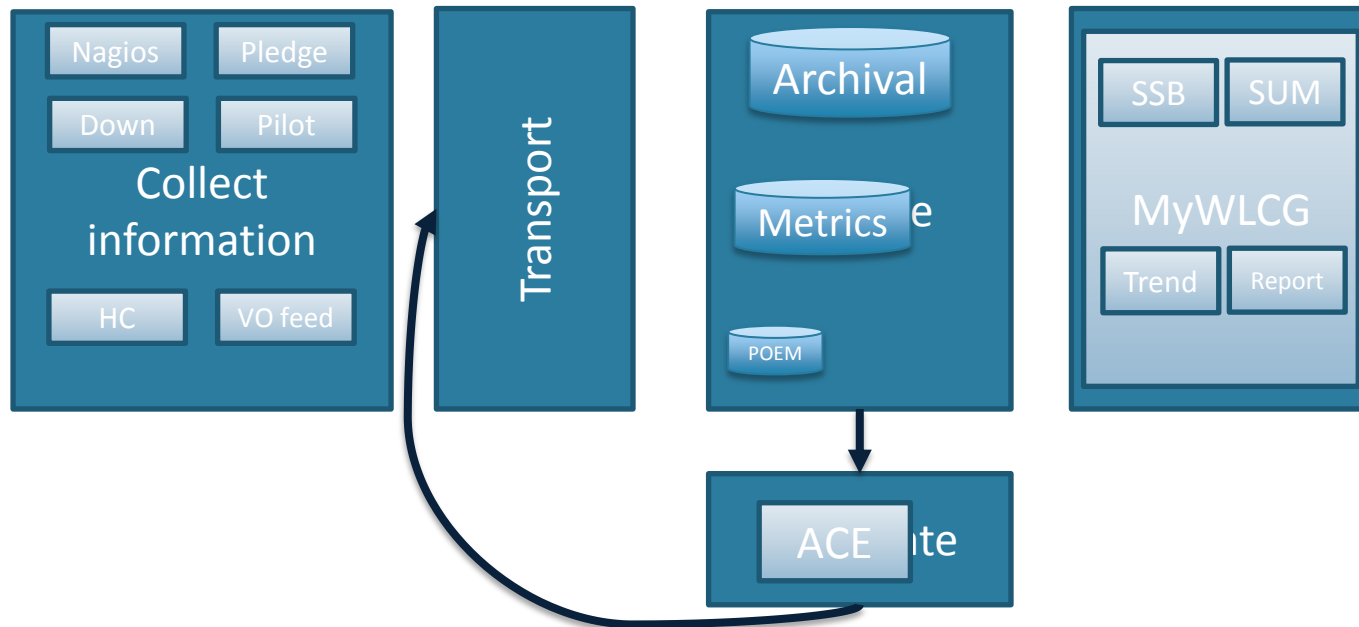


III. Infrastructure monitoring

Current situation

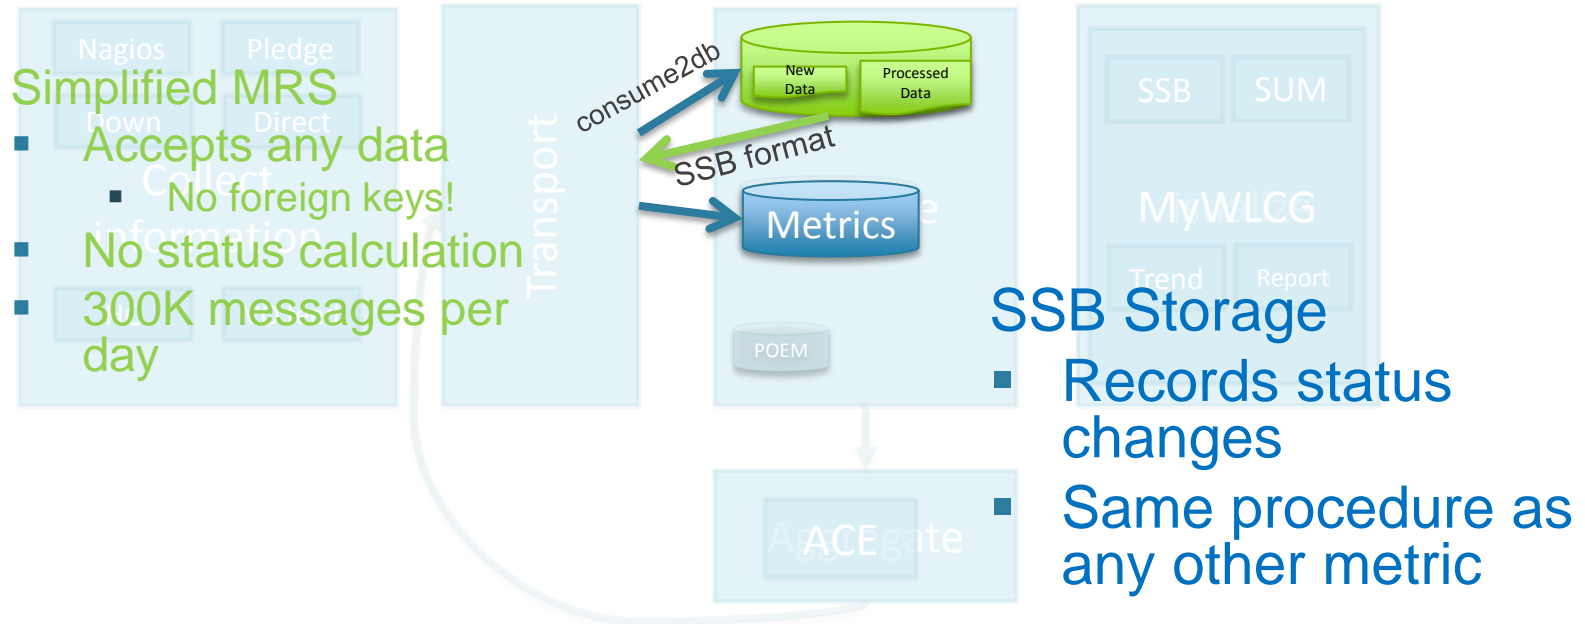
- Big system, difficult to maintain/evolve
- **Many internal dependencies**
- **Multiple schemas, aggregations:**
 - SSB, MRS, ACE
- Scope much bigger than what we need
 - Limit to WLCG
 - Usage of probes
- **Does not test what the experiments are doing!**
- Non-trivial deployment of new tests
- Based on technologies available at the time of the design
- New requests from experiments:
 - Test whatever they want
 - **Availability vs Usability**
- Combine Dashboard/SAM apps

Infrastructure monitoring



- Recurrent Tasks
- Documentation
- Deployment

And for the prototype...



All the data in storage have the same format:

- Instance, Metric, Time range, Value
- Source could be nagios, pilot framework, VO-defined metrics, availabilities

And now we can see metrics...


Show 200 entries Copy Print Save views CMS Search...

Site Name	org.cms.SRM-AIICMS (_cms_Role_production)	org.cms.SRM-GetPFMFromTFC (_cms_Role_production)	org.cms.SRM-VODEl (_cms_Role_production)	org.cms.SRM-VOGet (_cms_Role_production)	org.cms.SRM-VOGetURLs (_cms_Role_production)	org.cms.SRM-VOLs (_cms_Role_production)	org.cms.SRM-VOLsDir (_cms_Role_production)	org.cms.SRM-VOPut (_cms_Role_production)	org.cms.WN-analysis (_cms_Role_lcgadmin)	org.cms.WN-basic (_cms_Role_lcgadmin)	org.cms.WN-es (_cms_Role_lcgac)
OSG-CE red-gw 2.unl.edu	n/a	n/a	n/a	n/a	n/a	n/a	n/a	n/a	OK	OK	OK
OSG-CE red.unl.edu	n/a	n/a	n/a	n/a	n/a	n/a	n/a	n/a	OK	OK	OK
OSG-CE rossmann-osg.rcac.purdue.edu	n/a	n/a	n/a	n/a	n/a	n/a	n/a	n/a	OK	OK	OK
OSG-CE t3serv007.mit.edu	n/a	n/a	n/a	n/a	n/a	n/a	n/a	n/a	OK	OK	OK
OSG-CE top.ucr.edu	n/a	n/a	n/a	n/a	n/a	n/a	n/a	n/a	OK	OK	OK
OSG-CE umiss001.hep.olemiss.edu	n/a	n/a	n/a	n/a	n/a	n/a	n/a	n/a	CRITICAL	OK	OK
OSG-CE uscms1.ftech-grid3.ft.edu	n/a	n/a	n/a	n/a	n/a	n/a	n/a	n/a	n/a	n/a	n/a
OSG-SRMv2 bonner-grid.rice.edu	CRITICAL	OK	WARNING	WARNING	WARNING	WARNING	CRITICAL	CRITICAL	n/a	n/a	n/a
OSG-SRMv2 barm-1.i2.ucsd.edu	UNKNOWN	UNKNOWN	WARNING	WARNING	WARNING	WARNING	WARNING	WARNING	n/a	n/a	n/a
OSG-SRMv2 barm-3.i2.ucsd.edu	OK	OK	OK	OK	OK	OK	OK	OK	n/a	n/a	n/a
OSG-SRMv2 charm.ucr.edu	OK	OK	OK	OK	OK	OK	OK	OK	n/a	n/a	n/a
OSG-SRMv2 cil-se.ultraight.org	OK	OK	OK	OK	OK	OK	OK	OK	n/a	n/a	n/a
OSG-SRMv2 cms-se.hep.uprm.edu	OK	OK	OK	OK	OK	OK	OK	OK	n/a	n/a	n/a
OSG-SRMv2	OK	OK	OK	OK	OK	OK	OK	OK	n/a	n/a	n/a

Showing 201 to 310 of 310 entries DB query took 0.4681 s First Previous 1 2 Next Last

Status of org.cms.SRM-VODEl (_cms_Role_production)

24 Hours from 2013-08-13 06:31 to 2013-08-14 06:31



All_sites

Status	Site Name	Status	Site Name	Status	Site Name	Status	Site Name
●	ARC-CE arc-ce01.gridpp.rl.ac.uk	✓	CREAM-CE cream-ge-1-kit.gridka.de	✓	CREAM-CE lgoc2.shef.ac.uk	✓	SRMv2 dcache-se-atlas.desy.de
●	ARC-CE arc-ce02.gridpp.rl.ac.uk	✓	CREAM-CE cream-ge-2-kit.gridka.de	✓	CREAM-CE lgoc21.jinr.ru	✓	SRMv2 dcache-du.cesnet.cz
✓	CREAM-CE abaddon.hec.lancs.ac.uk	✓	CREAM-CE cream-ge-3-kit.gridka.de	✓	CREAM-CE lgoc3.shef.ac.uk	✓	SRMv2 dgridstrm-fzk.gridka.de
✓	CREAM-CE agcream1.atlas.unimeb.edu.au	✓	CREAM-CE cream-ge-4-kit.gridka.de	✓	CREAM-CE lgoc6a.dnp.fmph.uniba.sk	✓	SRMv2 dpm-cyf-kr.edu.pl
✓	CREAM-CE alice23.spbu.ru	✓	CREAM-CE cream-ge-5-kit.gridka.de	✓	CREAM-CE lgocreamce.dnp.fmph.uniba.sk	✓	SRMv2 dpm.hep.unibe.ch
●	CREAM-CE atlas-ce-01.roma1.infn.it	✓	CREAM-CE cream-ge-6-kit.gridka.de	✓	CREAM-CE lnucs-ce-01.cs.infn.it	✓	SRMv2 epgs1.ph.bham.ac.uk
✓	CREAM-CE atlas-ce-02.roma1.infn.it	✓	CREAM-CE cream-ge-7-kit.gridka.de	✓	CREAM-CE lnhe-cream.in2p3.fr	✓	SRMv2 f-dp001.grid.sinica.edu.tw
✓	CREAM-CE atlas-cream01.na.infn.it	✓	CREAM-CE cream-ge-8-kit.gridka.de	✓	CREAM-CE lpsc-ce.in2p3.fr	✓	SRMv2 fal-pygrid-30.lincs.ac.uk
✓	CREAM-CE atlas-cream01.roma1.infn.it	●	CREAM-CE cream.grid.cyf-kr.edu.pl	✓	CREAM-CE lpsc-cream-ce.in2p3.fr	✓	SRMv2 gfe02.grid.hep.ph.ic.ac.uk
✓	CREAM-CE atlas-cream02.roma1.infn.it	✓	CREAM-CE cream01-tic.ciemat.es	✓	CREAM-CE lpygrid01.in2p3.fr	✓	SRMv2 golias100.farm.particle.cz
✓	CREAM-CE atlasce02.scope.unina.it	✓	CREAM-CE cream01.grid.auth.gr	✓	CREAM-CE maigrd-ce1.physik.uni-mainz.de	✓	SRMv2 grid-cert-03.roma1.infn.it

Aggregation

- Combination of ACE +SSB Virtual Columns
- Two types:
 - **Horizontal**: $Ins_1 (M_1 \dots M_n) \rightarrow Ins_1 (M_p)$
 - **Vertical**: $M_1 (Ins_1 \dots Ins_n) \rightarrow Ins_p (M_2)$
- Initial options for “and”, “or” of current status
 - Later on, might be extended to ‘sliding window’
- Full description

Examples of aggregation

Show 200 entries Copy Print Save view: ATLAS CRITICAL

Site Name	SAM					
	org.atlas.SRM.VODef (_atlas_Role_production)	org.atlas.SRM.VOGet (_atlas_Role_production)	org.atlas.SRM.VOPut (_atlas_Role_production)	org.atlas.WN.swspace (_atlas_Role_lcadmin)	org.atlas.WN.swspace (_atlas_Role_pilot)	org.atlas.WN.ontag (_atlas_Role_lcadmin)
ARC-CE arc-ce01.gridpp.rl.ac.uk	n/a	n/a	n/a	OK	CRITICAL	CRITICAL
ARC-CE arc-ce02.gridpp.rl.ac.uk	n/a	n/a	n/a	CRITICAL	CRITICAL	CRITICAL
CREAM-CE abaxidon.hec.lancs.ac.uk	n/a	n/a	n/a	OK	OK	OK
CREAM-CE agcream1.atlas.unimelb.edu.au	n/a	n/a	n/a	OK	OK	OK
CREAM-CE atlas-ce-01.roma1.infn.it	n/a	n/a	n/a	CRITICAL	CRITICAL	WARNING
CREAM-CE atlas-ce-02.roma1.infn.it	n/a	n/a	n/a	OK	OK	OK
CREAM-CE atlas-cream01.na.infn.it	n/a	n/a	n/a	OK	OK	OK
CREAM-CE atlas-creamce-01.roma1.infn.it	n/a	n/a	n/a	OK	OK	OK
CREAM-CE atlas-creamce-02.roma1.infn.it	n/a	n/a	n/a	OK	OK	OK
CREAM-CE atlasce02.scope.unina.it	n/a	n/a	n/a	OK	n/a	OK
CREAM-CE atlasce1.infn.it	n/a	n/a	n/a	OK	OK	OK
CREAM-CE atlasce2.infn.it	n/a	n/a	n/a	OK	OK	OK
CREAM-CE atlasce3.infn.it	n/a	n/a	n/a	OK	OK	OK
CREAM-CE basf01.nipne.ro	n/a	n/a	n/a	OK	OK	OK
CREAM-CE bugaboo-hep.westgrid.ca	n/a	n/a	n/a	OK	OK	OK
CREAM-CE calc1.t1.grid.kiae.ru	n/a	n/a	n/a	OK	OK	OK

Showing 1 to 200 of 341 entries - DB query took 0.1018 s

Show 200 entries Copy Print Save view: ATLAS CRITICAL

Site Info		
Site Name	Tier	Cloud
INFN-FRASCATI	T2	IT

Showing 1 to 1 of 1 entries - DB query took 0.0186 s

ATLAS CRITICAL
 WN Site
 (expand this column)

Summary

- Lots of progress towards unified schema
- Data can be published from different sources
 - Nagios, VO-defined metrics, ACE, (HC, Job Pilots)
- Single schema for storage
- Components talk to each other through API
- Getting close to a “proof of concept”
 - Aggregation needs some work
 - Visualization might need adjusting
- Other tasks can go in parallel
 - NoSQL evaluation
 - Nagios configuration
 - Only active metrics