Experience in CMS with Analytic Tools for Data Flow and HLT Monitoring

Dainius Šimelevičius – Vilnius University (LT)
on behalf of CMS DAQ Group

ALICE, ATLAS, CMS & LHCb Second Joint Workshop on DAQ@LHC
12-14 April 2016, Château de Bossey, Switzerland

Main Authors

- *Emilio Meschi*
- *Srećko Morović*
- *Luciano Orsini*
- *Dainius Šimelevičius*

Introduction

Data Flow and HLT Monitoring

- Monitoring characteristics
  - DAQ and HLT are independent from monitoring
  - Monitoring runs in parallel with the DAQ and HLT processes
  - Shares resources with DAQ and HTL processes
- Monitoring functions
  - Collection of data from multiple sources
  - Transformation/Reduction (e.g., sum, average)
  - Presentation (e.g., raw data, HTML or JSON formatted data)
  - Storage
- Users
  - Shifters
  - DAQ on-call experts
  - Control System (e.g., perform automatic actions)
  - Diagnostic System (e.g., expert systems)

The Concept

- This is a follow-up work based on the presentation made in DAQ@LHC 2013 by Luciano Orsini "Dataflow Monitoring: Future Perspectives"
  - http://indico.cern.ch/event/217480/session/1/contribution/35/material/slides/0.pdf
  - The presentation highlighted the approach chosen to support CMS DAQ monitoring based on SaaS (Software as a Service)
    - Current architecture is service oriented (XaaS)
  - Indicated the cloud architecture well-suited for running the monitoring services
- Subsequent natural step would be to use off the shelf software as opposed to custom development
  - Reducing maintenance and support effort

Where We're Going

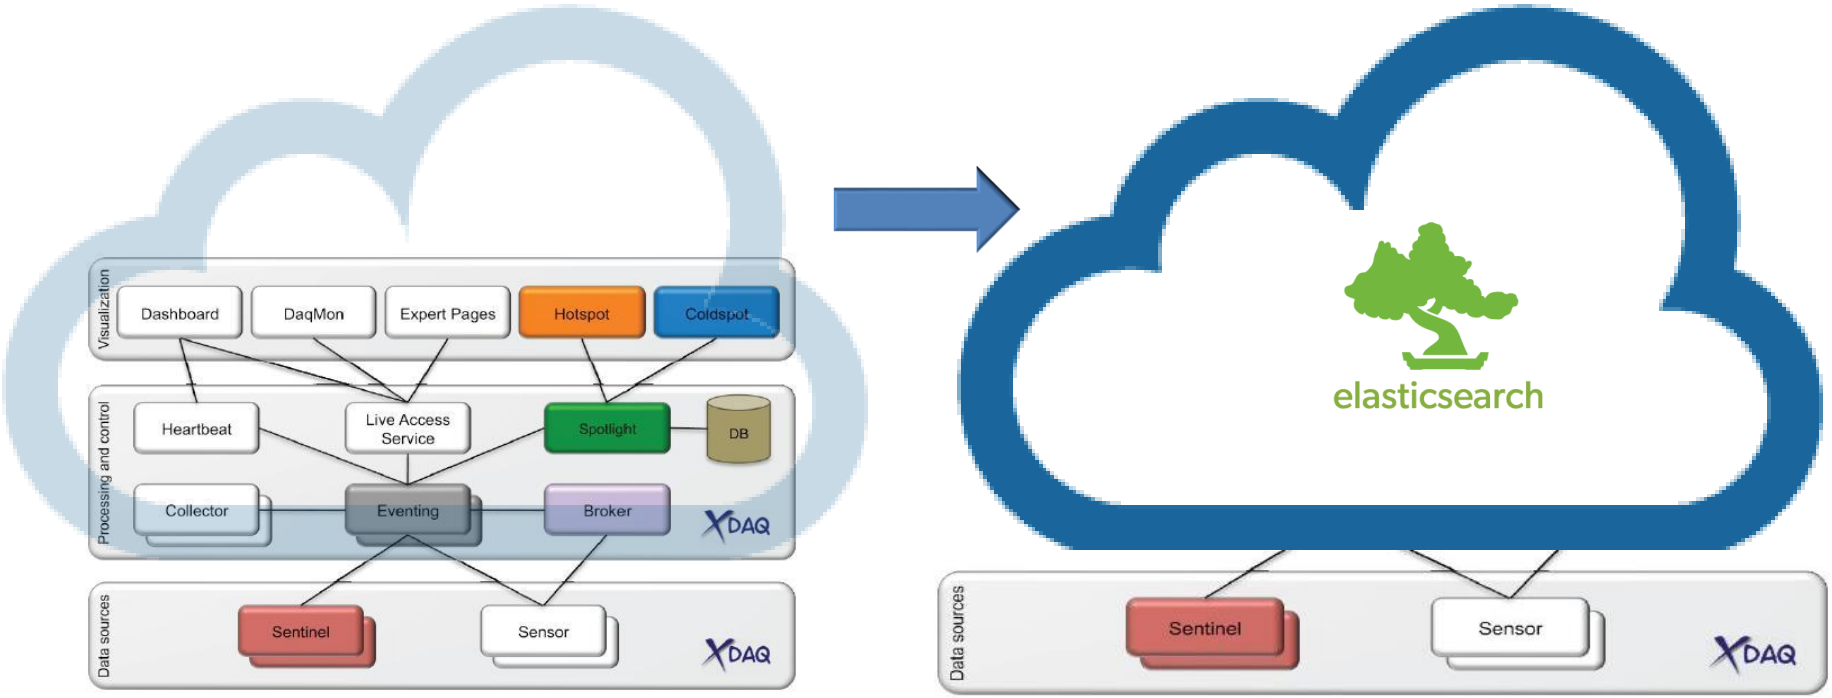

- Custom components
- Custom protocols

- Off The Shelf component
- Standard protocols

Proposed Approach

- Monitoring as an IT service
  - Cloud Administrator (sysadmin) – manage resources, provide the cloud
  - Cloud Facilitator (monitoring expert) – developers, sysadmins
    - Requires certain expertise
  - Consumers/Users (sub-detectors and/or central DAQ) – using the services in the cloud
- Cloud service could be private, IT or commercial
  - Currently P5 resources
  - Tested Openstack in IT
- Fully decoupled cloud services and DAQ resources
- No plugins or specialized programming but configuration only

Elasticsearch

- Elasticsearch chosen as a good candidate
  - Rank 11 in http://db-engines.com/en/ranking
- NoSQL database
- Search server based on Apache Lucene
- RESTful API with JSON over HTTP
- Scalable architecture
- Document search throughout the cluster with a single query

Elasticsearch

Elasticsearch Terminology

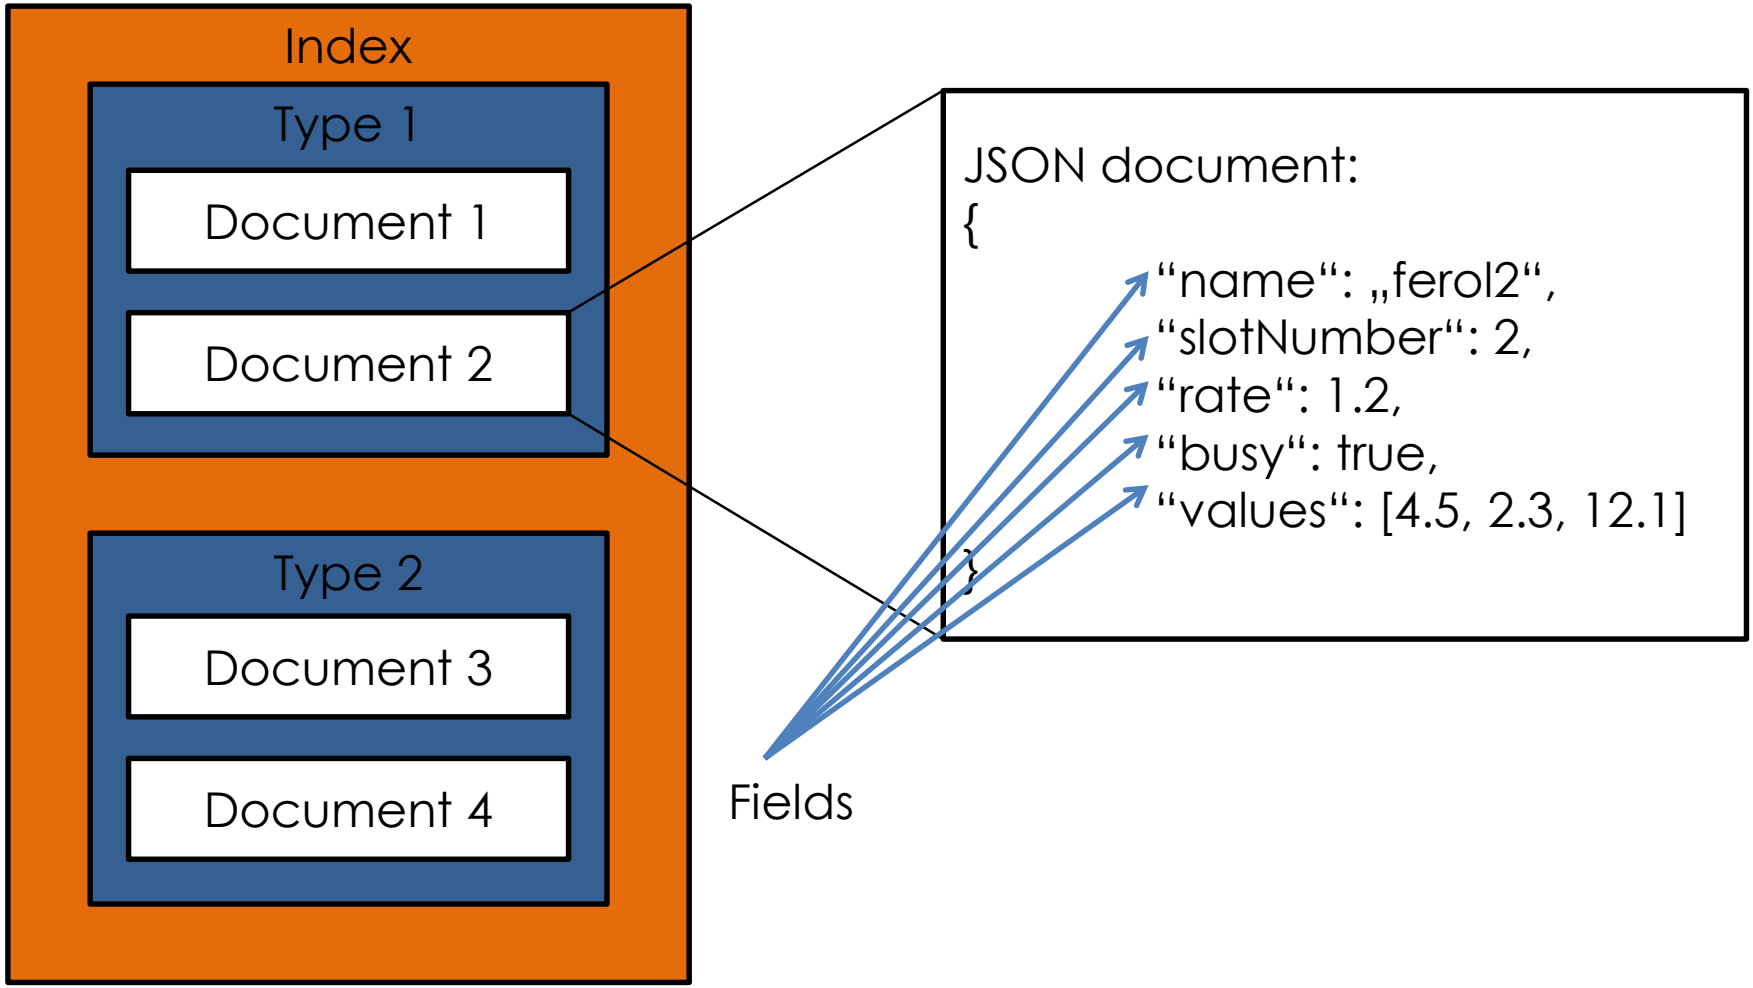

RESTful API with JSON over HTTP

- CRUD (create-retrieve-update-delete) part of API allows using Elasticsearch as a NoSQL database
  - curl -XPUT 'http://[hostname]:9200/flashlists/hostInfo/1'
    -d '{"context": "http://bu-c2f16-23-01.cms:9999", "cpuUsage": 9.68}'
  - curl -XGET 'http://[hostname]:9200/flashlists/hostInfo/1'
  - curl -XPOST 'http://[hostname]:9200/flashlists/hostInfo/1/_update'
    -d '{"doc": {" cpuUsage": 13.73}}'
  - curl -XDELETE 'http://[hostname]:9200/flashlists/hostInfo/1'
- Management of indices
  - curl -XPUT 'http://[hostname]:9200/shelflist'
- Data searching
  - curl -XGET 'http://[hostname]:9200/flashlists/hostInfo/_search'
    -d '{"query": {"range": {"cpuUsage": {"gt": 10}}}}'

Elasticsearch Scalability

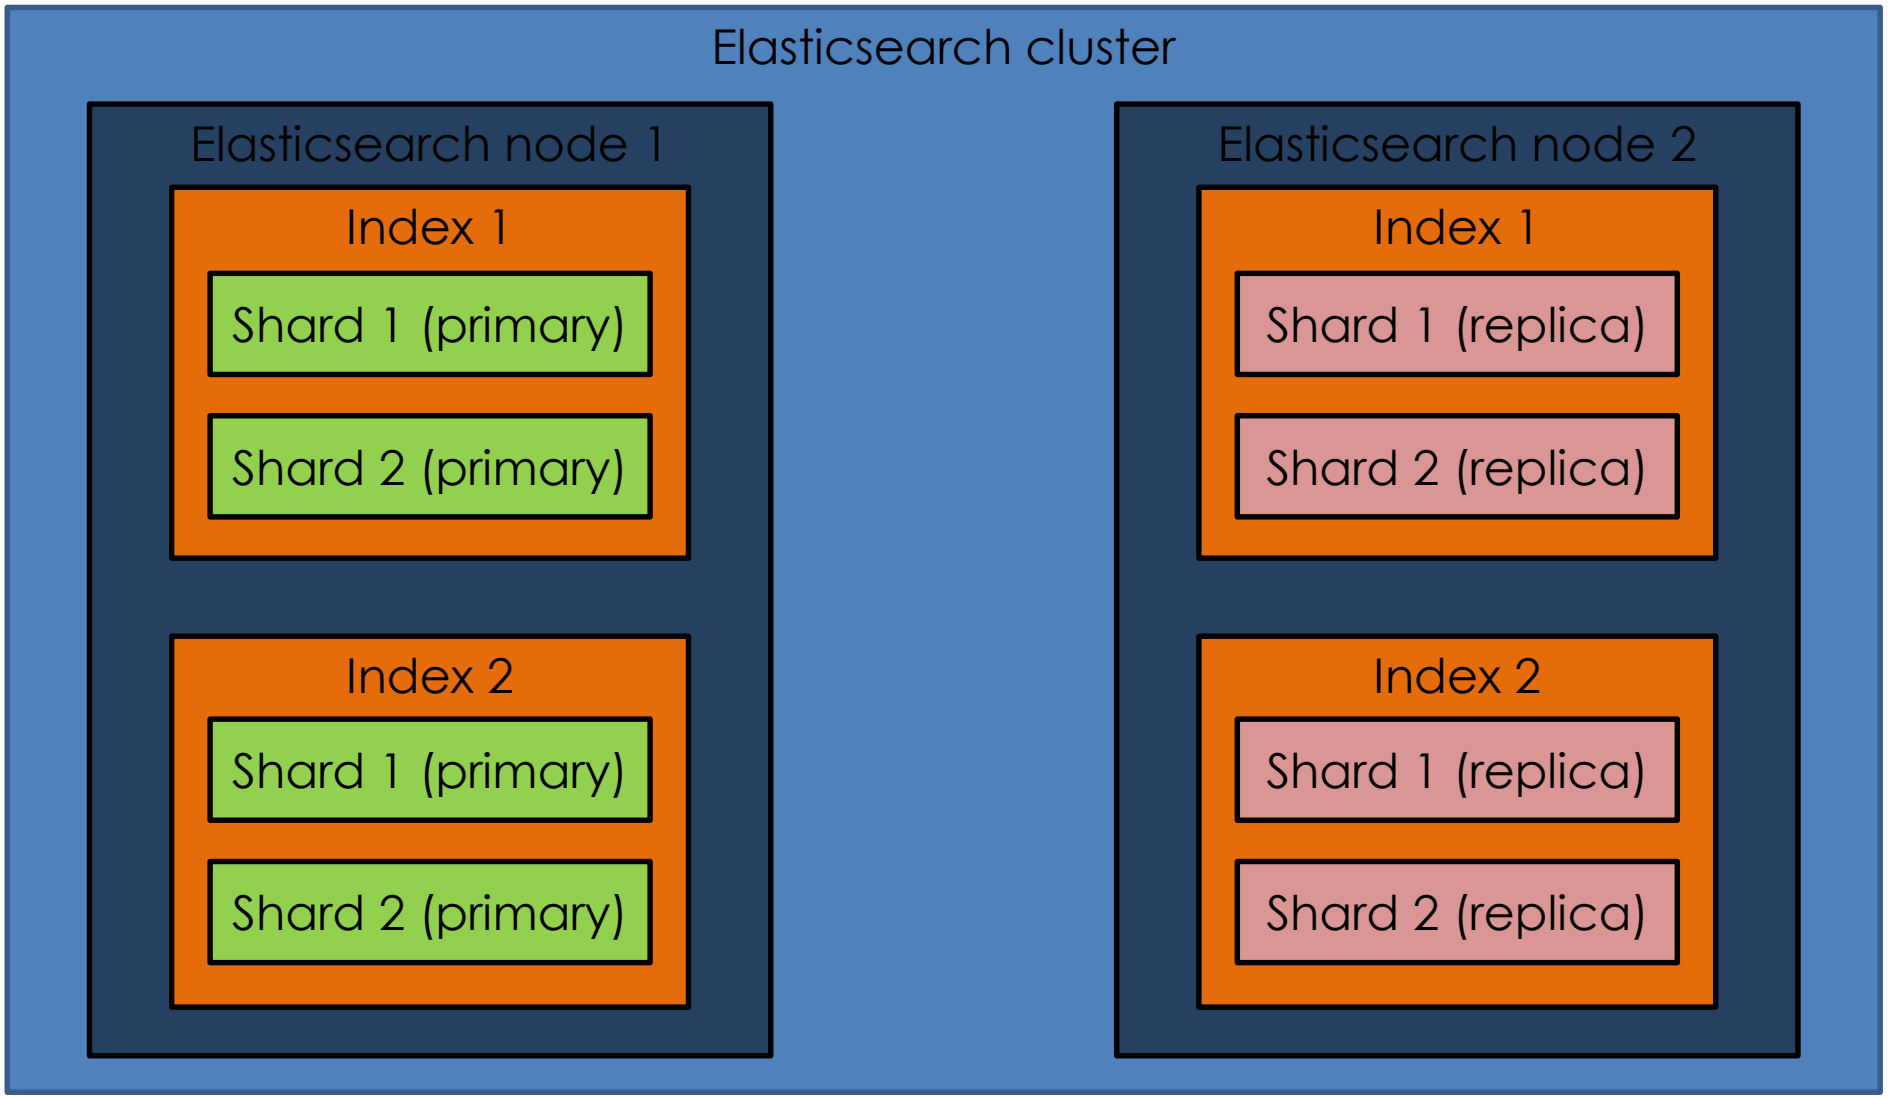

Limitations

- Does not allow fields of different types for the same name within the same index
- Elasticsearch only provides a limited number of data types
- Most integer type map to a single type rather than specialized

Supported	XData	Elasticsearch	Flashlist
*	Boolean	boolean	bool
*	Integer	integer[1]	int
	Integer32	integer	
	Integer64	long	
	Unsigned Integer	long	
*	Unsigned Integer32	long	unsigned int 32
*	Unsigned Integer64	long	unsigned int 64
*	Unsigned long	long[1]	unsigned long
*	Unsigned short	integer	unsigned short
*	String	string	string

Elasticsearch for DAQ and HLT Monitoring

Scenario for DAQ and HLT Monitoring

- Index creation
- Map creation
- Document creation
- Search for the document which was created last
- Documents are removed automatically by Elasticsearch according to TTL (Time To Live) values preset for the document types

Elasticsearch Usage Patterns

	DAQ monitoring	HLT monitoring
Data location	In memory	On disk
Data injector	C++ program	Python script
Source nodes	~200	~1000
Index operation rate	~3.4kHz	~5kHz
Data persistency	4 seconds to 2 minutes; some data stored permanently	Data stored permanently
Indices	2	One index per run + 3
Types	24	8 + 15
Fields	678 (16000 including sub-detectors)	87 + 180

DAQ Monitoring Option 1

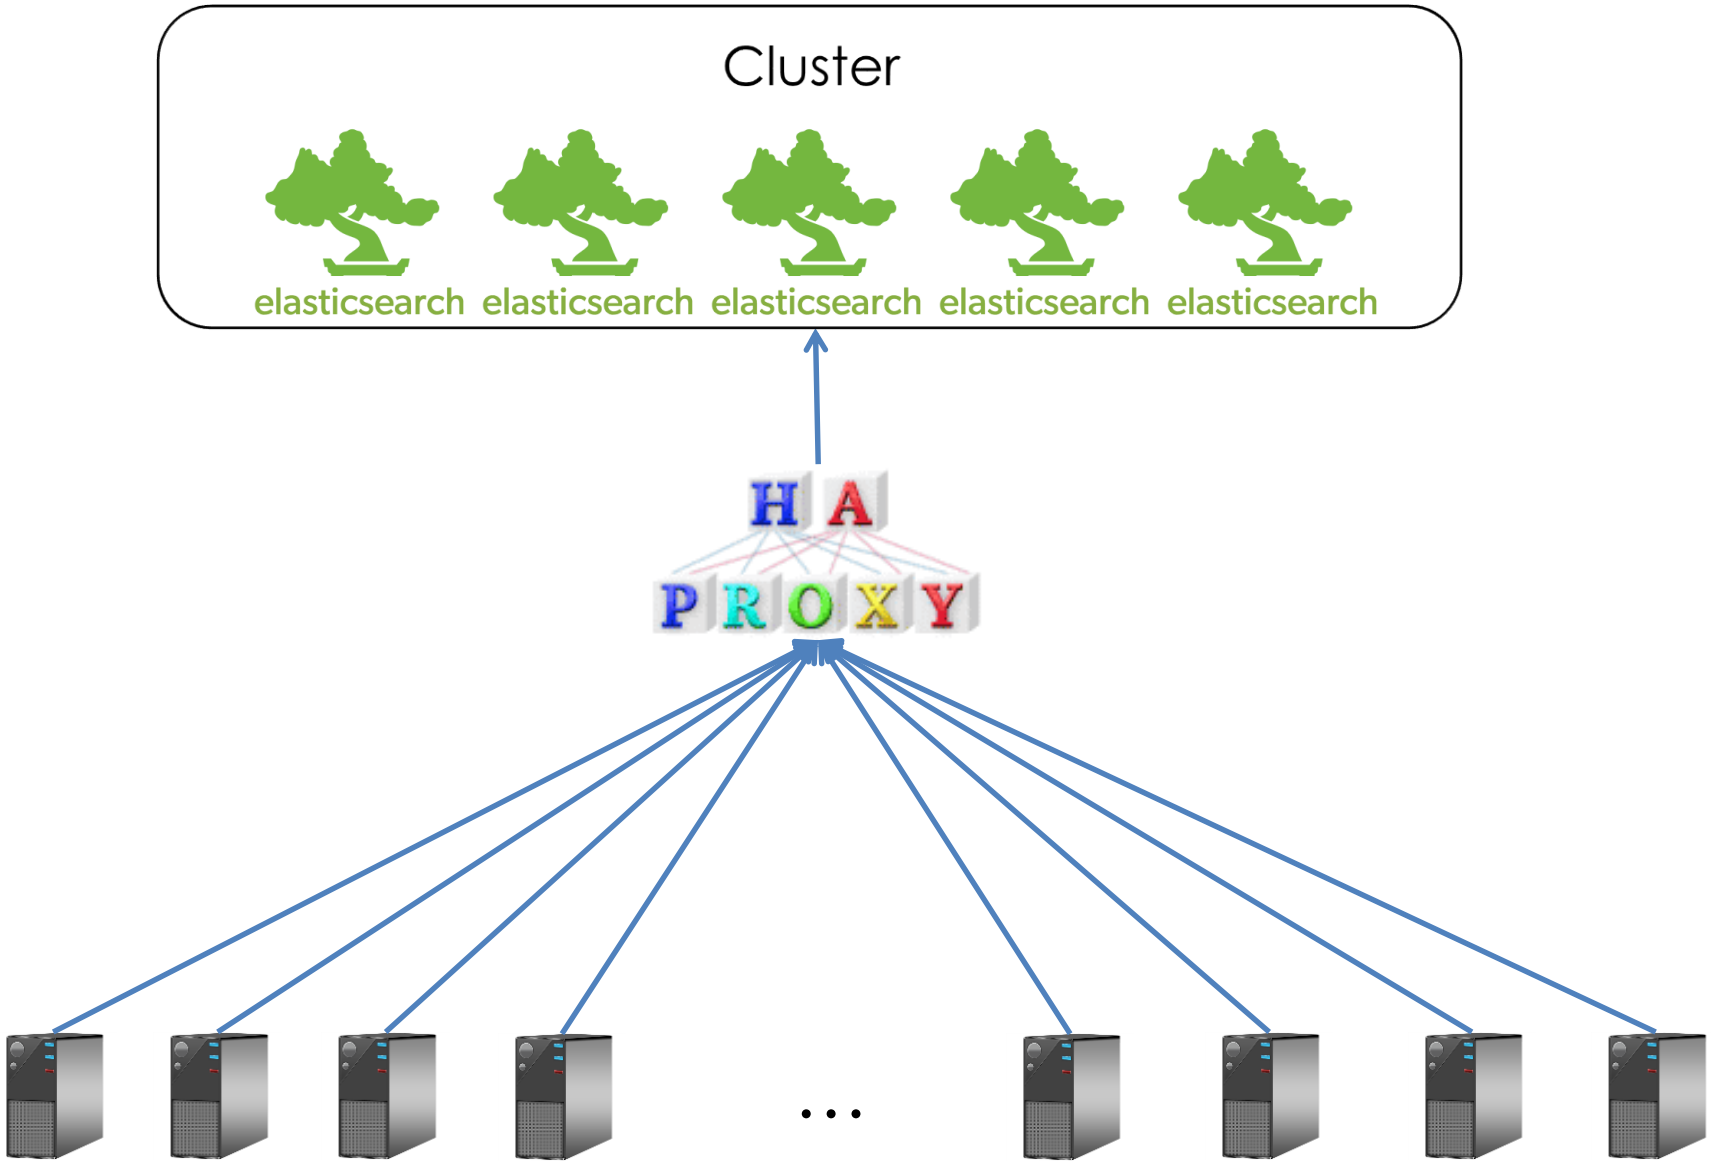

DAQ Monitoring Option 2

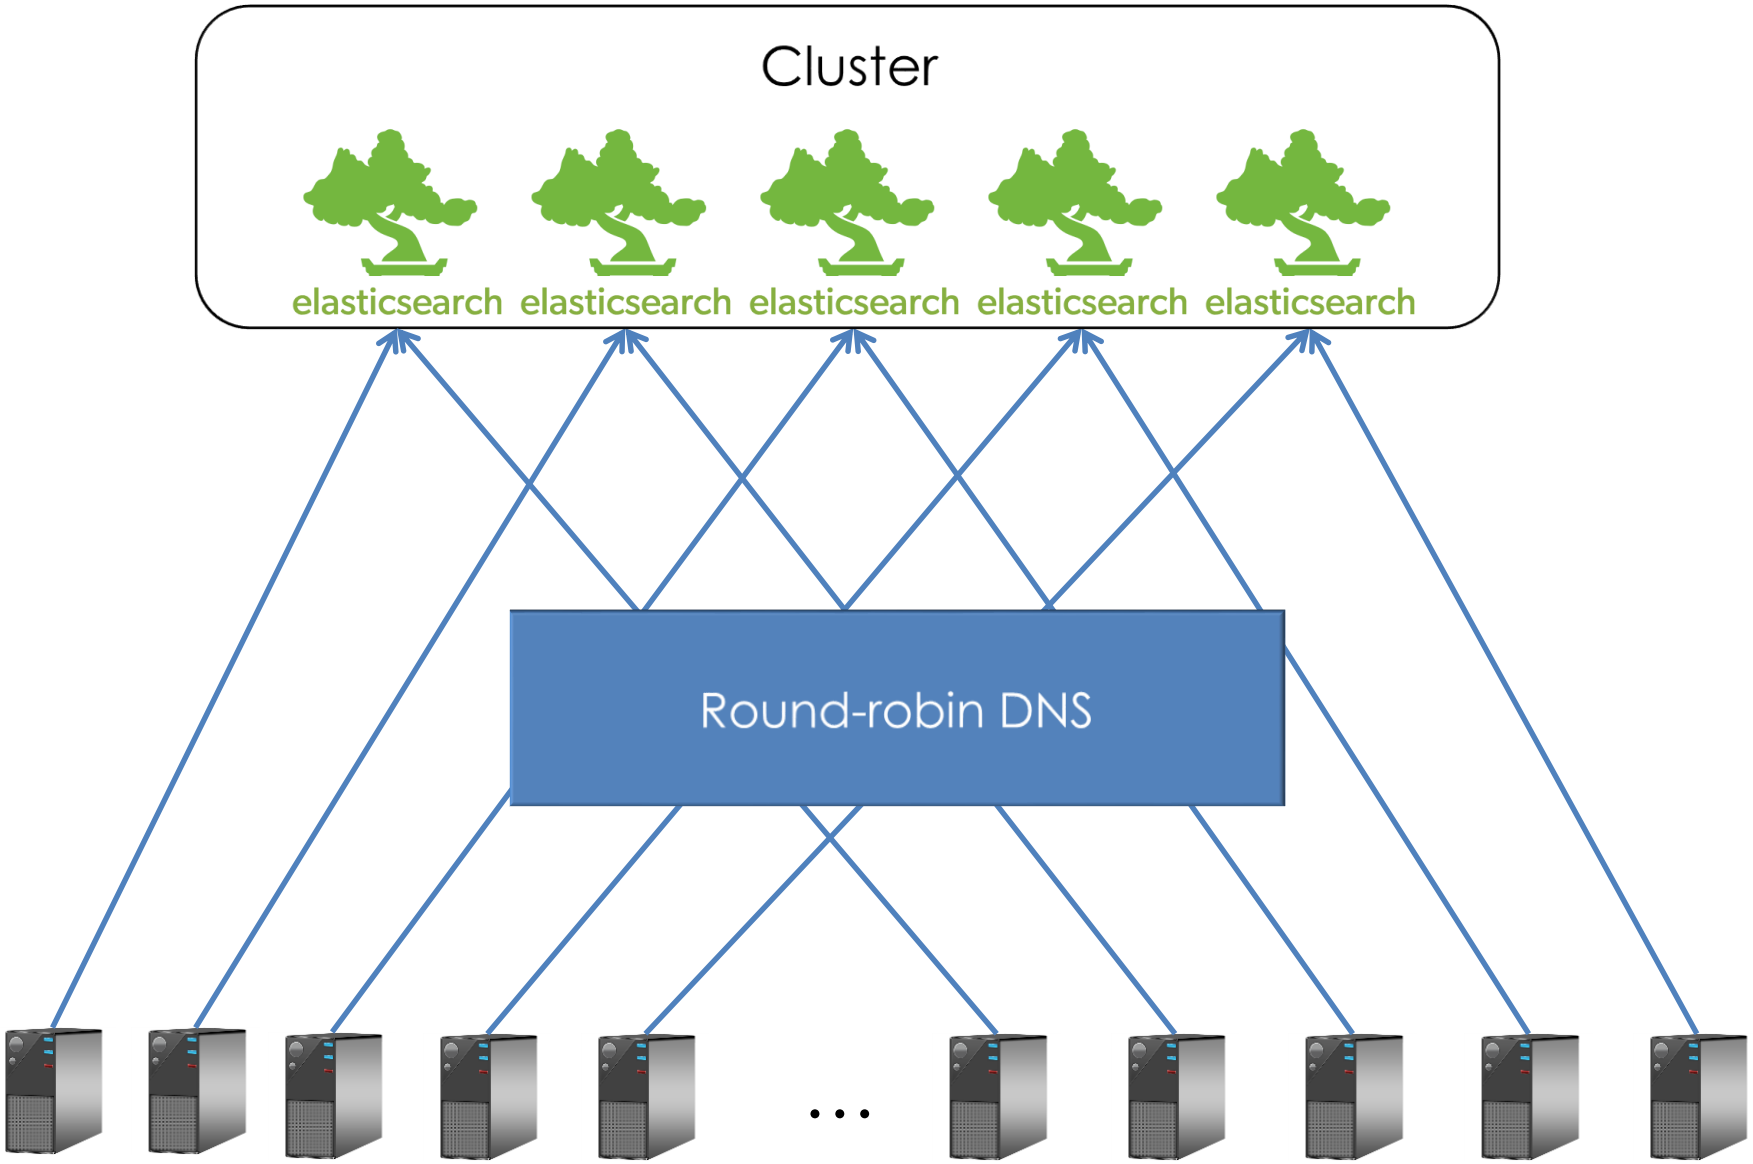

HLT monitoring

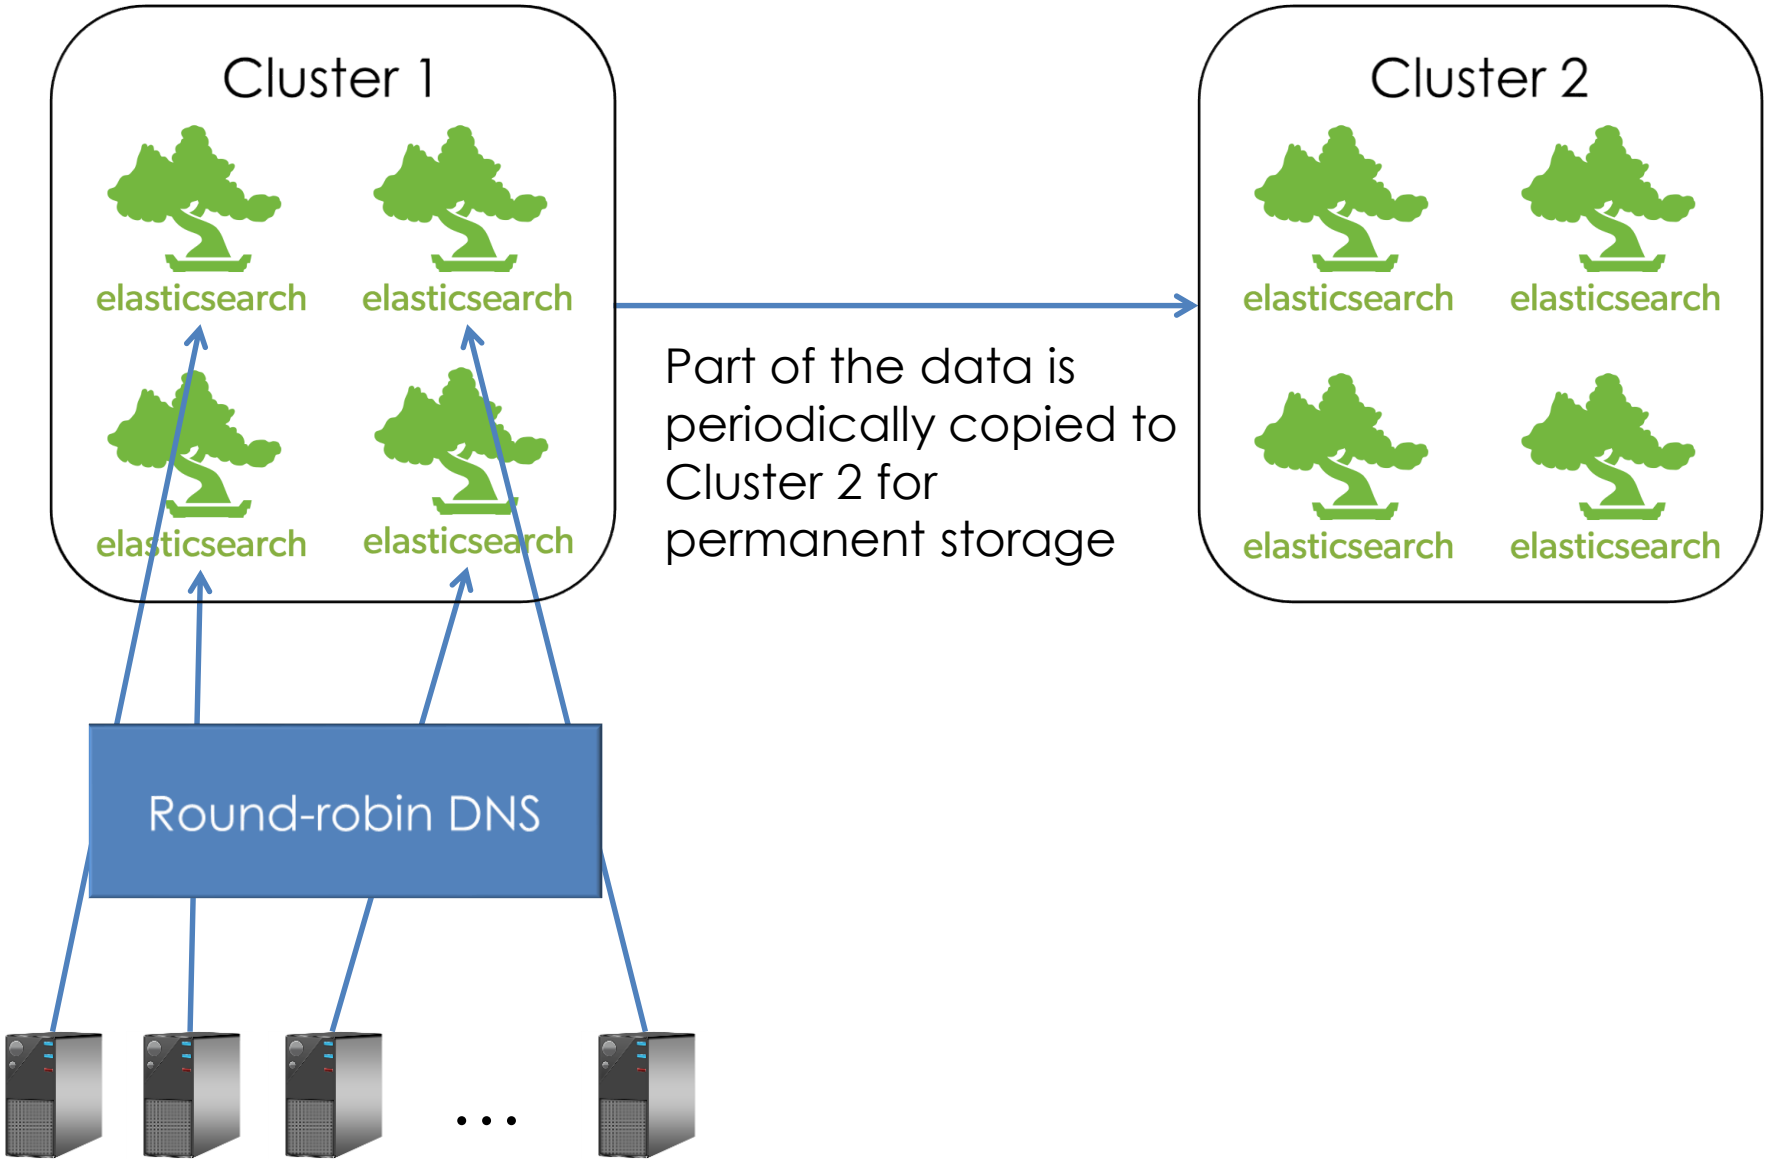

Using IT with Openstack

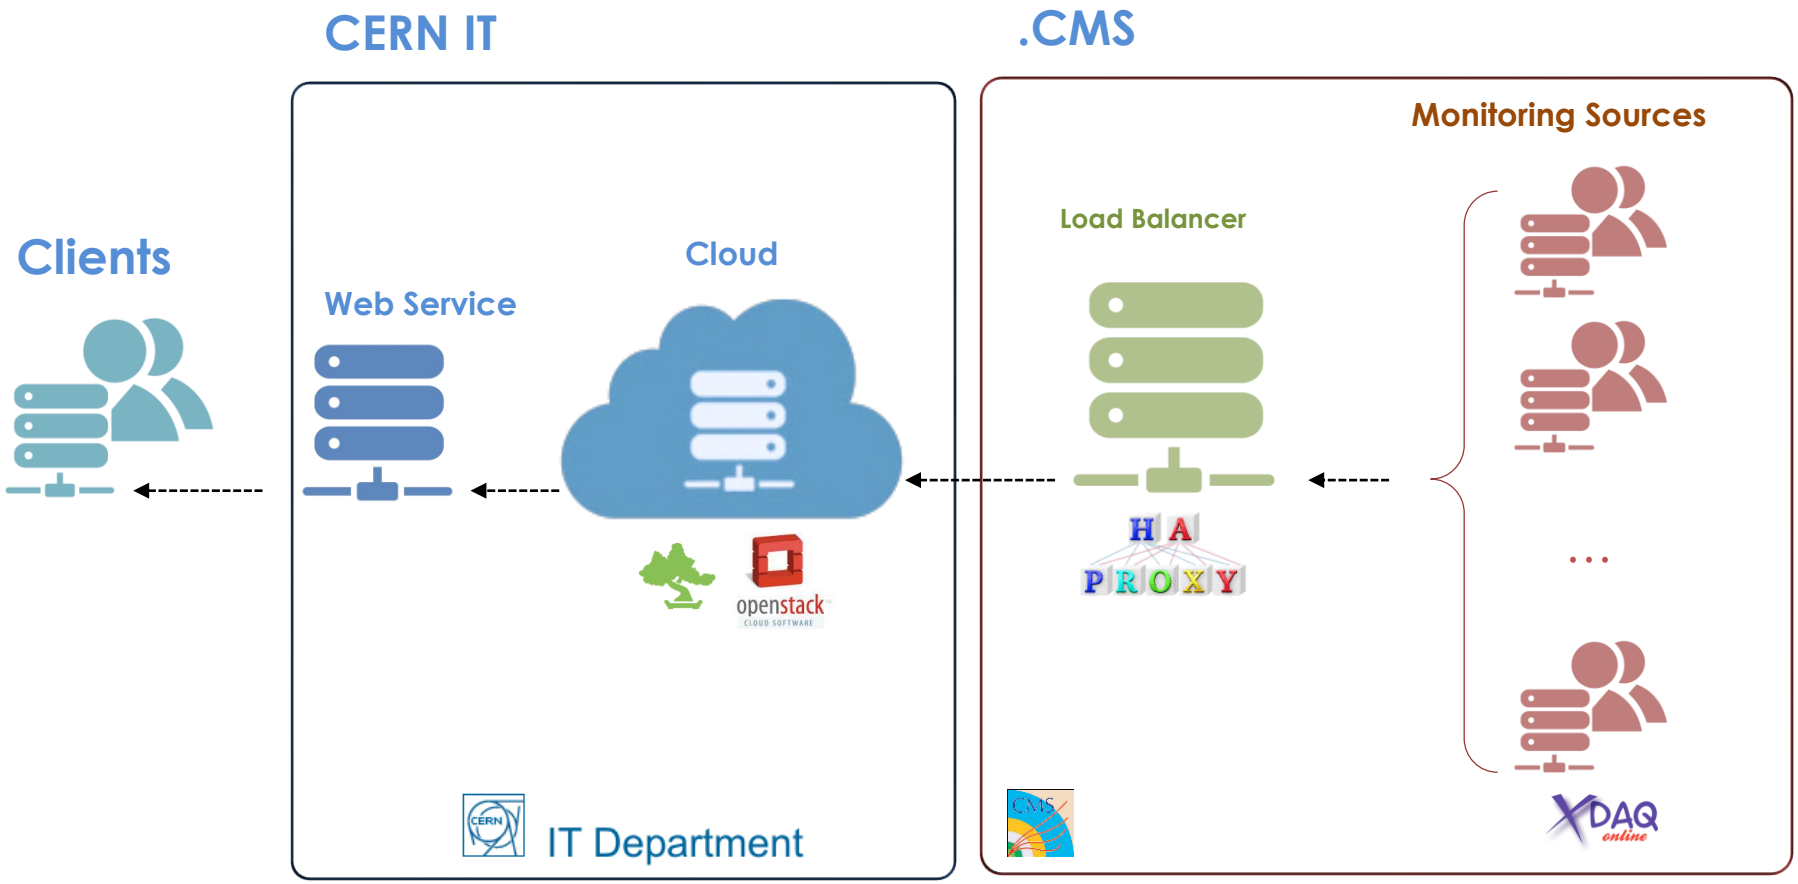

DAQMon

12/04/16 Tue 16:37:26	LHC: BEAM SETUP INJECTION PROBE	DAQ2 State Undefined	Run Number	Lv1 Rate 0.000 kHz	Ev.Size [Compact] 0.00 [12.64] kB	Dead Time [AB] 100.00 [0.00] %	Stream Physics 8.99 Hz	Accepted 0.0 %

Data Flow

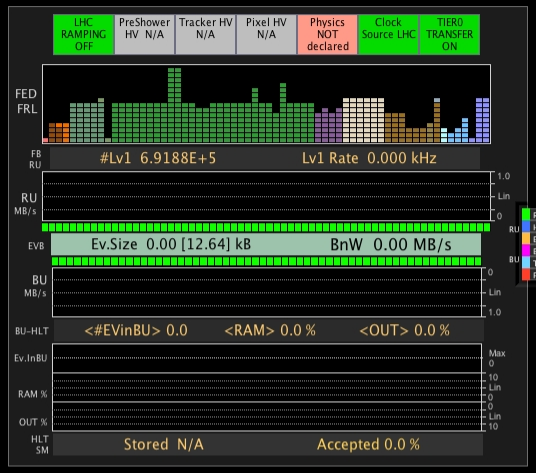

DAQ components

FMM	FED	FRL	EVM	RU	BU
676	456	1	63	58	0

Sub-System		State	FRL	FED	IN
TCDS	IN	Configured	1	1	1
TRG	IN	Configured	12	12	5
PIXEL	Out		40	40	0
PIXEL_UP	Out		2	2	0
TRACKER	Out		249	437	0
ES	Out		26	40	0
ECAL	IN	Configured	54	54	54
HCAL	Out		35	35	0
HF	Out		4	4	0
SCAL	IN	Configured	1	1	1
RPC	IN	Configured	3	3	3
DT	IN	Configured	10	10	10
CSC	IN	Configuring	19	37	37
DAQ	IN	Configured	0	0	0
DCS	IN	Connected	0	0	0
DQM	IN	Configured	0	0	0

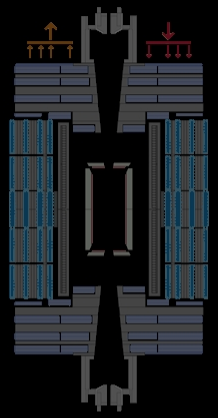

Stream	Tot.Events	Inst. Rate (Hz)
Physics	1.0290E+4	8.989
ALCALUMIPIXELS	6.2400E+2	1.356
ALCAPHISYM	0.0000E+0	0.000
Calibration	6.5901E+4	132.615
DQM	1.4610E+3	0.833
DQMCalibration	6.5980E+3	13.629
DQMEventDisplay	0.0000E+0	0.000
DQMHistograms	6.9188E+5	367.433
EcalCalibration	6.5901E+4	134.089
Error	0.0000E+0	0.000
ExpressCosmics	3.1670E+3	6.600
HLTRates	6.9188E+5	367.433
L1Rates	6.9188E+5	367.433
NanoDST	6.1961E+5	217.923
RPCMON	0.0000E+0	0.000

Beam state legende

States strips chart

Beam & DCS states & DAQ history

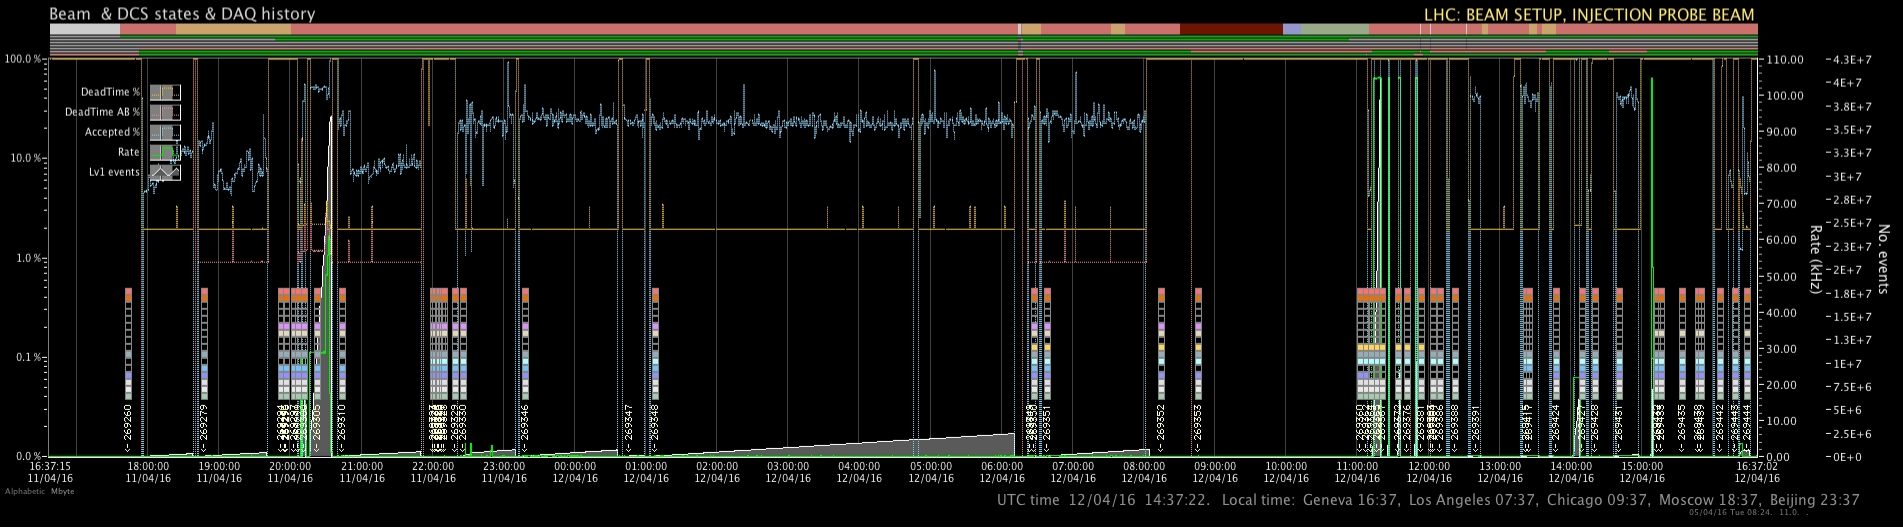

F³Mon GUI

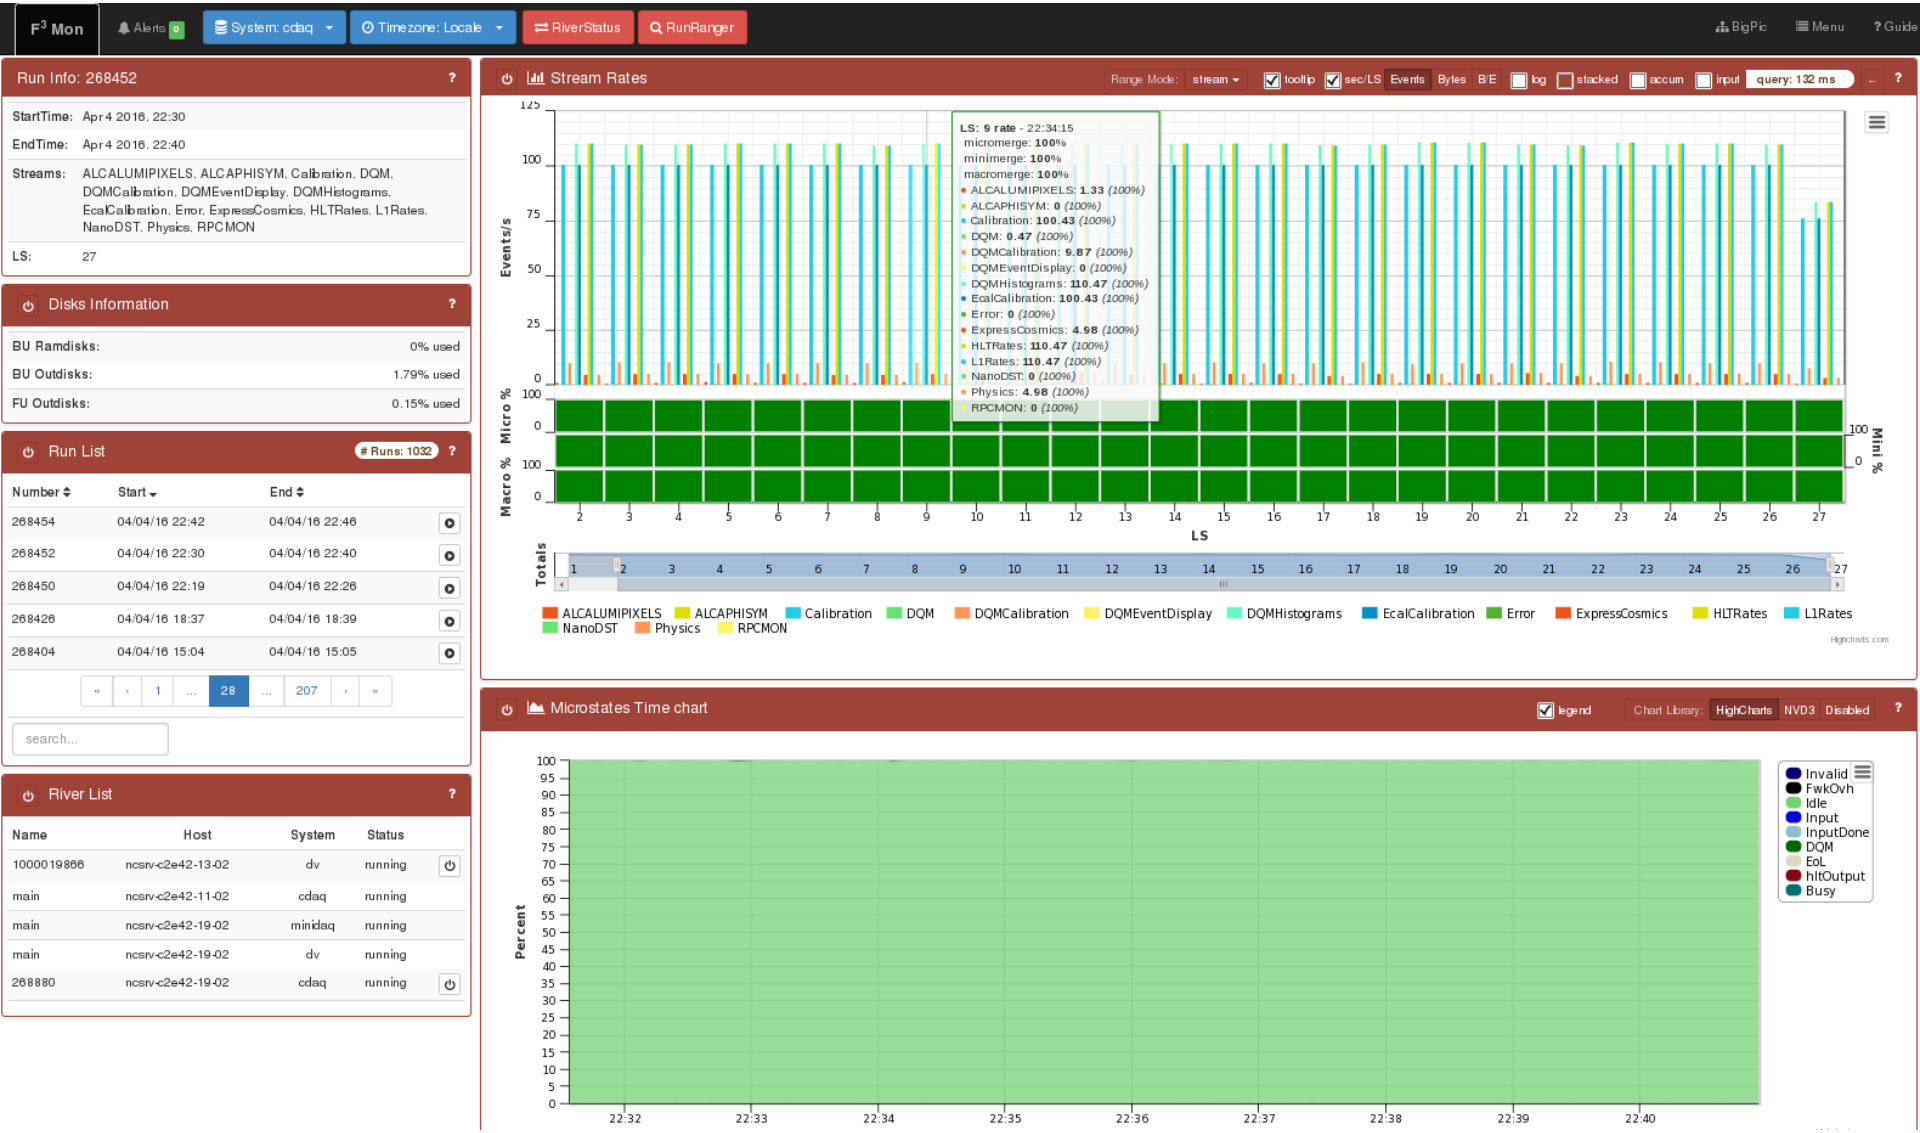
F³ Mon
Alerts 0
System: cdaq
Timezone: Locale
RiverStatus
RunRanger
BigPic
Menu
Guide
Run Info: 268452
StartTime: Apr 4 2016, 22:30
EndTime: Apr 4 2016, 22:40
Streams: ALCALUMIPIXELS, ALCAPHISYM, Calibration, DQM, DQMCalibration, DQMEventDisplay, DQMHistograms, EcalCalibration, Error, ExpressCosmics, HLTRates, L1Rates, NanoDST, Physics, RPCMON
LS: 27
Disks Information
BU Ramdisks: 0% used
BU Outdisks: 1.79% used
FU Outdisks: 0.15% used
Run List
Runs: 1032
Number
Start
End
268454 04/04/16 22:42 04/04/16 22:46
268452 04/04/16 22:30 04/04/16 22:40
268450 04/04/16 22:19 04/04/16 22:26
268426 04/04/16 18:37 04/04/16 18:39
268404 04/04/16 15:04 04/04/16 15:05
search...
River List
Name Host System Status
1000019866 ncsrv-c2e42-13-02 dv running
main ncsrv-c2e42-11-02 cdaq running
main ncsrv-c2e42-19-02 minidaq running
main ncsrv-c2e42-19-02 dv running
268880 ncsrv-c2e42-19-02 cdaq running
Stream Rates
Range Mode: stream
tooltip
sec/LS
Events
Bytes
B/E
log
stacked
accum
input
query: 132 ms
LS: 9 rate - 22:34:15
Events/s
Micro %
Macro %
Mini %
LS
Totals
ALCALUMIPIXELS
ALCAPHISYM
Calibration
DQM
DQMCalibration
DQMEventDisplay
DQMHistograms
EcalCalibration
Error
ExpressCosmics
HLTRates
L1Rates
NanoDST
Physics
RPCMON
Microstates Time chart
legend
Chart Library:
HighCharts
NVD3
Disabled
Percent
Invalid
FwkOvh
Idle
Input
InputDone
DQM
EoL
hltOutput
Busy

Summary

- Elasticsearch is highly scalable and robust system
- Elasticsearch nicely fits requirements of DAQ and HLT monitoring
- Working to bring Elasticsearch based DAQ monitoring to production
- Future perspective is to use Elasticsearch for other purposes besides monitoring: exception storing, log storing, etc.
- Include all sub-detector monitoring
- Persistency of all data (e.g. 24h, 48h, 1 week) for play-back

Thank you for your attention