


# Construction of real-time monitoring system for Grid services based on log analysis at the Tokyo Tier-2 center

**Tomoe Kishimoto**

ICEPP, The University of Tokyo

Mar. 22 2018



**ICEPP**  
The University of Tokyo



# Contents

- ✓ Introduction about ICEPP
- ✓ Status and configuration of the Tokyo Tier-2 center
  - Resource overview and WLCG pledge
  - Site status in the ATLAS experiment
  - Grid middleware
- ✓ Construction of new monitoring system for Grid services
  - Test cluster using the ELK stack
  - Example of new monitors for Grid services
  - Bandwidth monitoring using beat library
- ✓ Summary

# International Center for Elementary Particle Physics




東京大学  
素粒子物理国際研究センター  
International Center for Elementary Particle Physics  
The University of Tokyo


## ✓ Main projects in ICEPP



**ATLAS experiment  
at LHC**



**MEG experiment  
at PSI  
( $\mu \rightarrow e\gamma$  rare decay)**



**R & D for ILC**

## ✓ ATLAS–Japan group

- 17 institutes and ~150 members
- Tokyo Tier2 is the only WLCG site in ATLAS–Japan



# ICEPP regional analysis center

## ✓ Resource overview

- Support only ATLAS VO in WLCG (**Tier2**) and provide ATLAS–Japan dedicated resources (local use)
- Hardwares are leased, and are replaced in every three years
- ~10000 CPU cores including service instances and ~10 PB disk storage (T2 + local use)
  - ▶ 18.11HS06/core (Intel Xenon E5–2680 v3)

## 4th system (2016–2018)

## Single VO and uniform architecture

## ✓ Operation team

- H.Sakamoto (will retire in this Mar. ),  
J.Tanaka, T.Mashimo, N.Tomoaki,  
T.Kishimoto, N.Matsui



# WLCG pledge


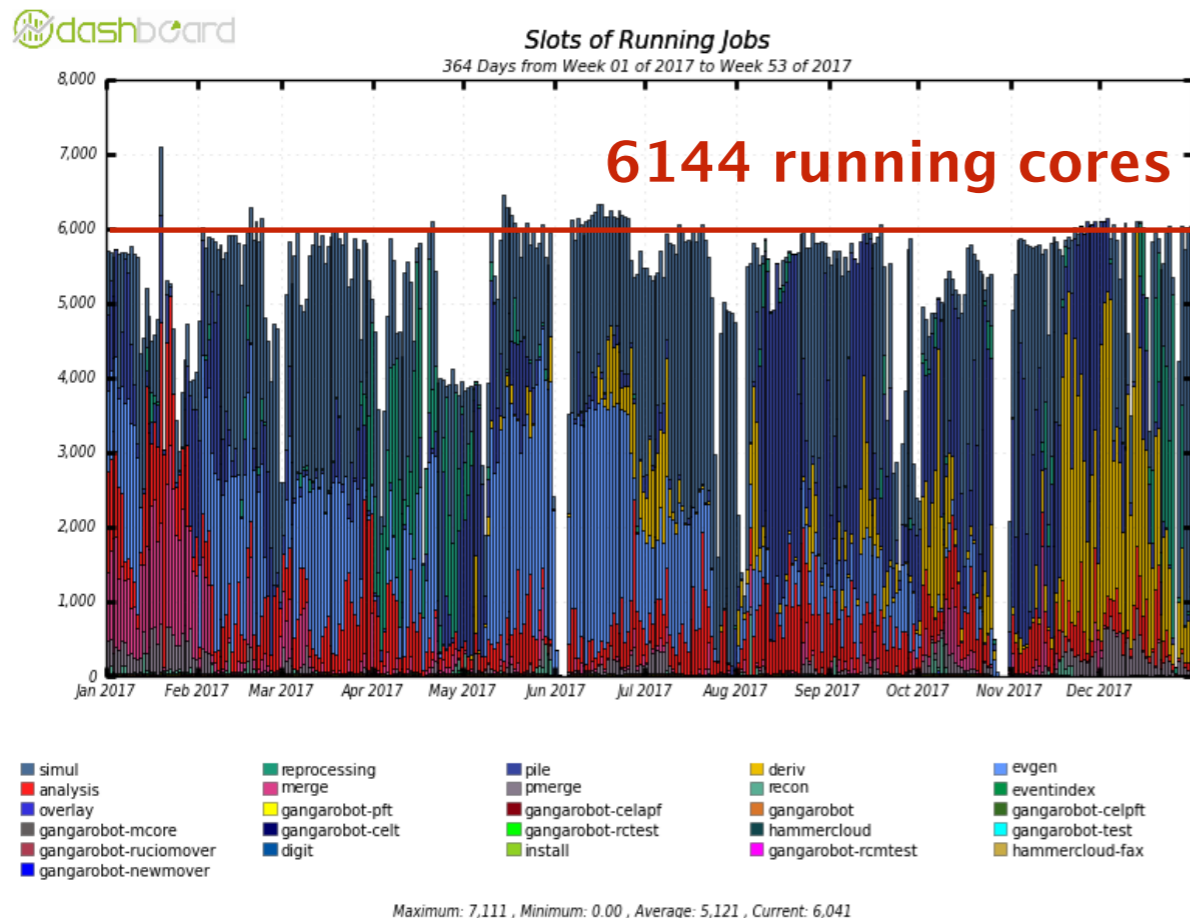
		CPU [HS06]	DISK [TB]	(*)LOCALGROUPDISK [TB]
2017	Pledge	34,000	4,000	-
	Deployed	111,268 (6144 CPU cores)	4,000	1,000
2018	Pledge	40,000	4,800	-
	Deployed	111,268 (6144 CPU cores)	4,800	1,000

(\*) Grid disks for ATLAS–Japan group

## ✓ Tier2 resources

- The current system (4th system) satisfies 2018 WLCG pledge
- New system will be provided for 2019–2021
  - ▶ (Need to migrate 5.8 PB data to the new system...)

# Status in ATLAS experiment



✓ Fraction of # of completed jobs for the last year:

– Production: **4.0% (Tier2)**, 2.2% (All)

– Analysis: **6.3% (Tier2)**, 4.1% (All)

← Good contributions

# of ATLAS-J authors ~ 150

# of ATLAS authors ~ 3000

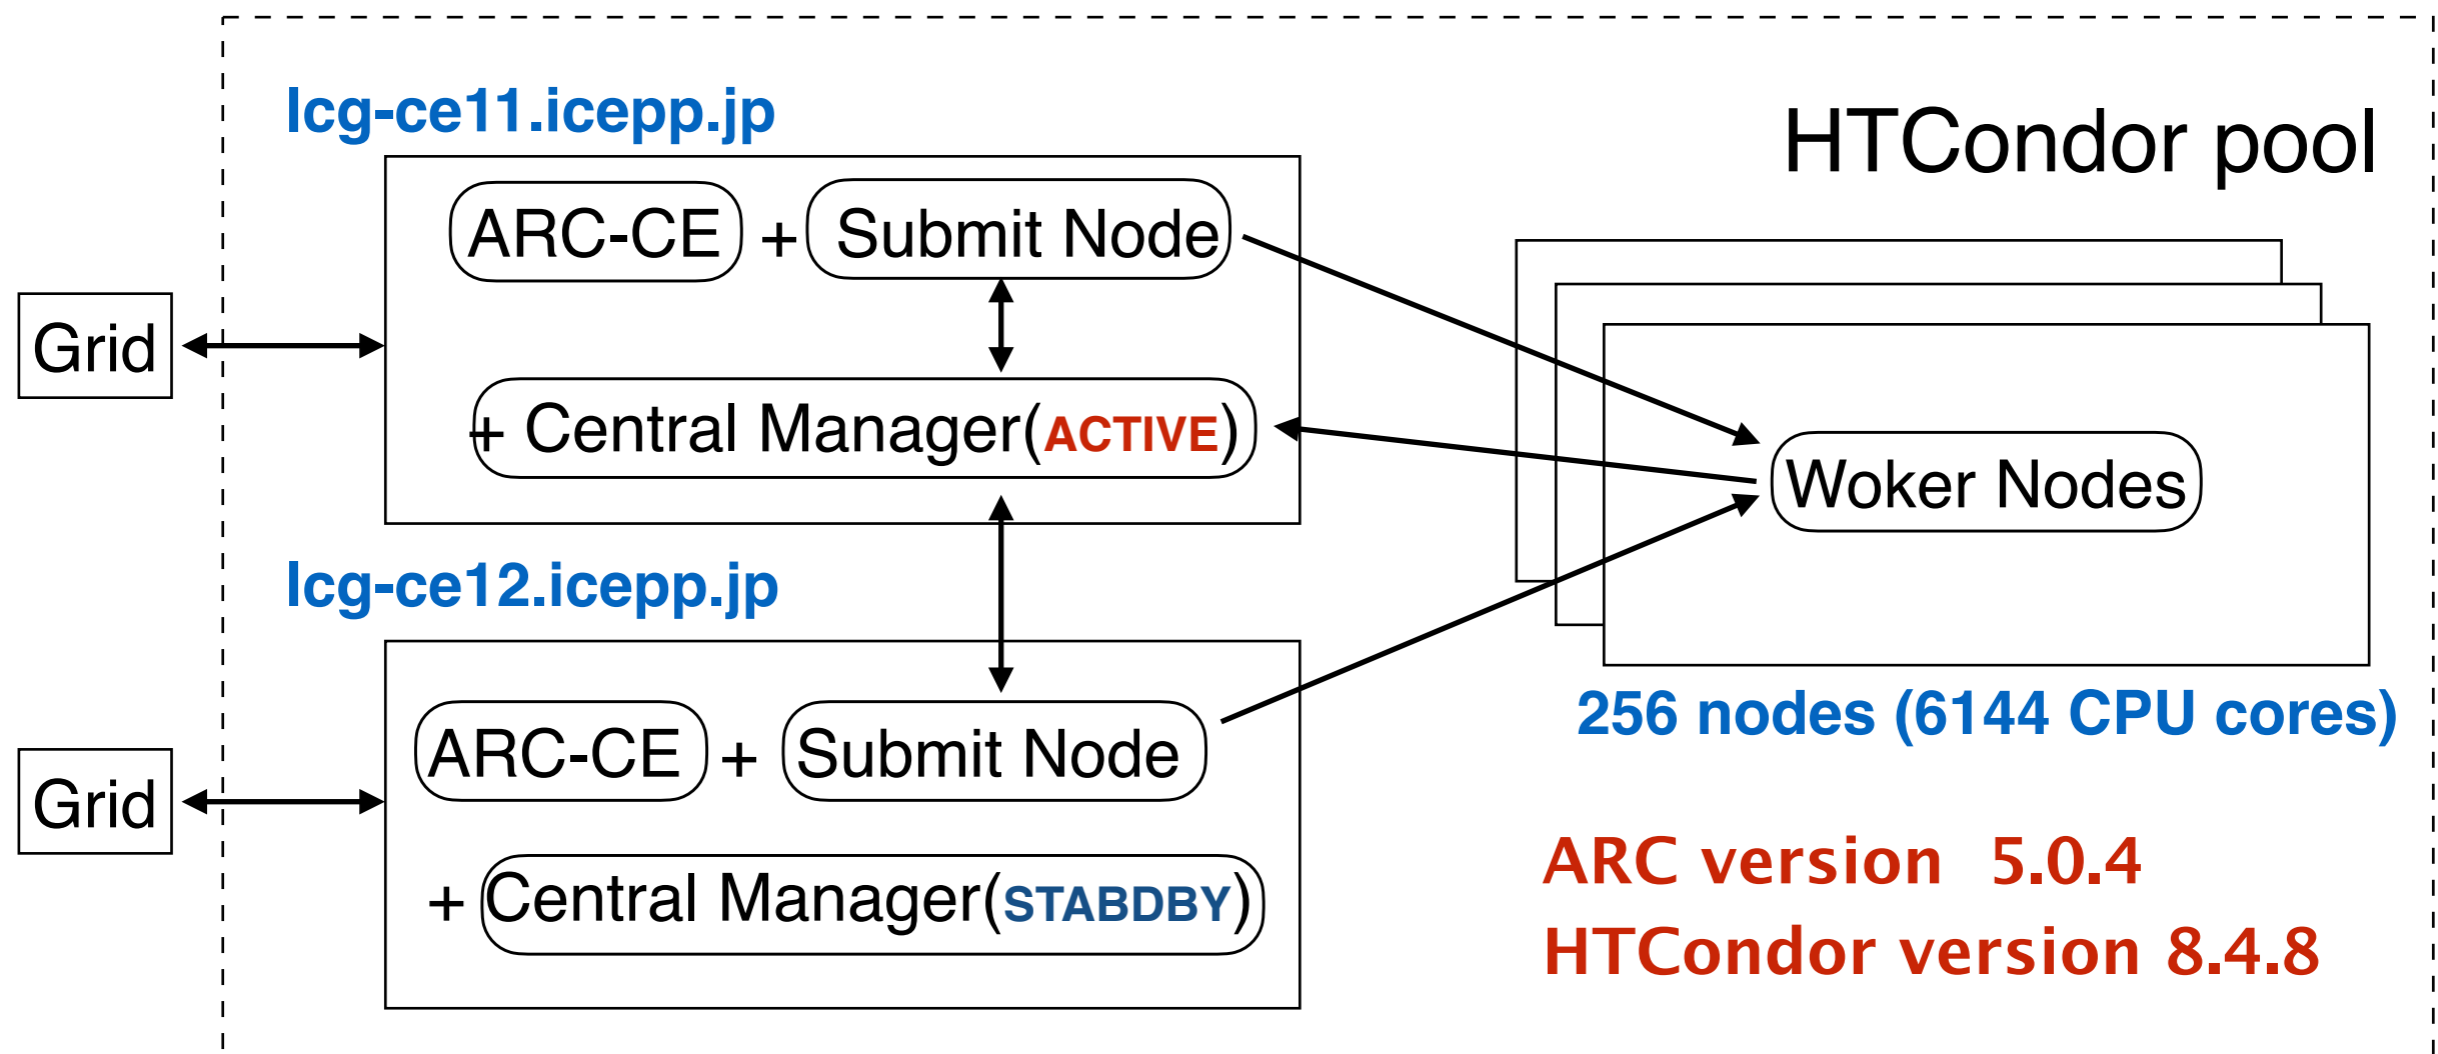
✓ > 99% site availability has been achieved using the 4th system (for 2 years)

# Grid services – CE



## ✓ Computing element:

- Two **ARC-CEs** for redundancy
- **HTCondor** as the local resource batch system
- High availability of central managers

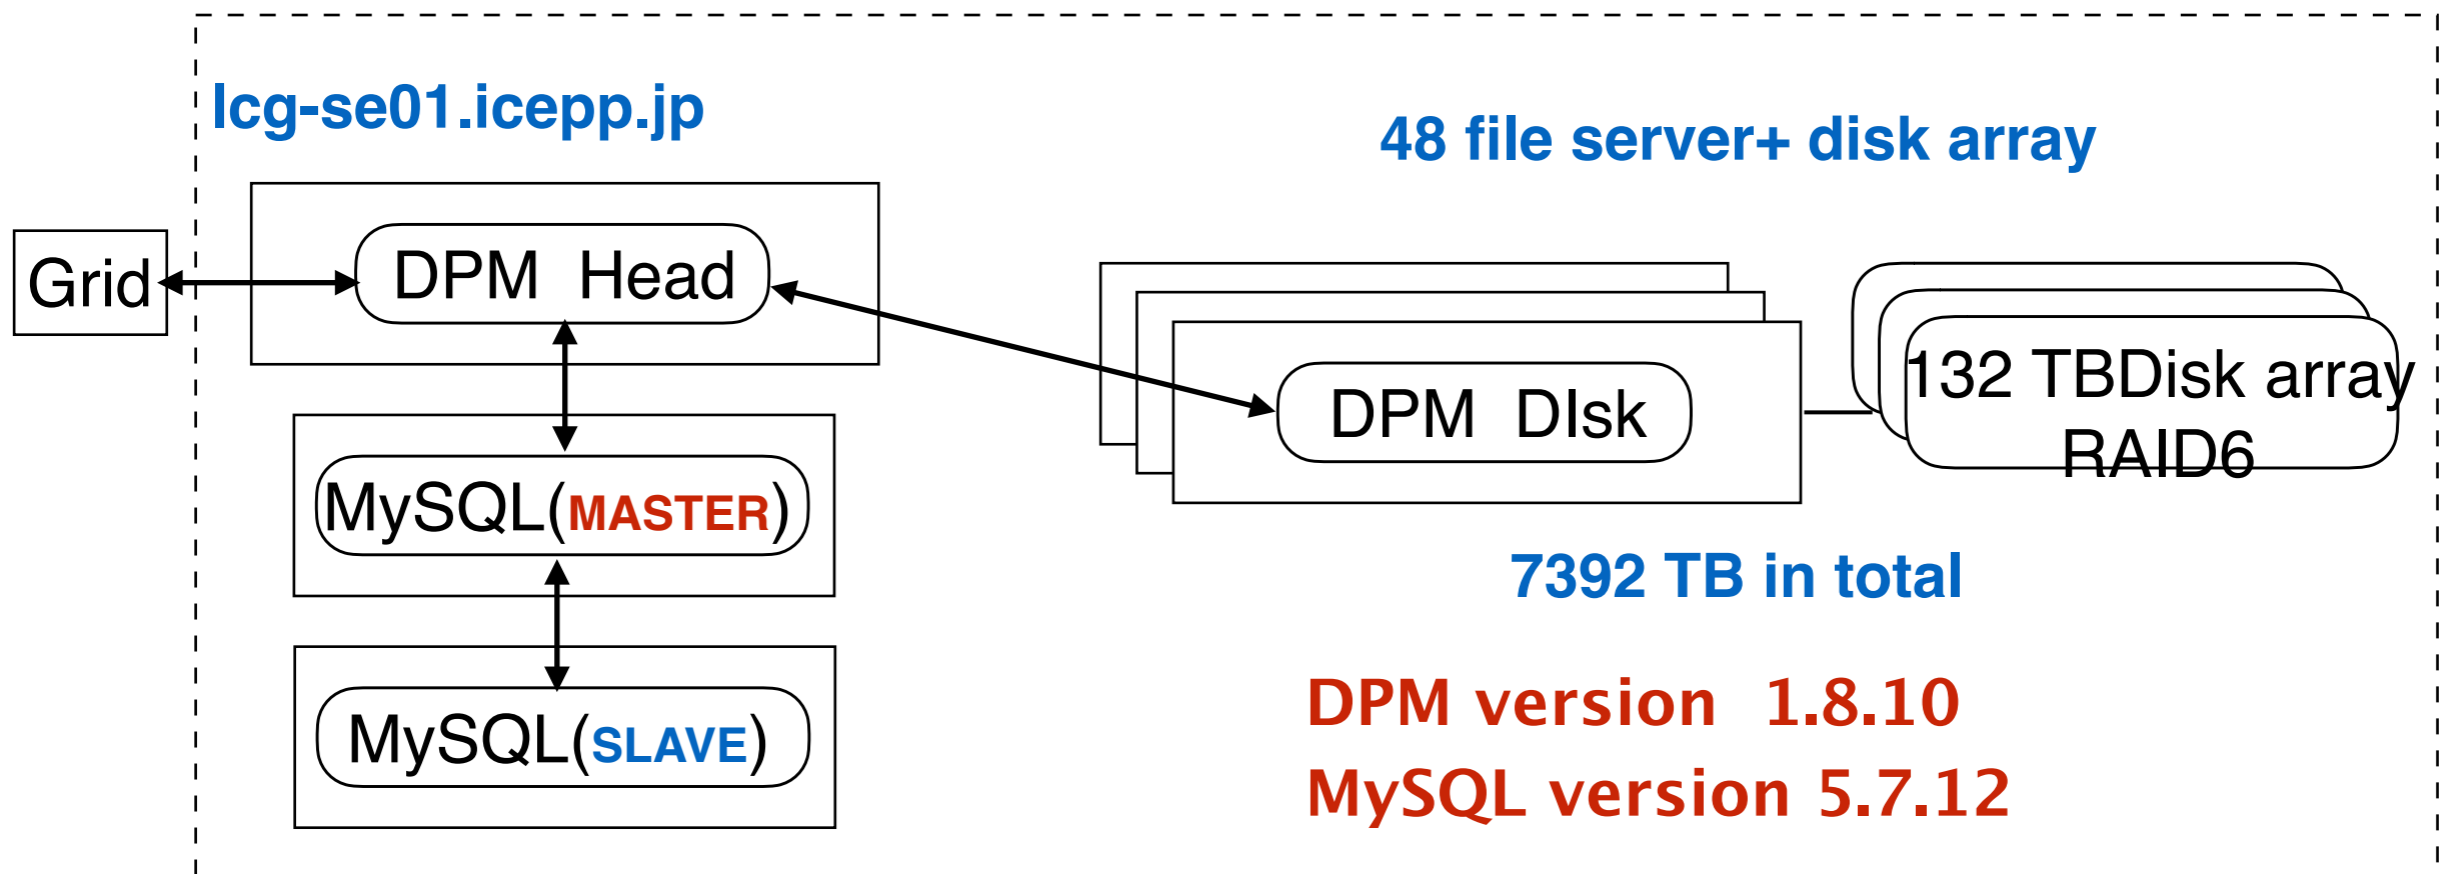


# Grid services – SE and other



## ✓ Storage element:

- DPM speaks GridFTP, XRootD, https, SRM..
- Replication of MySQL DB



## ✓ perfSONAR, BDII, Argus, APEL..





# Grid service logs

- ✓ “Logs” produced by the Grid services provide useful information to check health of the services
  - ← Important for a stable and reliable operation

✓ e.g. ARC-CE log:

Time	Severity	Job global ID	State transition
[2018-02-05 04:05:23]	[Arc] [INFO] [6734/2]	roaNDm7Cn2rnG457in1NHNUqABFKDmABFKDmVi1XDmABFKDmpGQQXn:	State: ACCEPTED: parsing job description
[2018-02-05 04:05:23]	[Arc] [INFO] [6734/2]	roaNDm7Cn2rnG457in1NHNUqABFKDmABFKDmVi1XDmABFKDmpGQQXn:	State: ACCEPTED: moving to PREPARING
[2018-02-05 04:05:23]	[Arc] [INFO] [6734/2]	roaNDm7Cn2rnG457in1NHNUqABFKDmABFKDmVi1XDmABFKDmpGQQXn:	State: PREPARING from ACCEPTED
[2018-02-05 04:05:31]	[Arc.DataStaging.DTR] [INFO] [6734/7]	roaNDm7Cn2rnG457in1NHNUqABFKDmABFKDmVi1XDmABFKDmpGQQXn:	All downloads finished successfully
[2018-02-05 04:05:31]	[Arc] [INFO] [6734/2]	roaNDm7Cn2rnG457in1NHNUqABFKDmABFKDmVi1XDmABFKDmpGQQXn:	State: SUBMIT from PREPARING
[2018-02-05 04:05:31]	[Arc] [INFO] [6734/2]	roaNDm7Cn2rnG457in1NHNUqABFKDmABFKDmVi1XDmABFKDmpGQQXn:	state SUBMIT: starting child: /usr/share/arc/submit-condor-job
[2018-02-05 04:05:32]	[Arc] [INFO] [6734/2]	roaNDm7Cn2rnG457in1NHNUqABFKDmABFKDmVi1XDmABFKDmpGQQXn:	state SUBMIT: child exited with code 0
[2018-02-05 04:05:32]	[Arc] [INFO] [6734/2]	roaNDm7Cn2rnG457in1NHNUqABFKDmABFKDmVi1XDmABFKDmpGQQXn:	State: INLRMS from SUBMIT
[2018-02-05 04:22:01]	[Arc] [INFO] [6734/2]	roaNDm7Cn2rnG457in1NHNUqABFKDmABFKDmVi1XDmABFKDmpGQQXn:	Job finished
[2018-02-05 04:22:01]	[Arc] [INFO] [6734/2]	roaNDm7Cn2rnG457in1NHNUqABFKDmABFKDmVi1XDmABFKDmpGQQXn:	State: FINISHING from INLRMS
[2018-02-05 04:22:03]	[Arc] [INFO] [6734/2]	roaNDm7Cn2rnG457in1NHNUqABFKDmABFKDmVi1XDmABFKDmpGQQXn:	State: FINISHED from FINISHING

- We can judge that the job was processed at CE without errors

# Monitoring system

✓ Constructing a new real-time monitoring system based on log analysis using the **ELK stack**

- ELK stack provides efficient way of log processing, storing, query and visualization



## Log processing

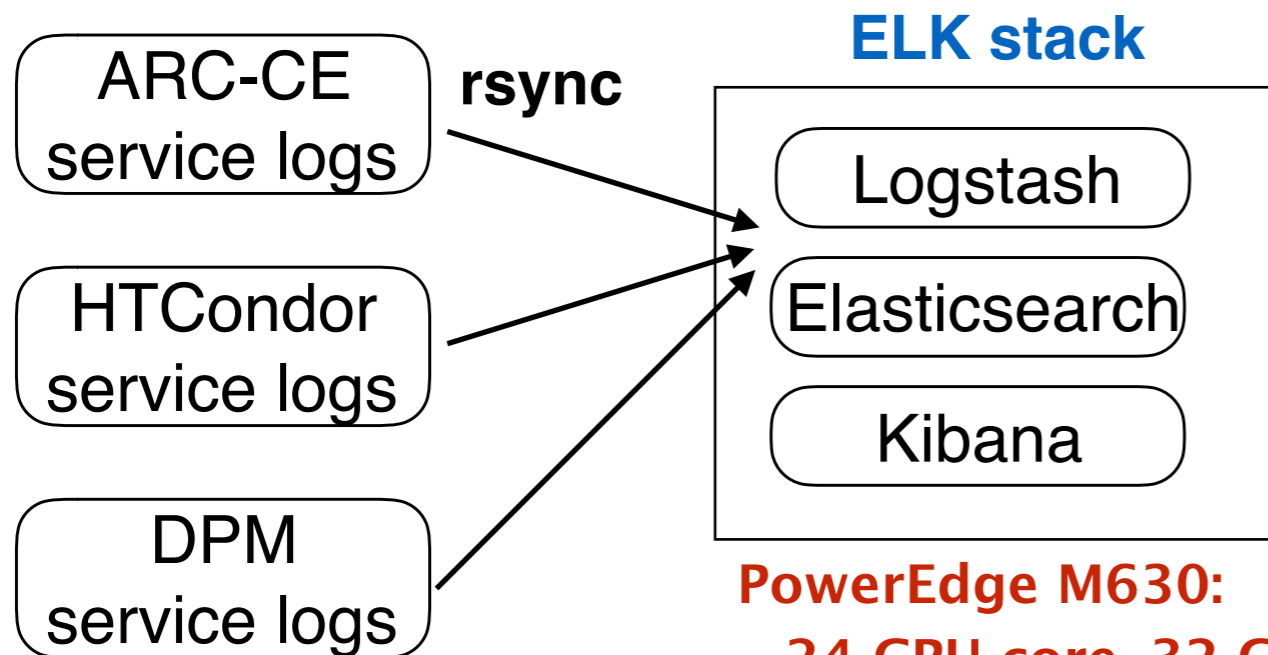
```
[2018-02-05 08:39:45] [Arc] [INFO] [6734/2]
UVZMDmQUr2rnG457in1NHNUqABFKDmABFKDmWy8aDmA
BFKdMwQcvdn: State: INLRMS from SUBMIT
```

**1 log line = 1 record in DB**

@timestamp	February 5th 2018, 17:40:04.696
@version	1
_id	AWF1H0qG019fJ1Qw3GDC
_index	logstash-arc-service-2018.02.05
_score	-
_type	arc.service
data	UVZMDmQUr2rnG457in1NHNUqABFKDmABFKDmWy8aDmA
date	February 5th 2018, 17:39:45.000
host	lcg-elk01.icepp.jp
path	/data/arc-service/lcg-ce11.icepp.jp/grid-manager.log
process	INLRMS
server	lcg-ce11.icepp.jp
service	Arc
severity	INFO
type	arc.service

# Test environment

- ✓ Simple configuration: 1 host for the ELK stack



- Currently, rsync is used to collect logs (can be replaced with Filebeat)
- Dedicated log analysis (logstash conf. ) for each service

**PowerEdge M630:**  
24 CPU core, 32 GB memory  
ELK version 5.6

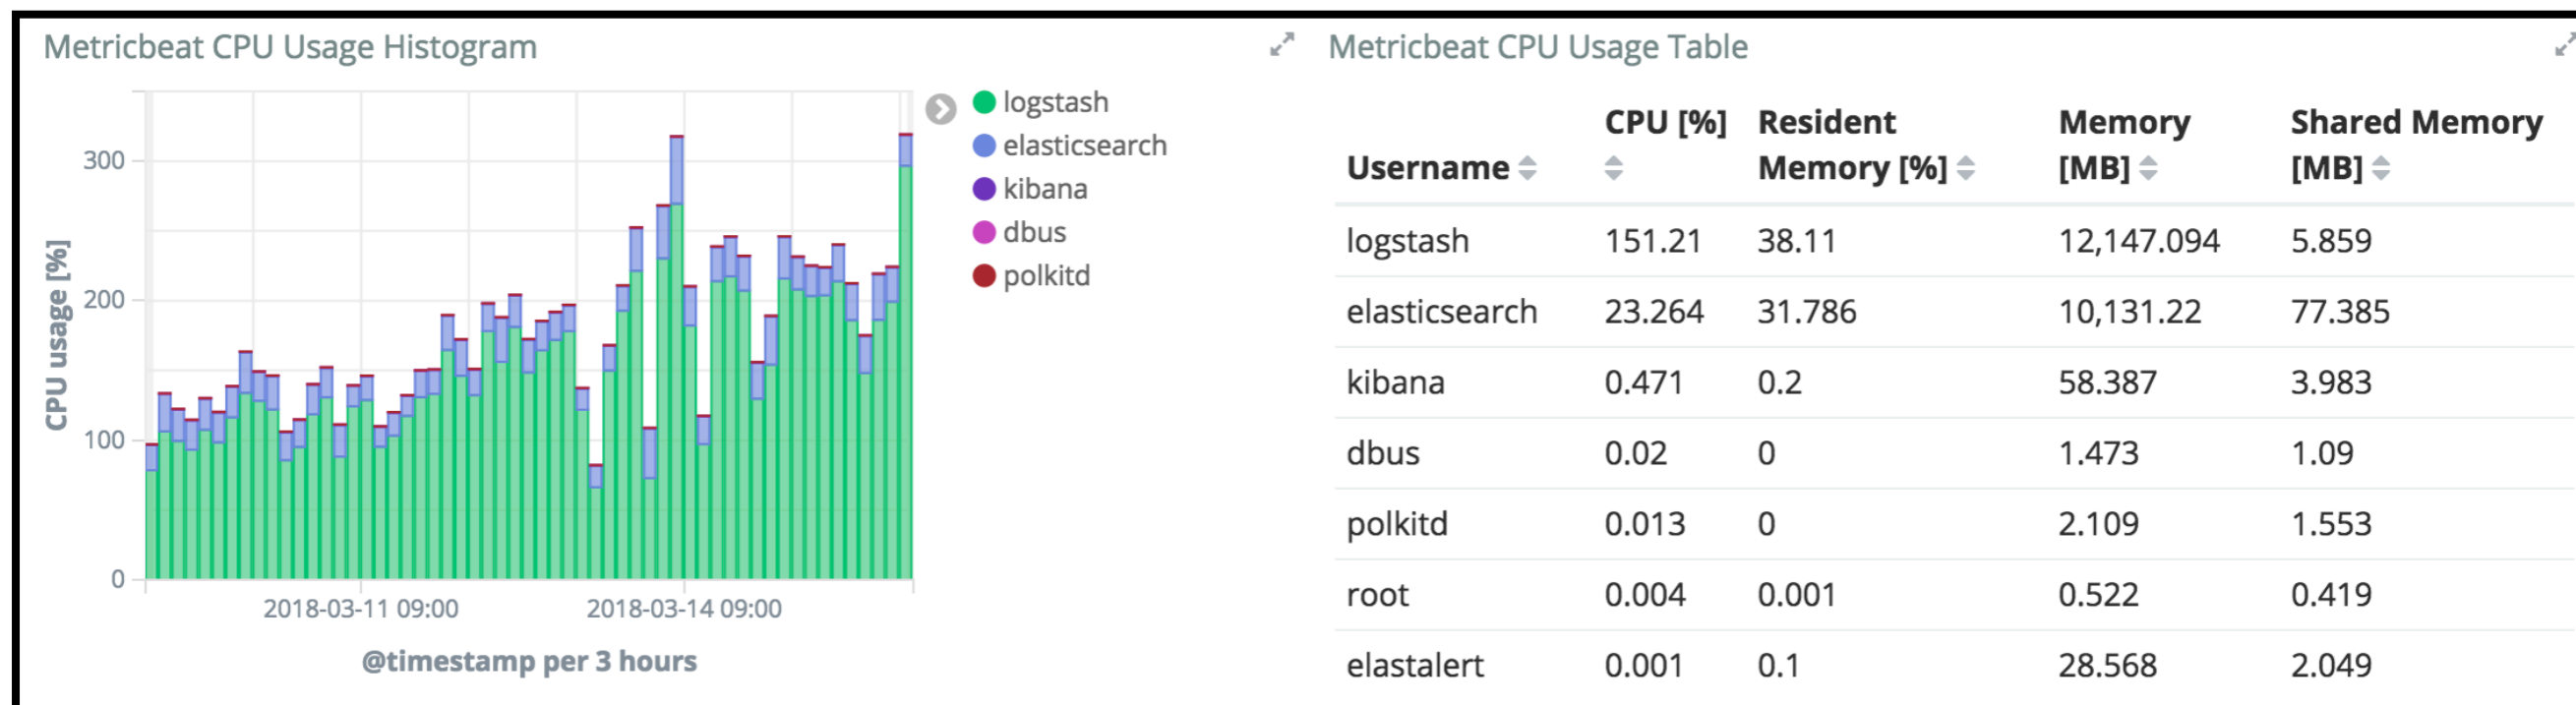
- ✓ Log frequency

	ARC-CE	HTCondor	DPM
Records (log lines)	~300/min	~5,000/min	~50,000/min

- DPM is the dominant sources of records

# Resource usage of ELK host

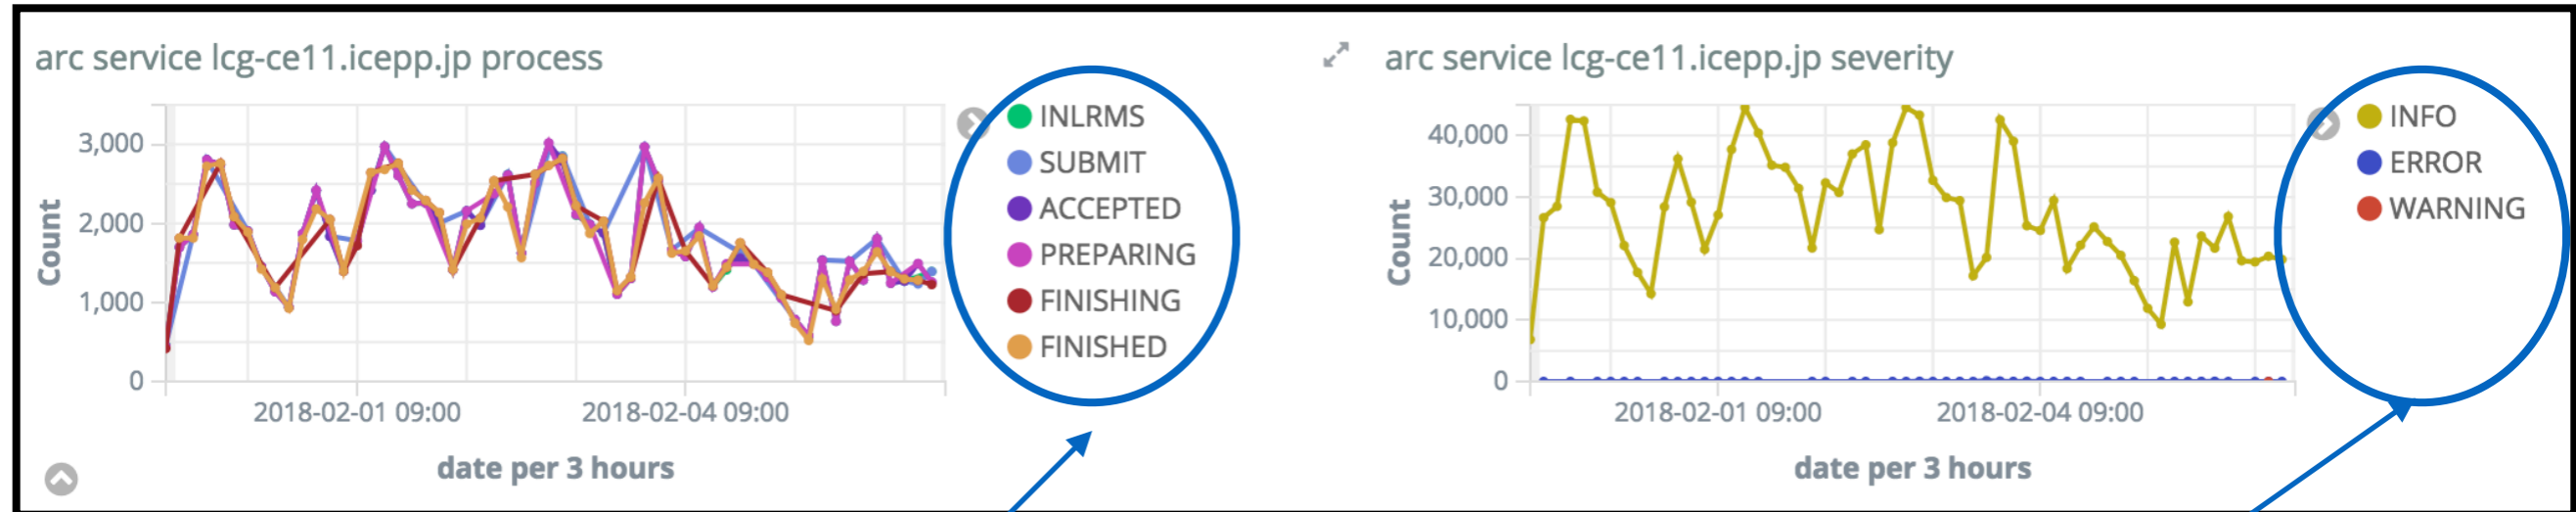
Data is collected by Metricbeat



- Average CPU usage is ~200% (24 cores = max 2400%), logstash is the dominant source
  - ▶ Correlated with DPM loads
- JVM heap size is 12GB for logstash and elasticsearch
- Elasticsearch DB increases ~20 GB per day
  - ▶ Delete records, which is passed 1 month

# ARC-CE and HTCondor monitoring

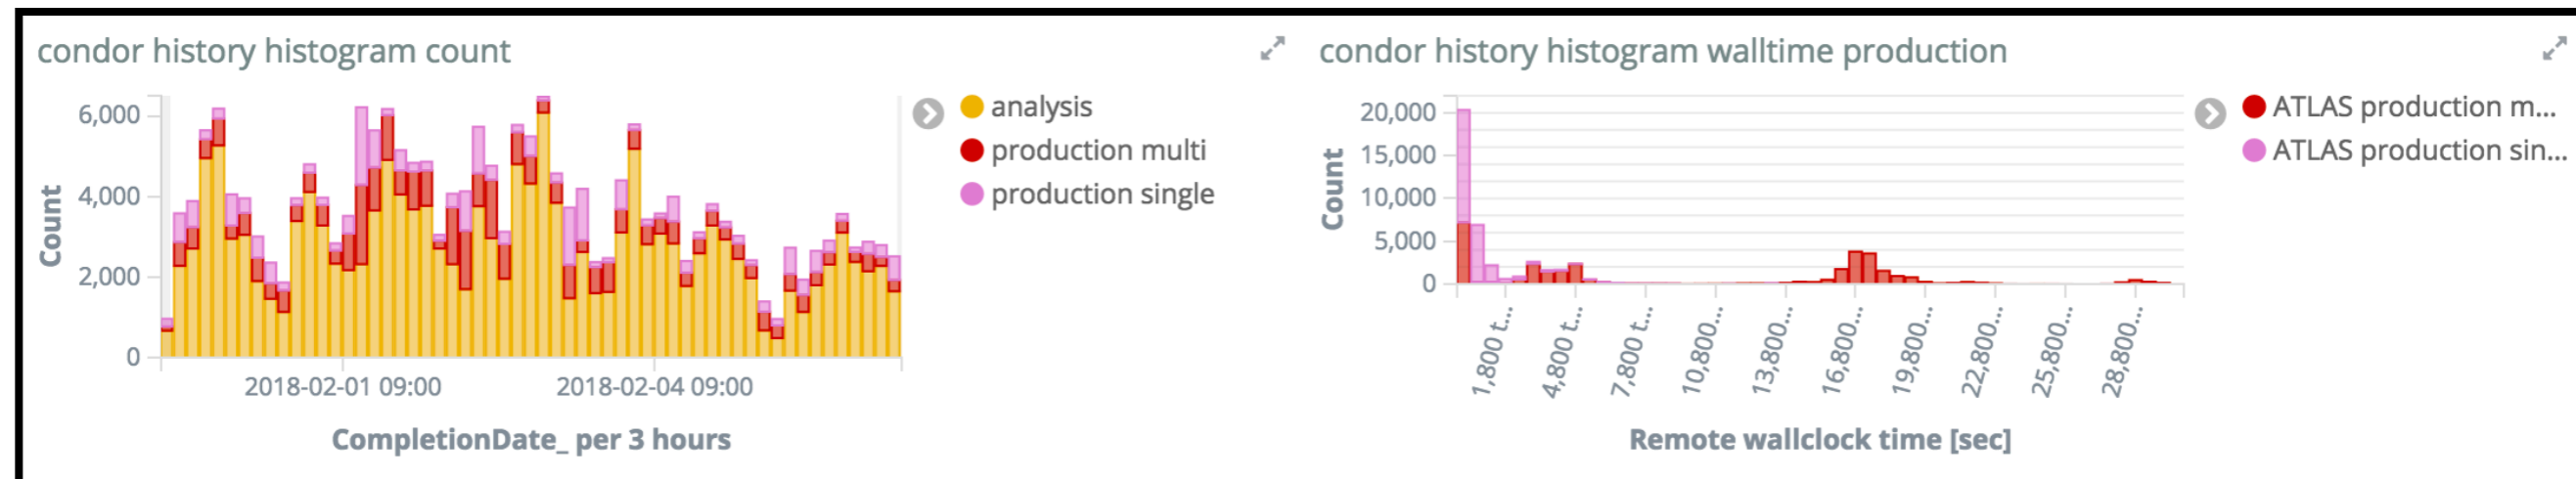
## ARC-CE



Counts of job status  
→ Can detect pending status

log severity

## HTCondor

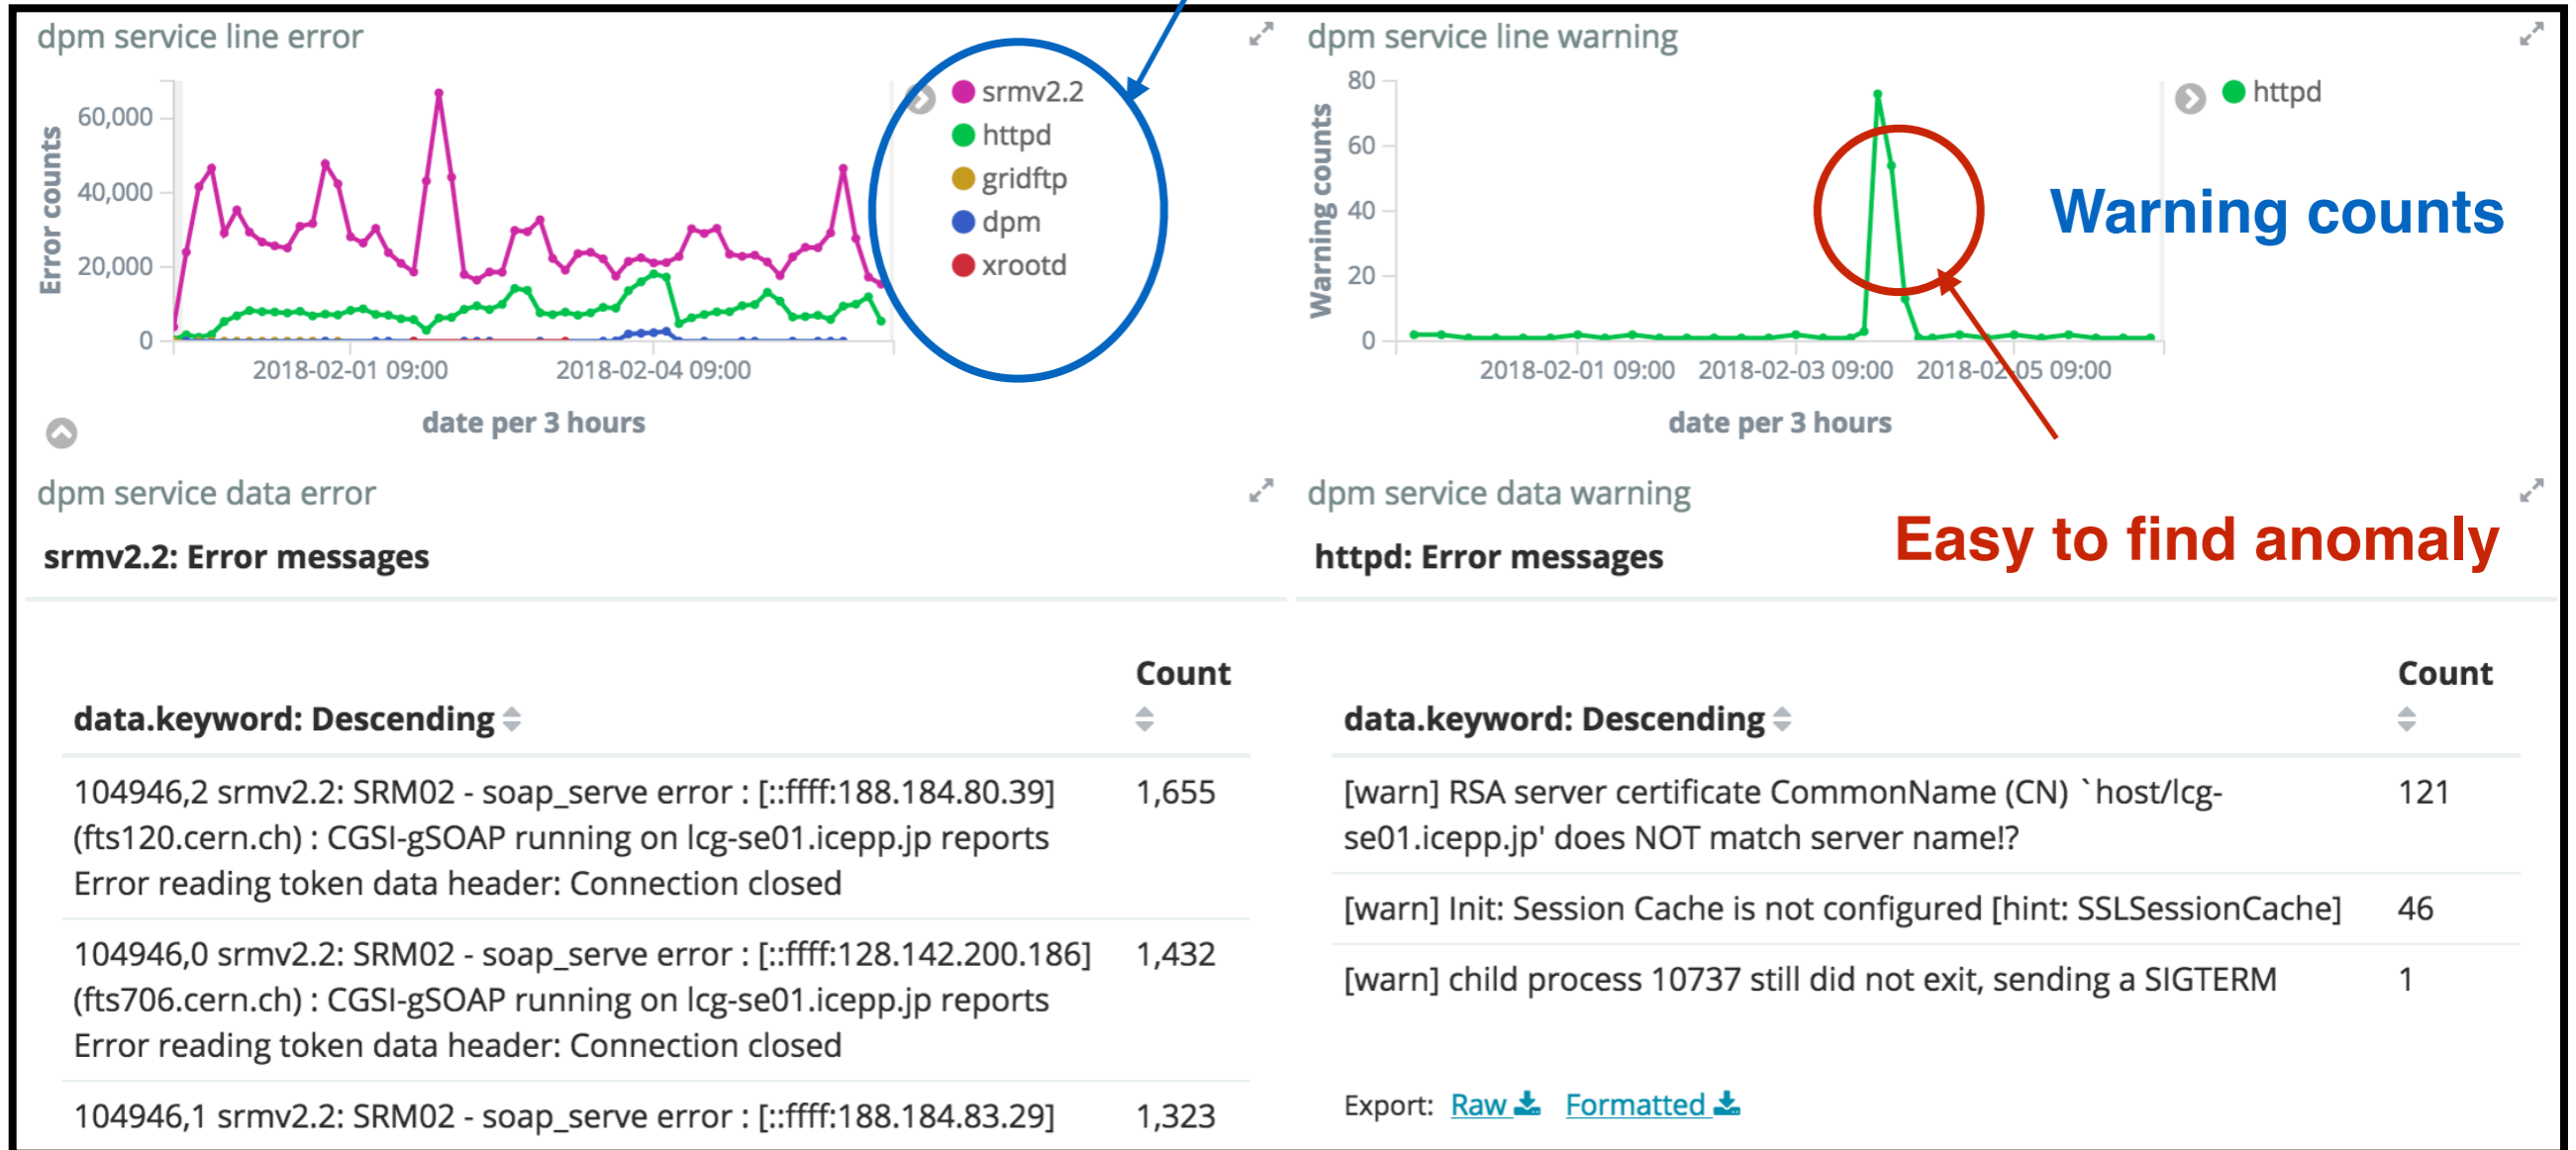


Job throughput

Job duration

# DPM monitoring

## Error counts for each protocol



- ✓ Grid service logs are well visualized.
- ✓ Alerting is implemented using 'elastalert'
  - <https://github.com/Yelp/elastalert>

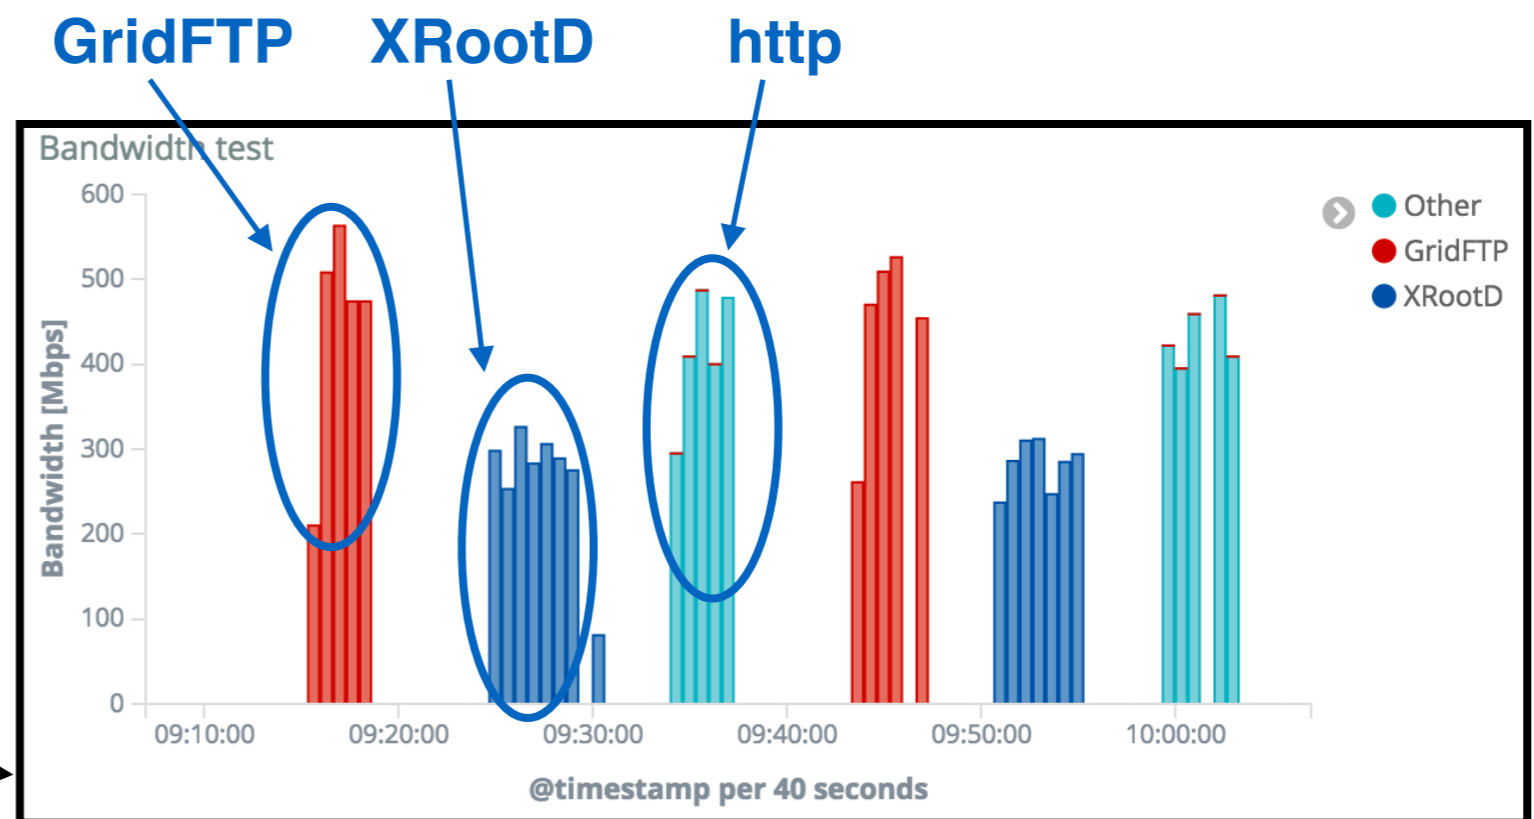
# iftopbeat

- ✓ DPM supports several protocols for data transfer
- ✓ Developed 'iftopbeat' to measure real-time throughput of data transfer for each protocol

Beats: lightweight data shipper



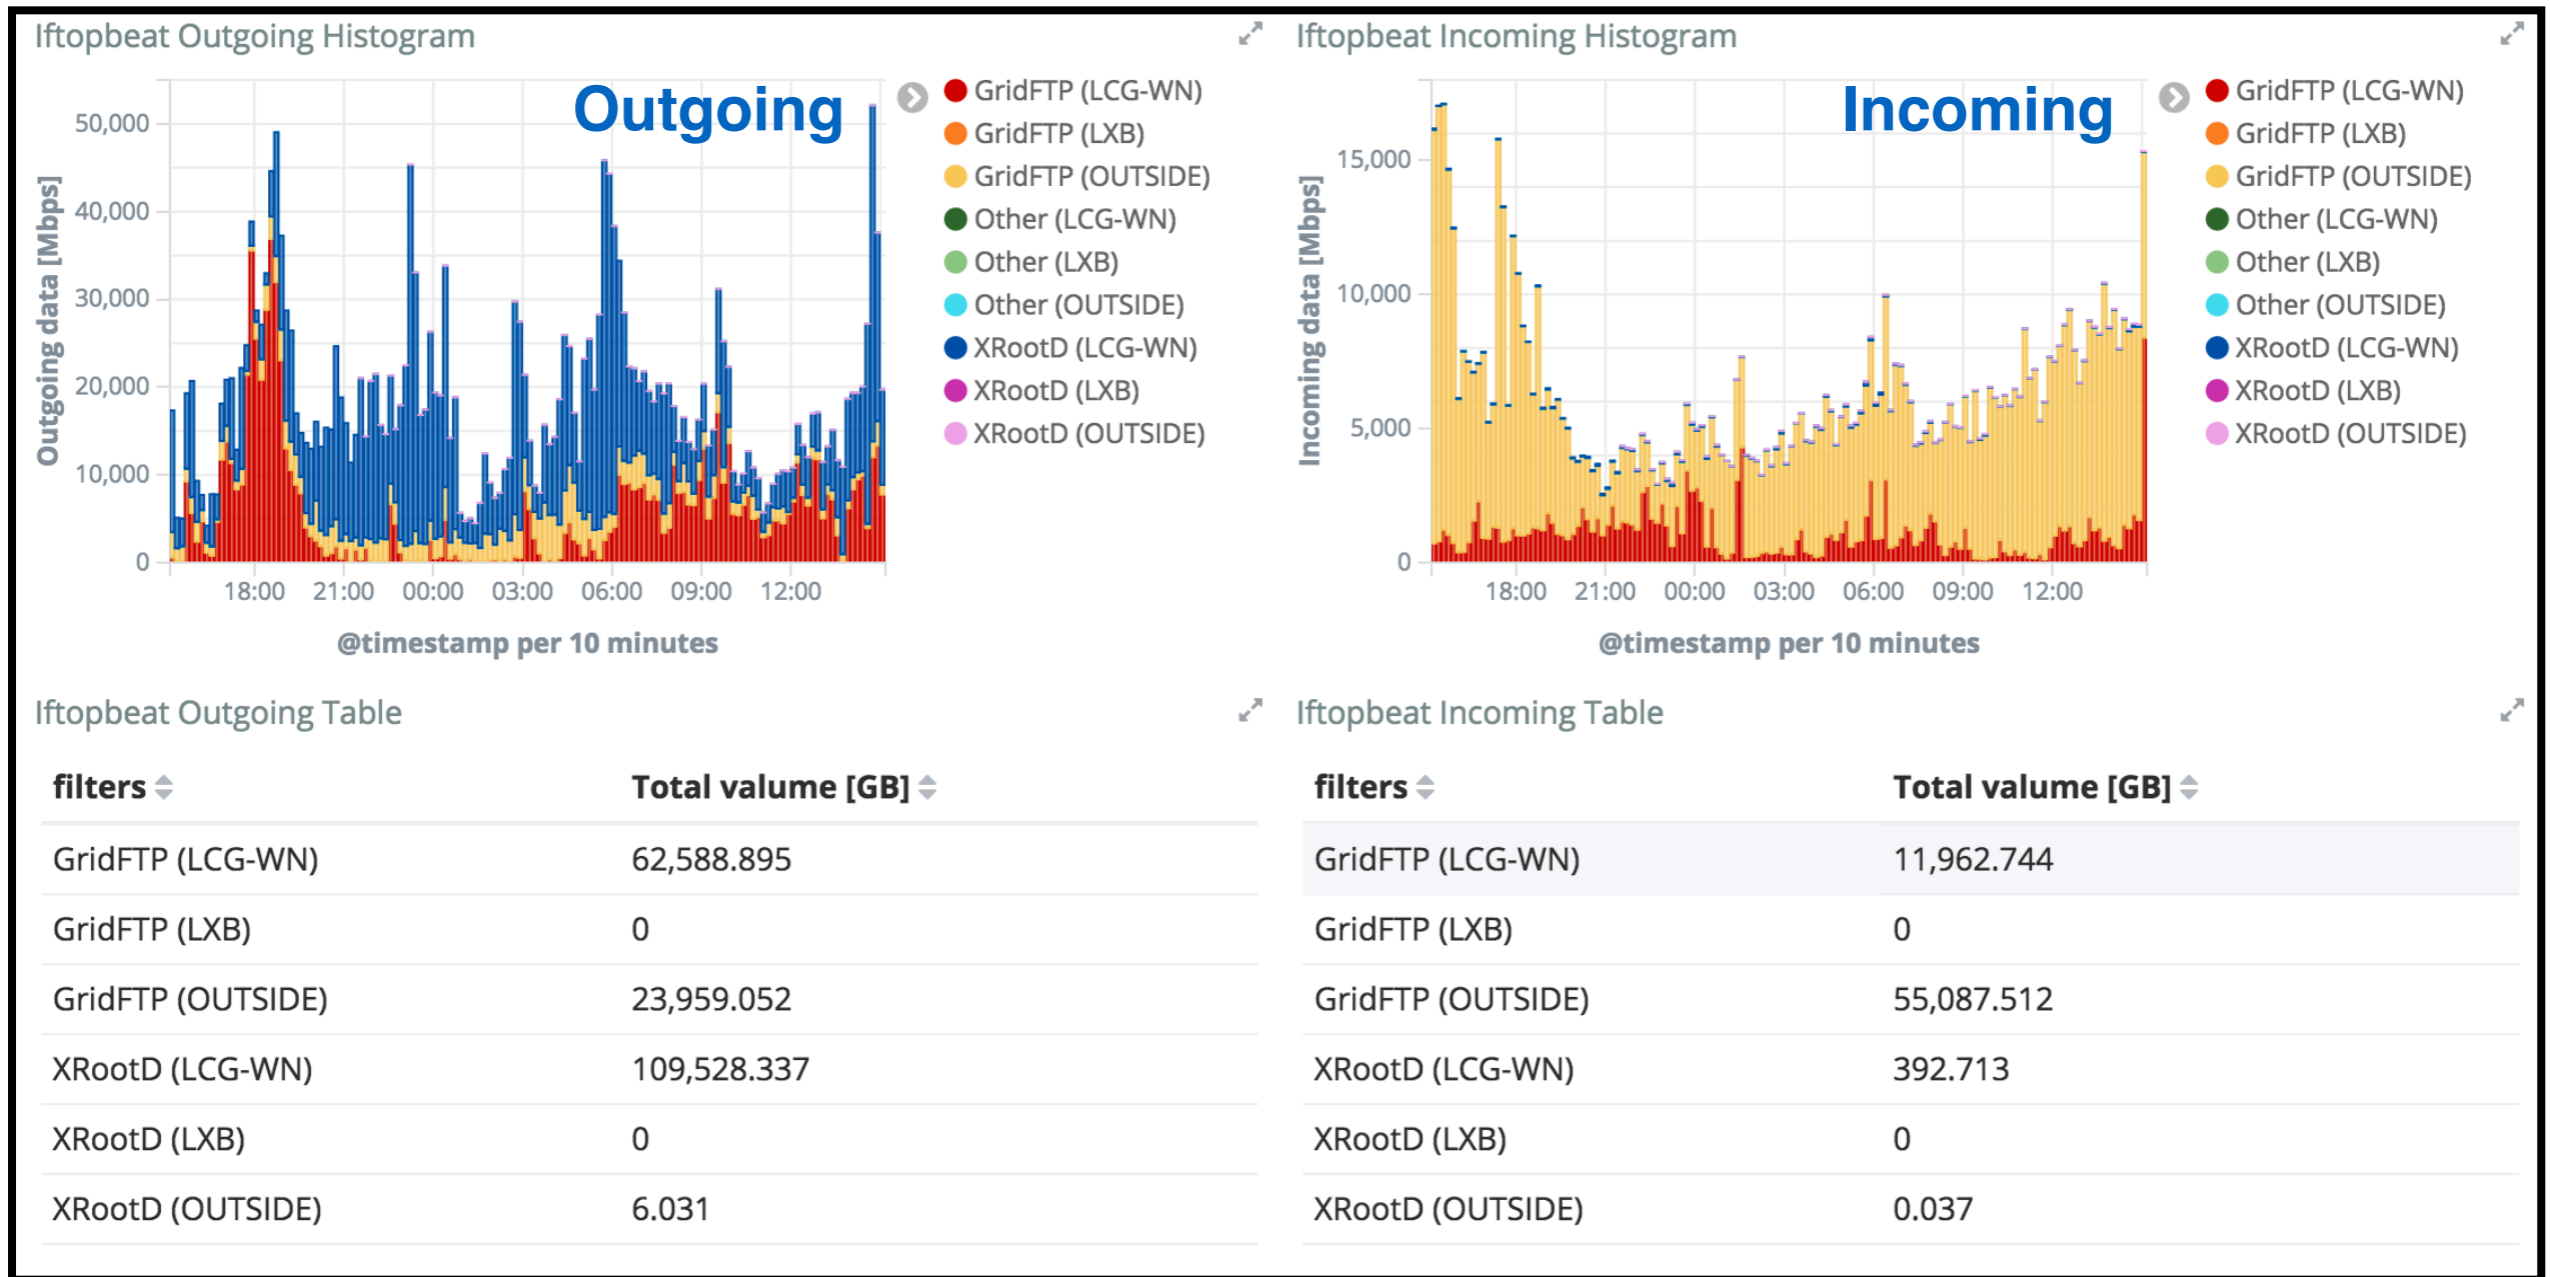
- Data are measured by iftop command, and are directly sent to elasticserach using libbeat platform
- Written by golang
- Test of a 10GB file transfer



Well measured/visualized

# iftopbeat

## ✓ Monitoring for real data transfers



Easy to check correlation between service logs and transfer loads



# Summary

- ✓ Tokyo Tier2 center with the 4th system is running
  - Providing enough computing resources for ATLAS
  - > 99% site availability is achieved
- ✓ Constructing a new monitoring system for Grid services using the ELK stack
  - ELK stack can handle the logs produced at TOKYO Tier2 center so far
  - Logs are well visualized, and easy to check correlations between the Grid services
  - Beats libraries provides efficient way to ship data to elasticsearch
    - ▶ Developed iftopbeat as an example