LHCb Online Log Analysis and Maintenance System

Jean-Christophe Garnier

13 October 2011

Introduction

- Writing every line of log of the LHCb Online Cluster in a central place
  - 2000 Linux machines
    - ★ Event Filter Farm and other data acquisition systems
    - ★ User nodes
    - ★ Storage
  - 200 Windows machines
    - ★ Experiment Control System device drivers
    - ★ Domain Controllers
    - ★ Numerous services
  - 60 switches and routers

➔ $O(10000)$ log sources

Outline

Outline

Requirements

- Usual requirements:
  - Reliable
  - Available
  - Scalable
  - Fault tolerant
  - Standard -> use well known open source tools
  - Cheap
- Analysis
  - Fast
  - Logs provide complementary information to snmp
  - Efficient alert system
- Distribution to users for custom analysis
  - Read only

Log servers

- Linux based system
  - $O(10000)$ sources for 8 severities and 22 syslog facilities
- Use Rsyslog
  - Rules can contain regular expressions
  - About 100 rules to write 12000 files
  - Push files to syslog protocol
- Keeping logs for a few months
  - Using standard logrotate
  - Compression ratio: $1/200$
  - Security relevant logs are kept longer
- Distribute log as a file system hierarchy via NFS and Samba

Architecture

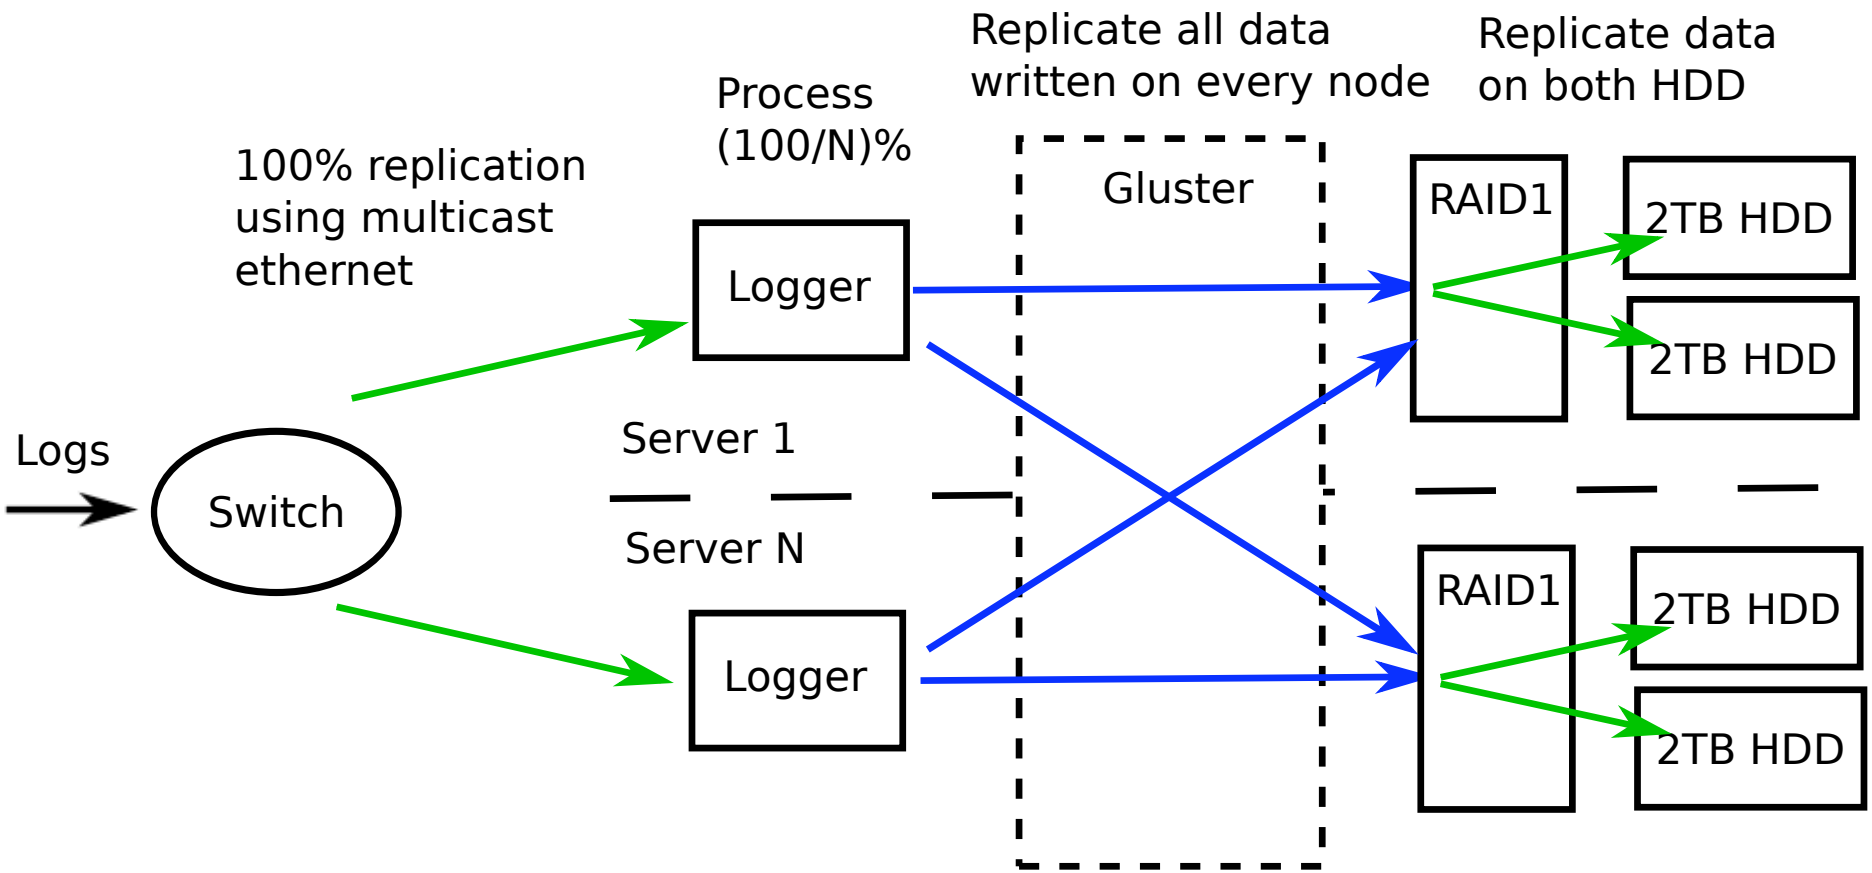

Figure: Third version of the cluster using Pacemaker and Corosync

- Switch: Single point of failure
- Cheap storage solution

Log clients

- Syslog over UDP for most of the sources
  - Linux systems and most of the application
  - Windows Event logs using *Snare*
  - Files using the rsyslog module *imfile* or *Epilog* on Windows; e.g. PVSS.
  - Network devices
- Legacy custom log protocol for Data Acquisition software
  - Over TCP/IP
  - Cannot be distributed over the cluster

Outline

OSSEC

- Host based Intrusion Detection System
- Limited use yet: log analysis
  - Log digest
  - Alerts
  - Active responses
- Hierarchical rule set to improve analysis time
  - Attack signature
  - Problems in data acquisition
    - ★ Router Line card crash -> Get all information for the tech support
- Running on the exported file system
  - Single node

OSSEC rules

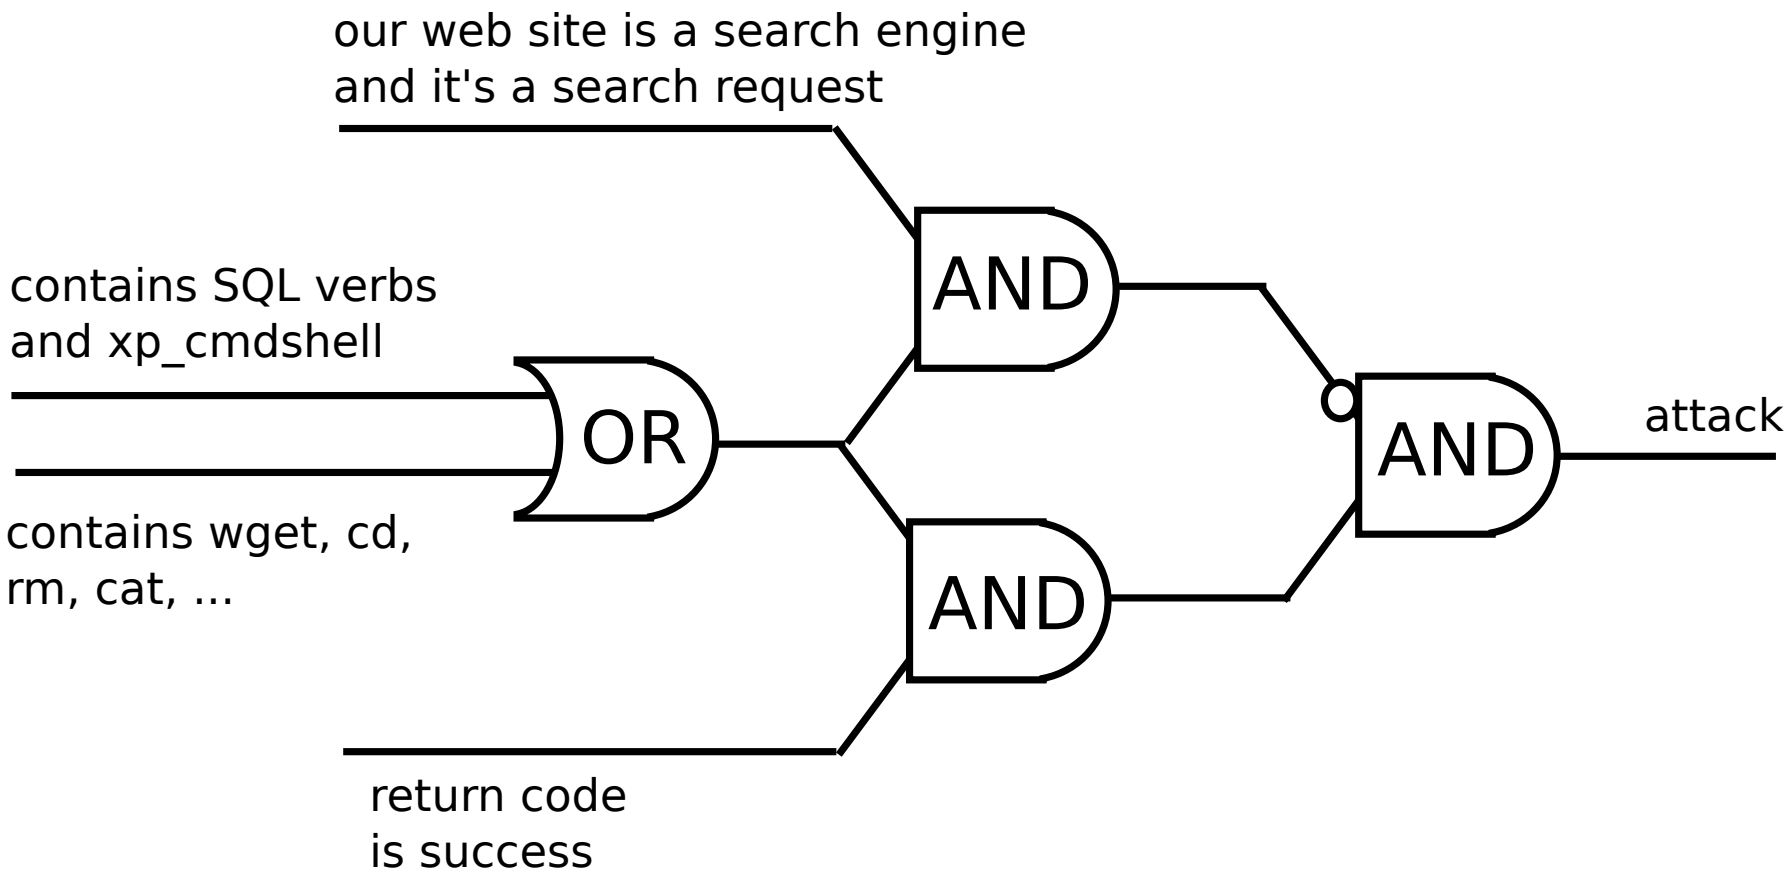

Figure: web log analysis

Outline

Network usage

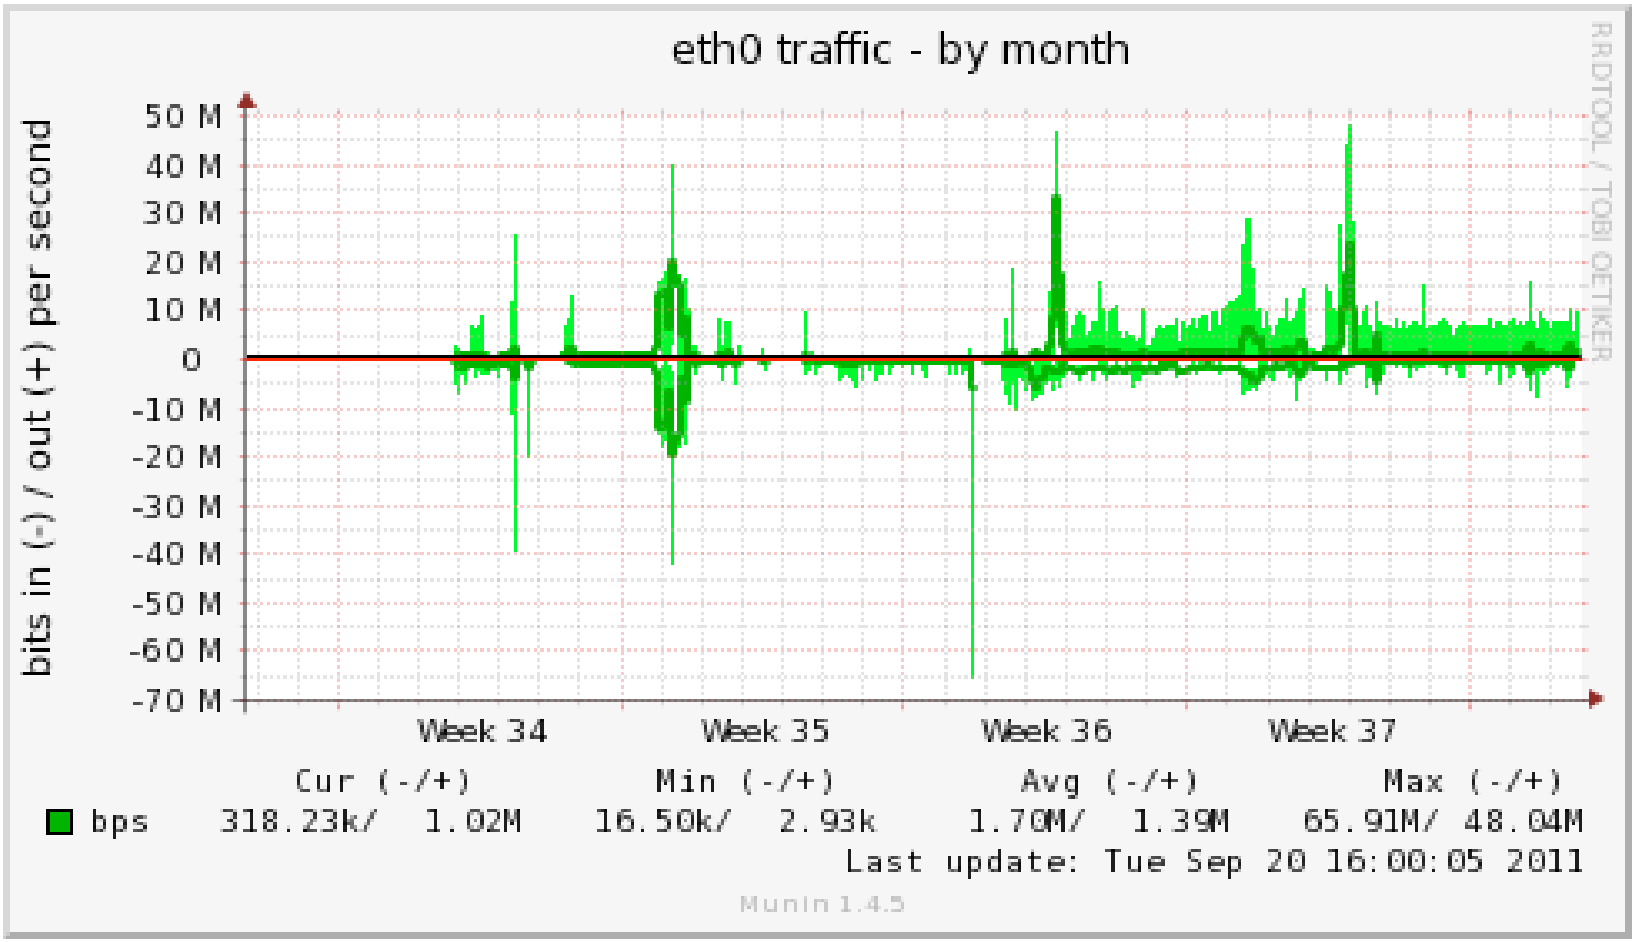

Figure: For one node serving NFS in parallel

Split brain recovery

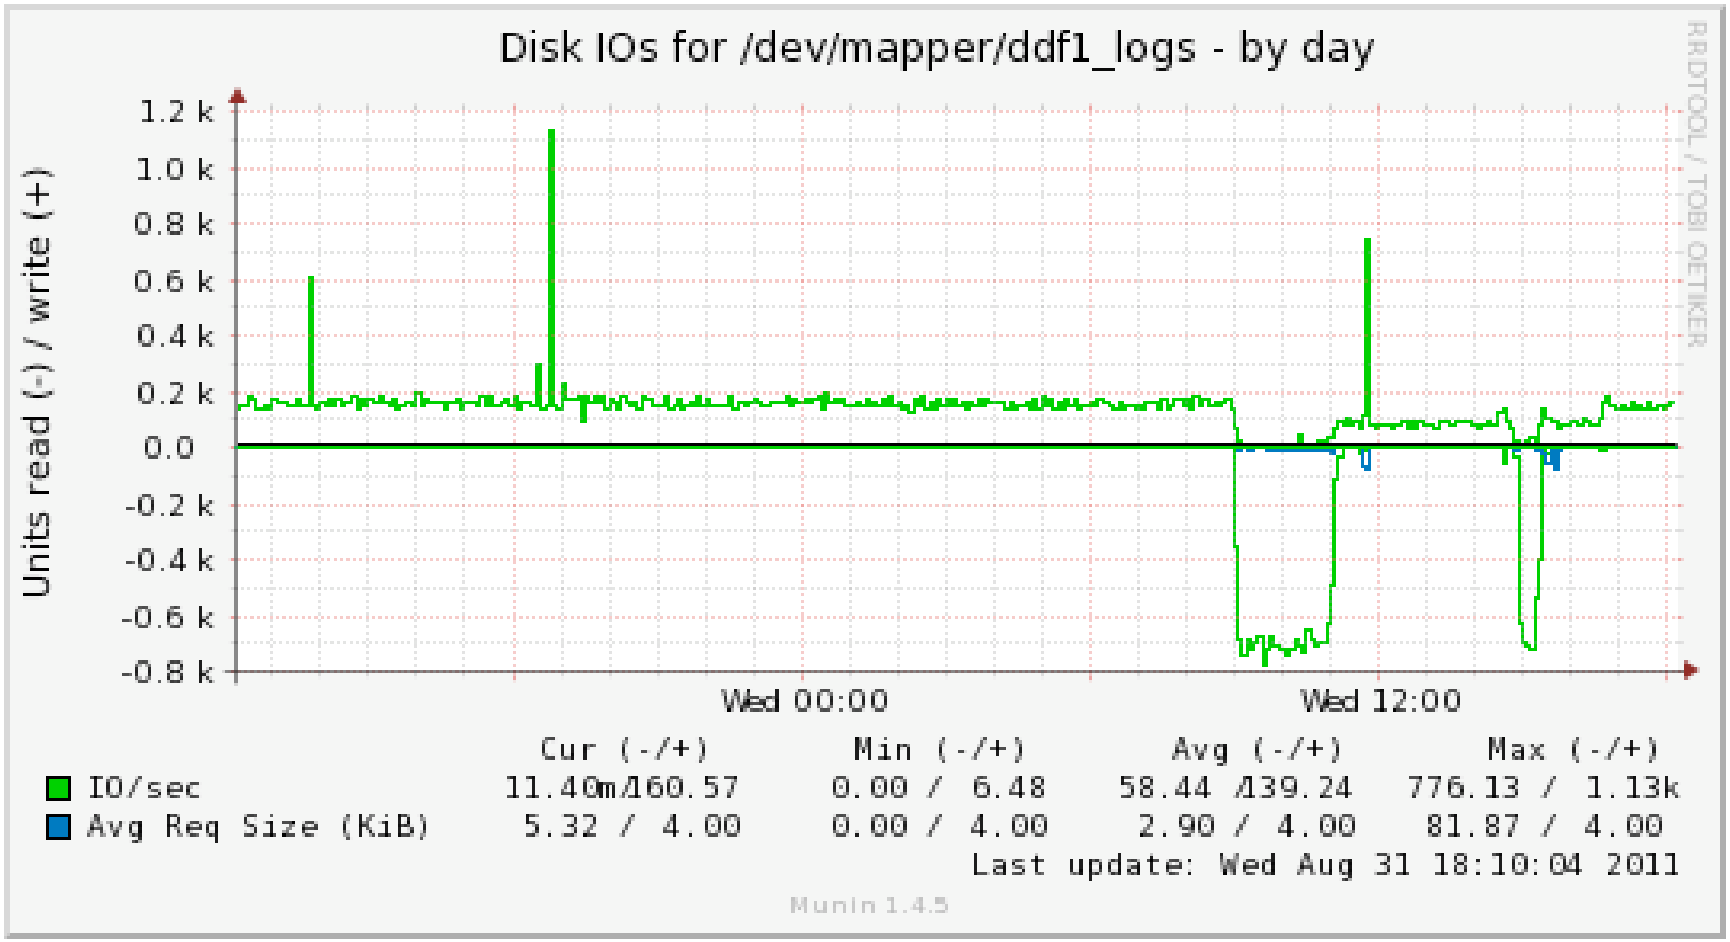

Figure: 2 hour recovery, all IO spent reading the disk

Conclusion

- Third version of the system in production for about one month
- Final version stores currently 400 GB of logs
- Alarms about most of important issues in our system, good complement to Icinga
- Work in progress
  - Issue with rsyslog multi-threads multi-queues, take the time to contribute
  - Implement more analysis and active responses in OSSEC


Thank You
Questions?