FACULTY
OF INFORMATICS
Masaryk University

Effective computer infrastructure monitoring

For Cental Management Service and beyond

Lukáš Daubner

Bedřichov, 2018

Motivation

- If you can't measure it, you can't manage it
  - Is the study room overused or underused?
  - Where is the bottleneck?
  - Do we need to buy new a new storage/CPU/memory?
- Security
  - Are there any threats?
  - Are our data safe?
- Troubleshooting
  - What happened and where?
  - Which systems are affected?
- Higher-level management decisions
  - Prioritizing
  - Annual reports

Environment

- Central Management Service (CMS)
- Large computer infrastructure
  - ~3000 Workstations
  - ~250 Servers
- Initial monitoring state
  - Nagios
  - Custom task-specific tools
  - Distributed logs
  - Some parts aren’t even checked at all

Task

- Centralized monitoring solution
- Data for supporting services
- Scalable and extensible solution
- Basis for data analysis
- Provide data for CSIRT

Task - Use Cases

- Monitor service failures and behavior
- Resource utilization monitoring for servers
- Monitoring of PC study rooms
- Active Directory auditing
- Application performance
- Ease the problem tracing

Scope

- New monitoring model
- Based on model by J. Spring
- Incorporates specifics of Central Management Service

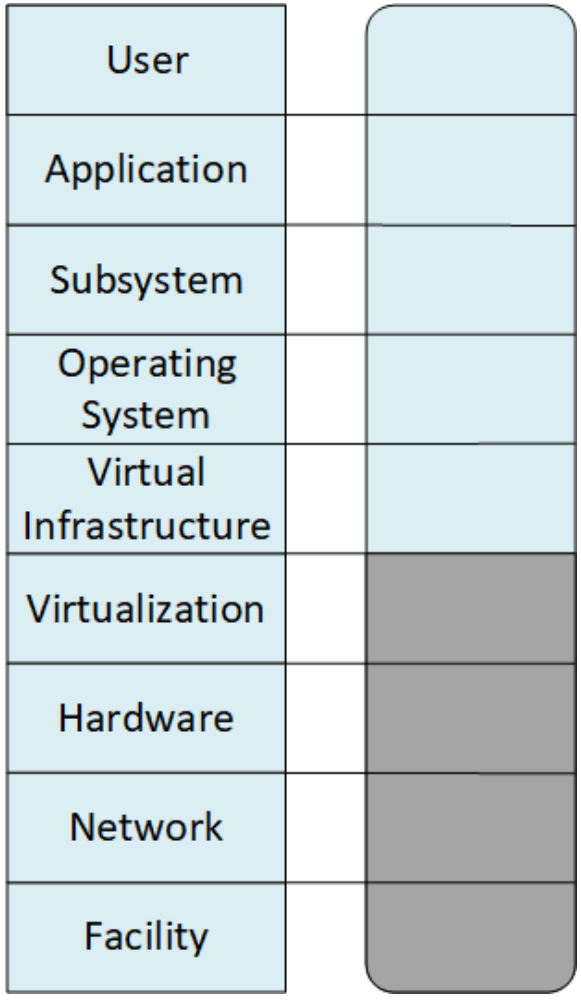

Monitoring through logs

- Logging is common practice
- Logs contain detailed information
- Metric-oriented monitoring can be reduced to log-based
- Logs are Big
  - Volume
  - Variety
  - Velocity
  - Veracity

Elastic Stack

- Advantages
  - Well-known toolset
  - Open source
  - Strong community and support
  - Large gallery of modules and extensions
  - Rather simple horizontal scalability
- Disadvantages
  - DIY approach
  - Elasticsearch is primarily a search engine, not data store

Architecture

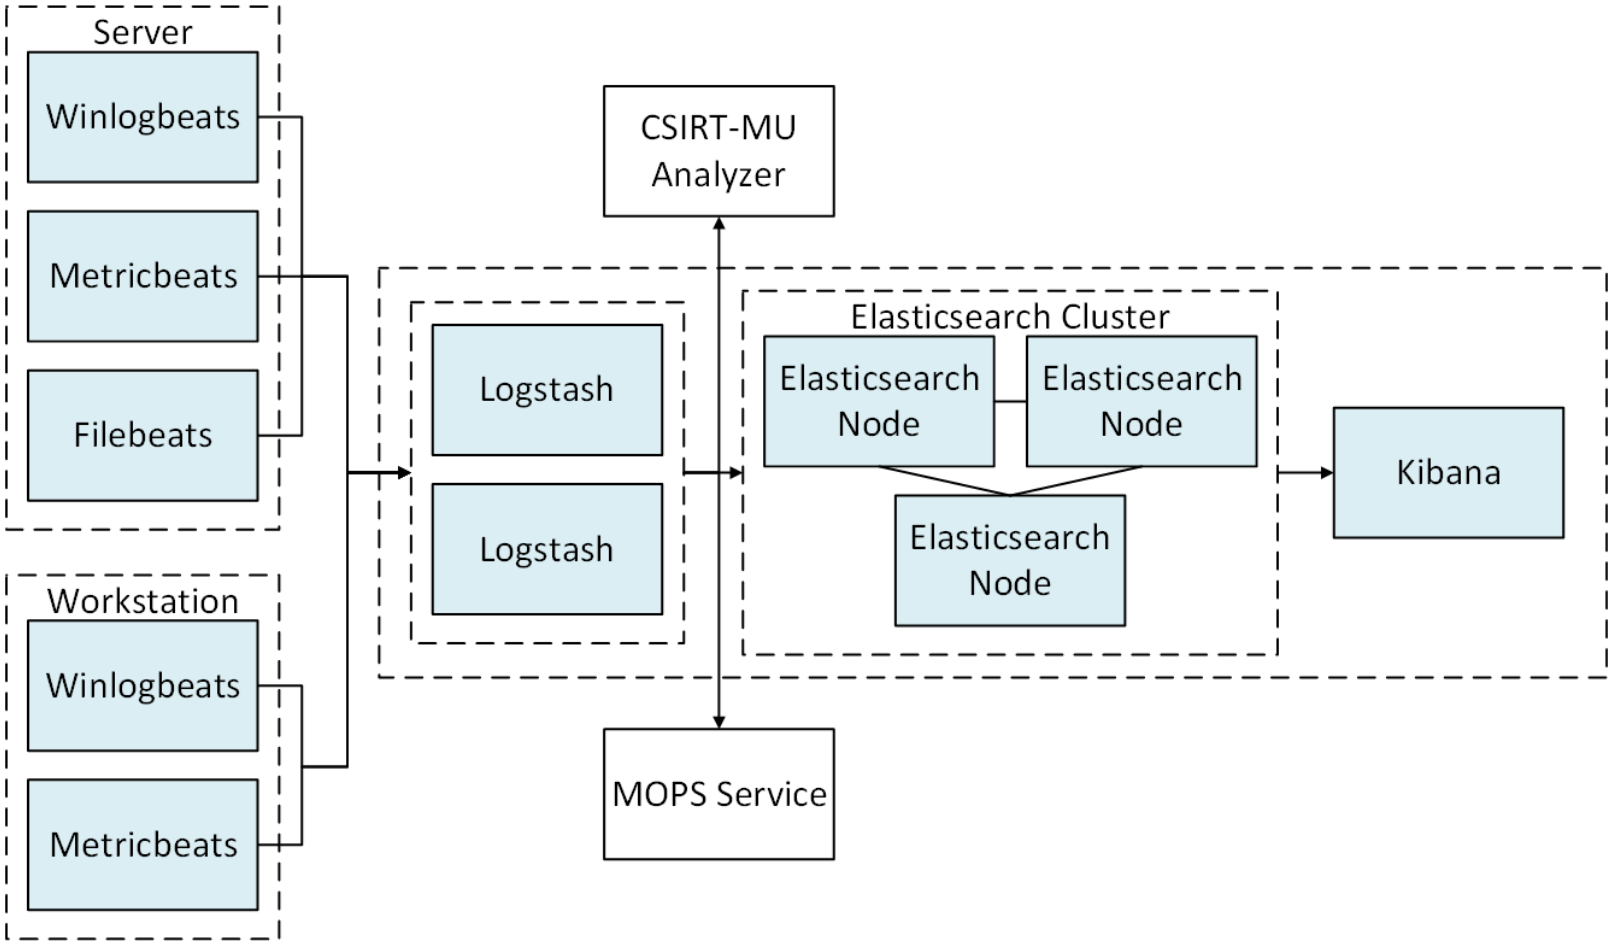

Next Steps

- Deploy collectors to the whole infrastructure
  - Choosing the right logs
  - Knowing the context
  - Maintenance
- Taxonomy of logs
  - How do logs translate to events?
  - Which logs are useful?
- Analysis
  - Finding the incidents
  - Provide problem tracing and explanation