CHEP 2015

Analysis of CERN Computing Infrastructure and Monitoring Data

Christian Nieke, CERN IT / Technische Universität Braunschweig

On behalf of the CERN IT Analytics Working Group

IT Analytics Working Group

- Goals:
  - Coordinate analysis and trending of application/service usage data
    - E.g. batch computing, data storage, network…
  - At different stages of maturity
    - Getting a quantitative understanding of a service (exploratory)
    - Informing strategy or planning decisions (hypothesis check)
    - Developing & validating predictive models

Data Sources - Before

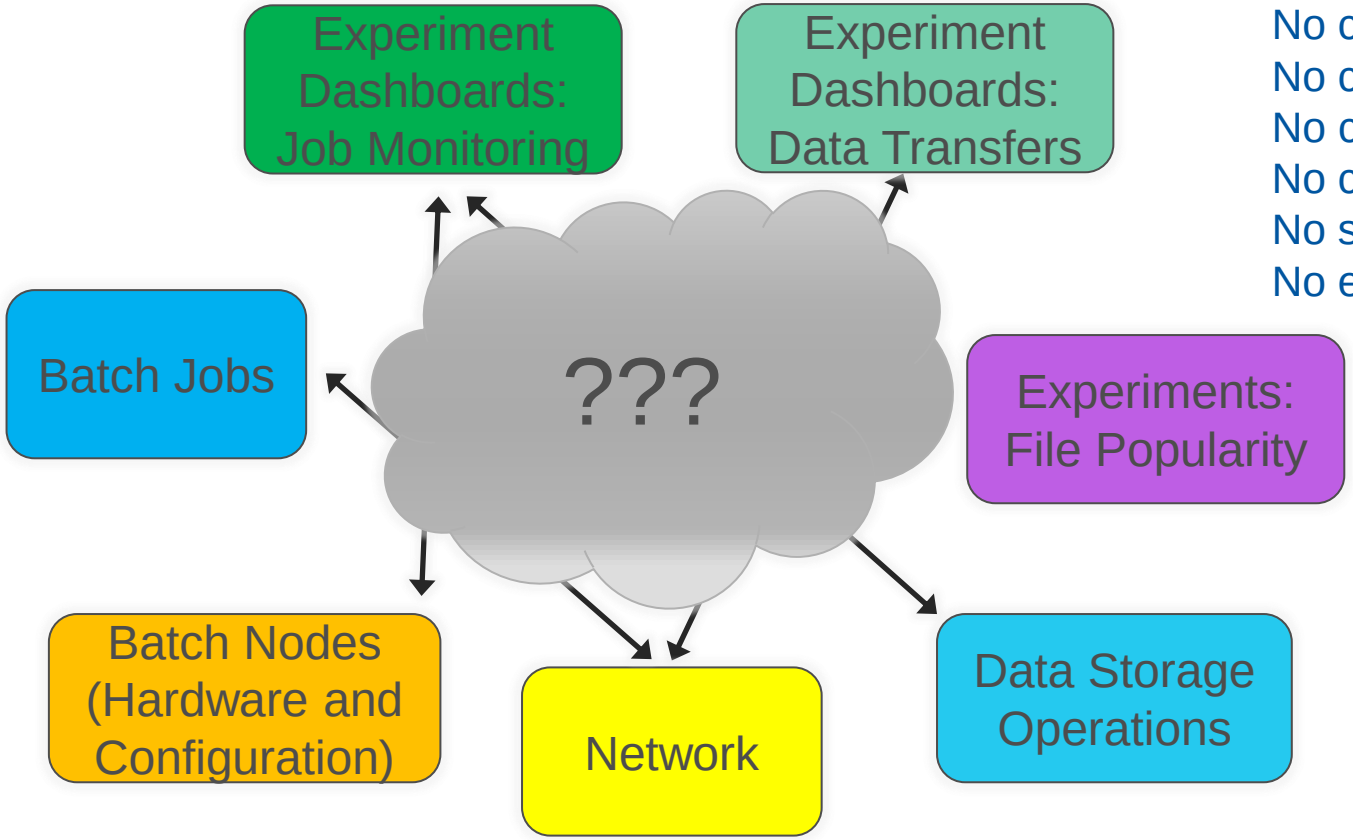

No common analysis goal
No common schema
No common format
No common repository
No shared documentation
No easy way of joining

Getting the Big Picture

- Combined Activity
  - Enable integrated studies crossing single data source / service boundaries
  - Using a common base repository of prepared input data
  - Provide an exchange forum for discussion on analysis methods, tools and result validation

Common Repository

- Data Warehouse
  - Write once, read many
- Hadoop cluster
  - Raw files in any format
  - Using Hadoop jobs for cleaning and pre-processing
  - Export in CSV, Avro, Parquet, … for Analysis

Data Source Documentation

- Example: EOS file system operations

Processed Parameters

Name	Type	Link	Description	Purpose	Remarks
td	String	-	client trace identifier (<username>.<pid>@<client-host>)	This identifier can be used to aggregate sessions (of the same user, with the same job on the same machine)	
path	String	-	Full namespace path to the file	The folder structure in the path could be usefull to match the file to a project/user/programm etc.	
ruid	Int	-	mapped unix user id	Allows aggregation by user	ruid==1 means root process, which is usually an internal process like rebalancing, draining etc.
rgid	Int	-	mapped unix group id	Allows aggregation by group	
host	String	LanDB : name	Name of the disk server serving the file	Allows aggregation by user host and links to LanDB	
fid	Int	-	EOS file id	Allows aggregation by file	unique per EOS instance (e.g. eosatlas, eoscms,...)
fsid		-	EOS file system id, e.g. disk	Allows aggregation of files by file system (disk)	
ots	Date	-	File open time as unix timestamp (in seconds since January 1st, 1970 at UTC)	Aggregation by time	

Data Sources - Federation

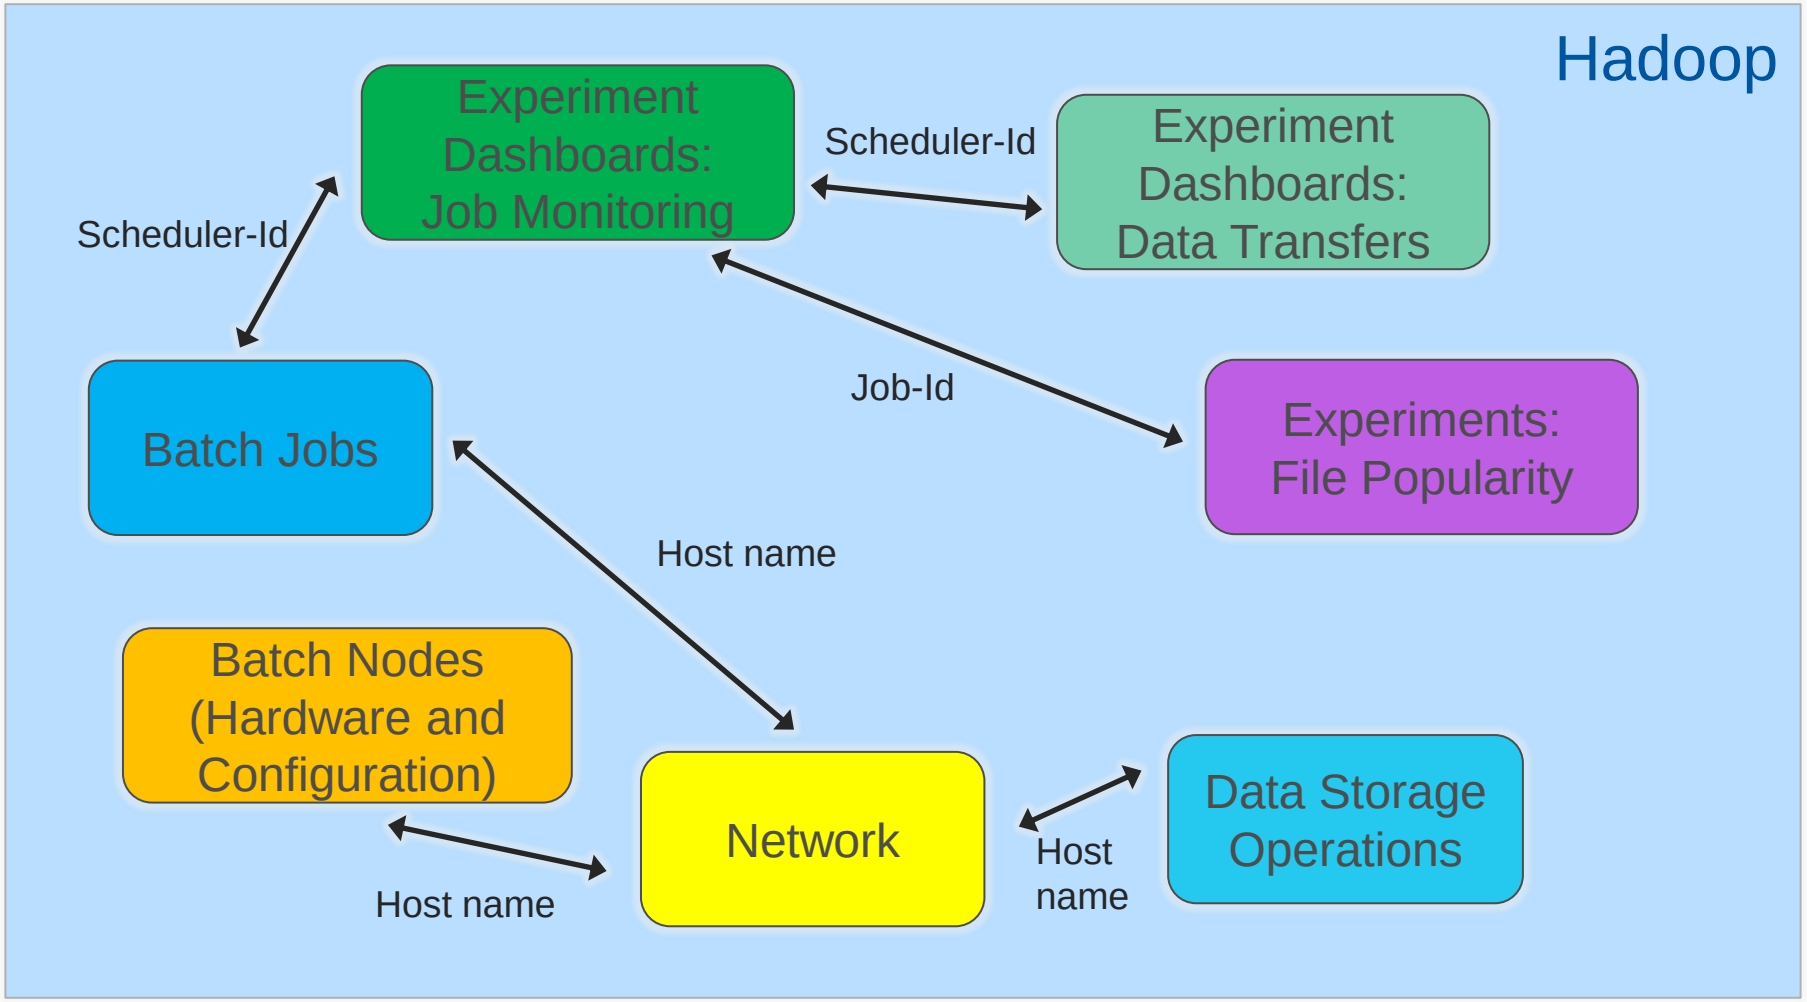

Example Analysis Workflow

- Job Performance: Geneva vs. Budapest
  - Different computing centers
  - Different hardware
    - CPU, Memory, Network, ….
  - Do we get the same performance?
    - Compare CPU time used per job

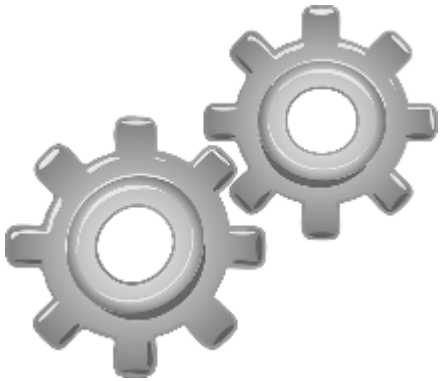

CPU Time and Location

- Based on batch computing logs and network configuration

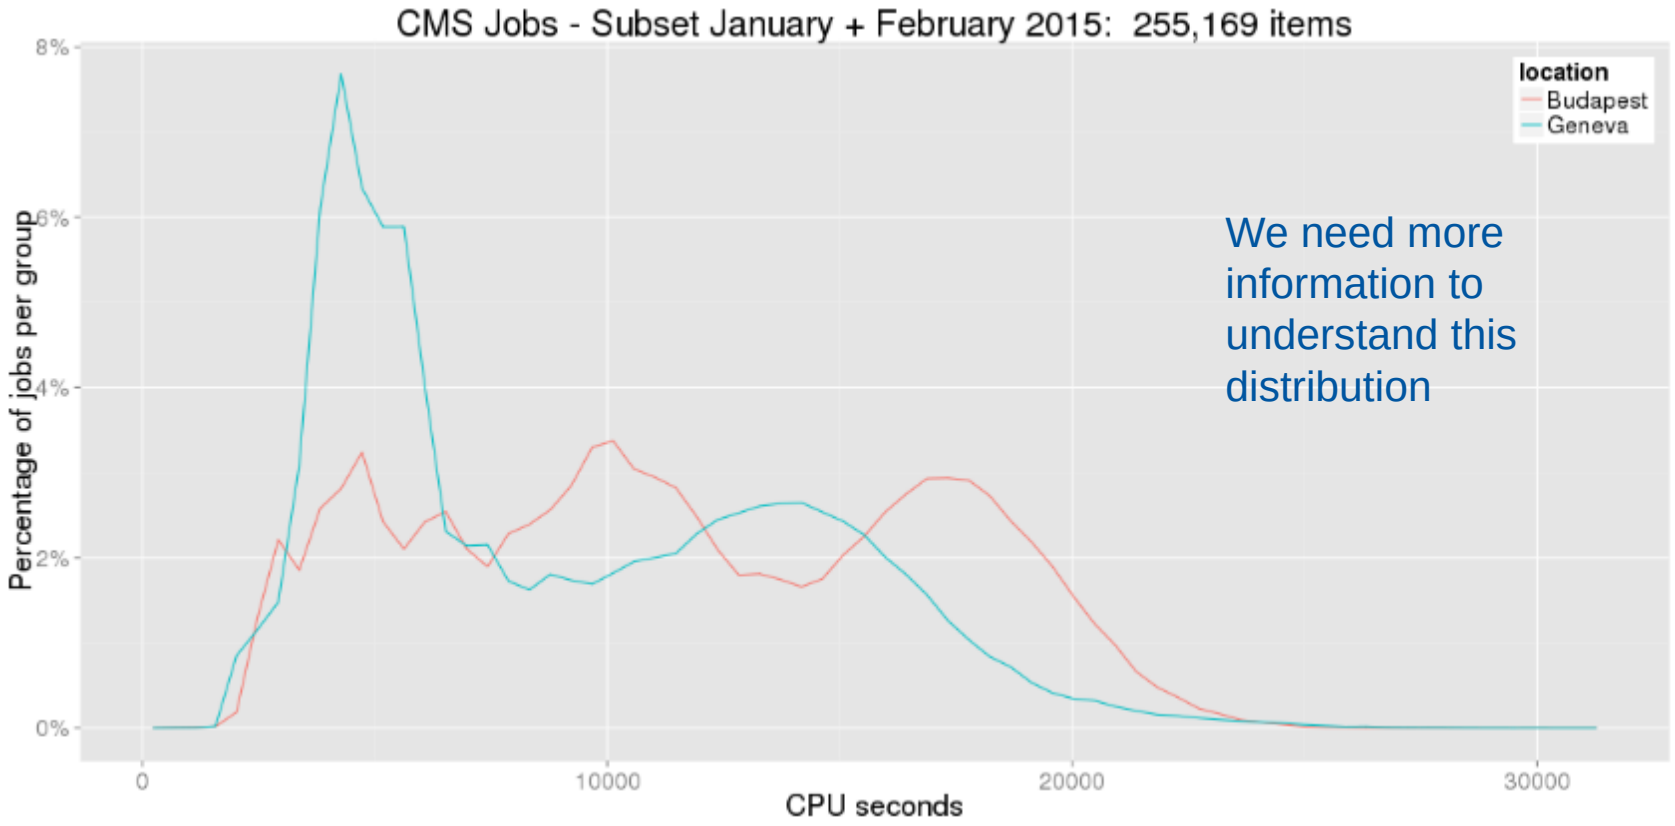

Tasks

- Based on experiment job dashboard

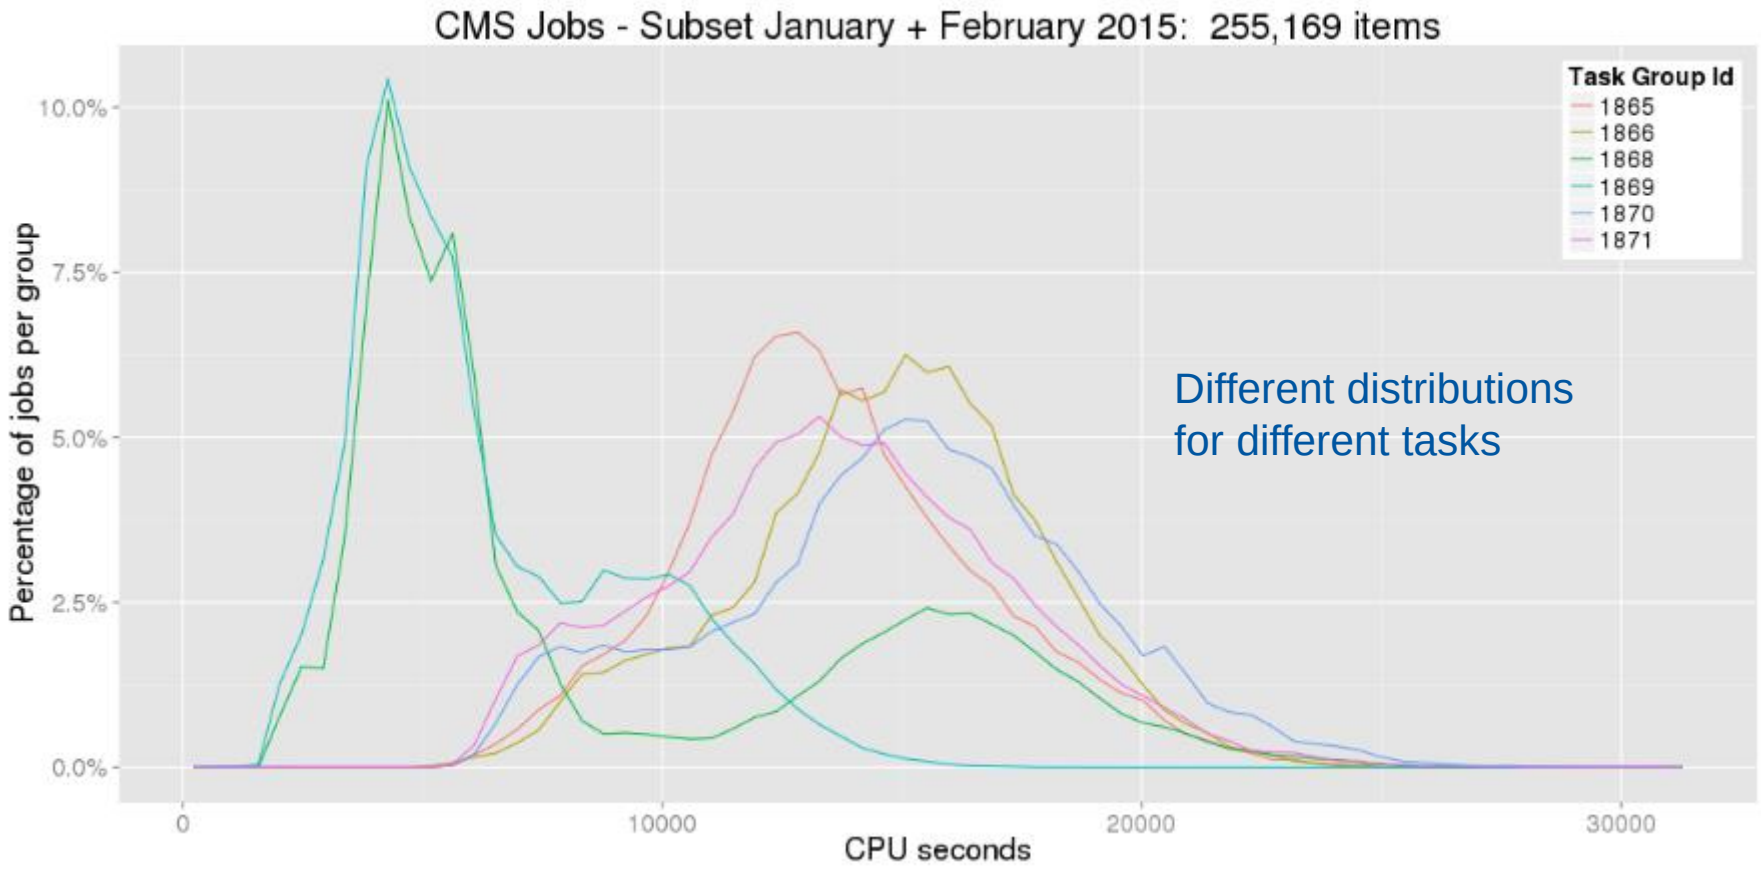

Tasks

- Selecting a single task

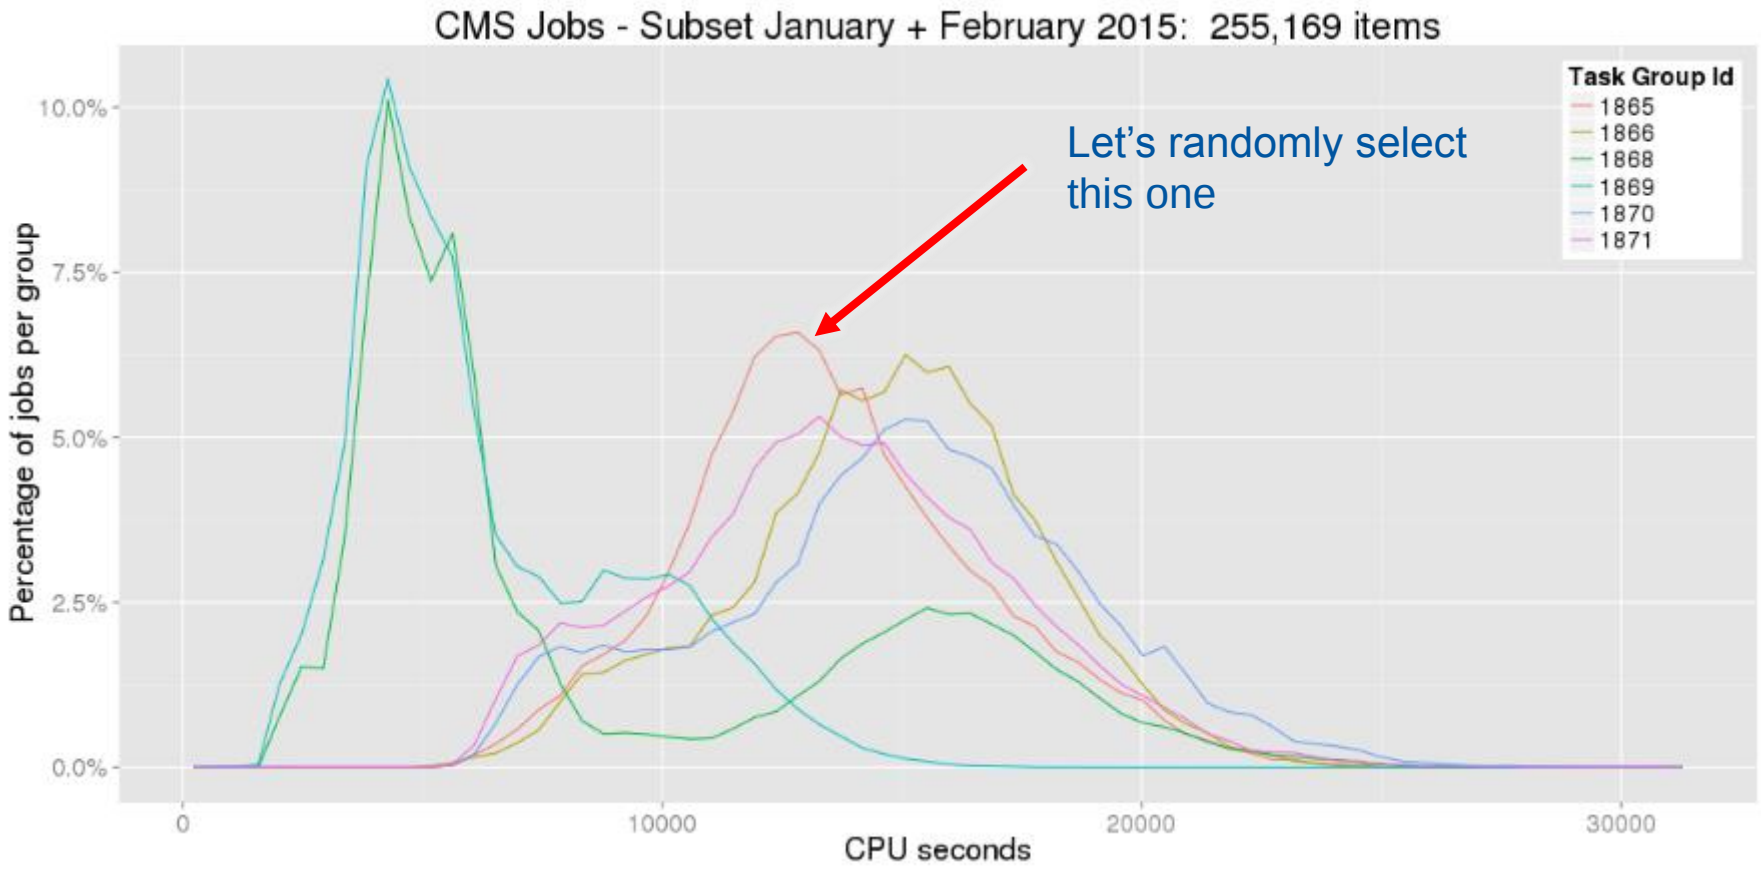

Tasks

- It seems like there are still more underlying effects

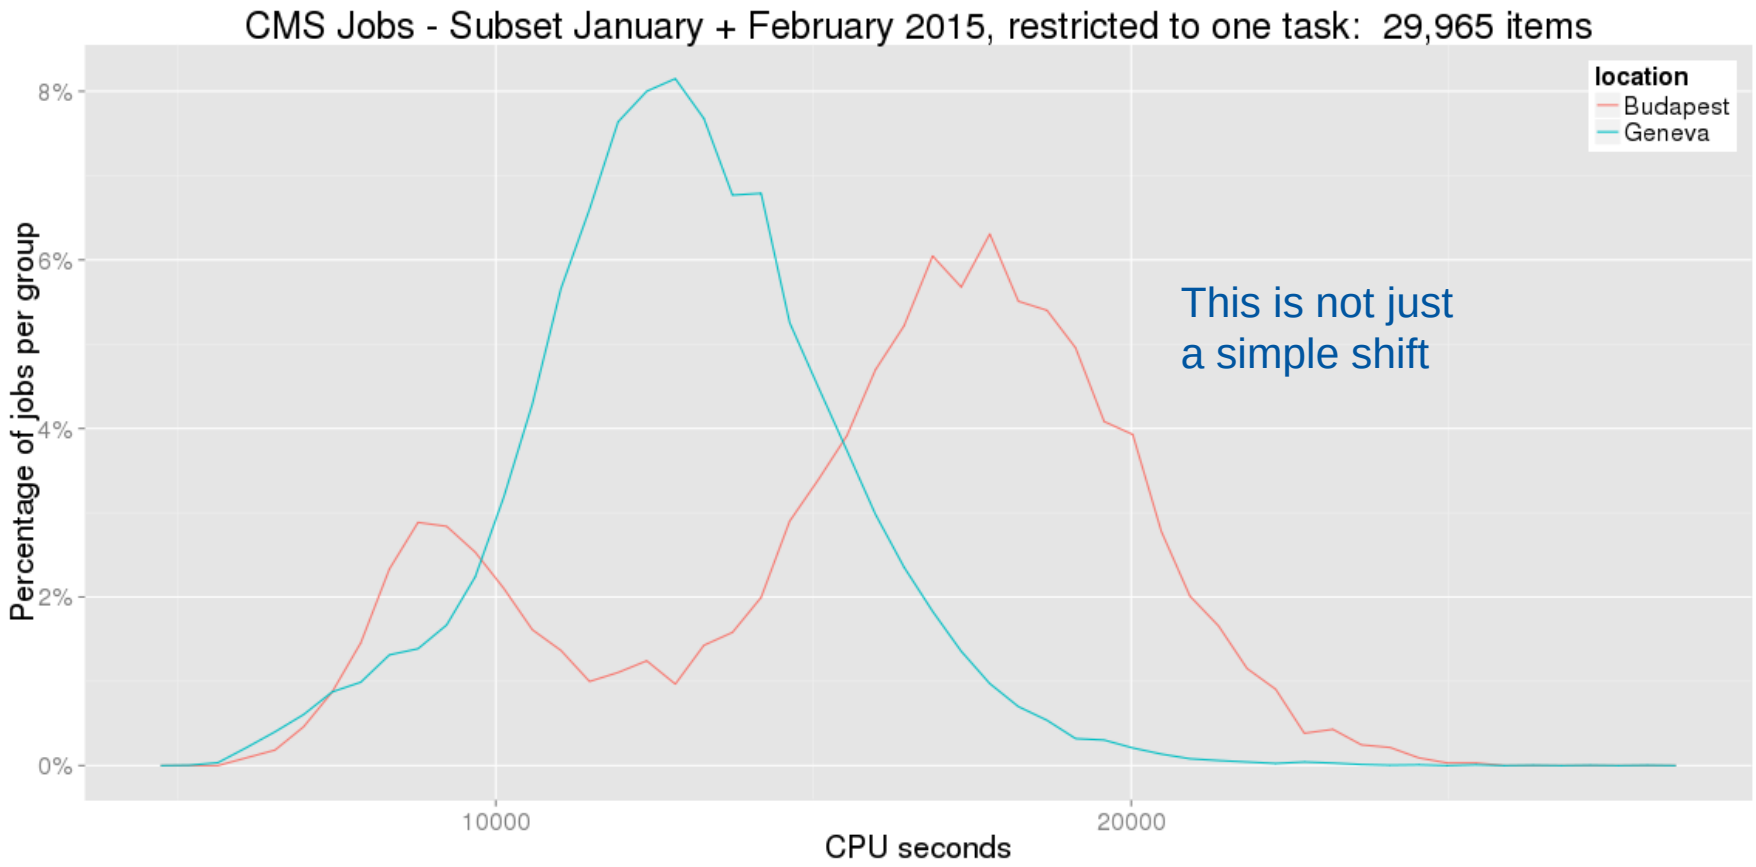

HepSpec Benchmark

- HepSpec Factor based on batch benchmarks

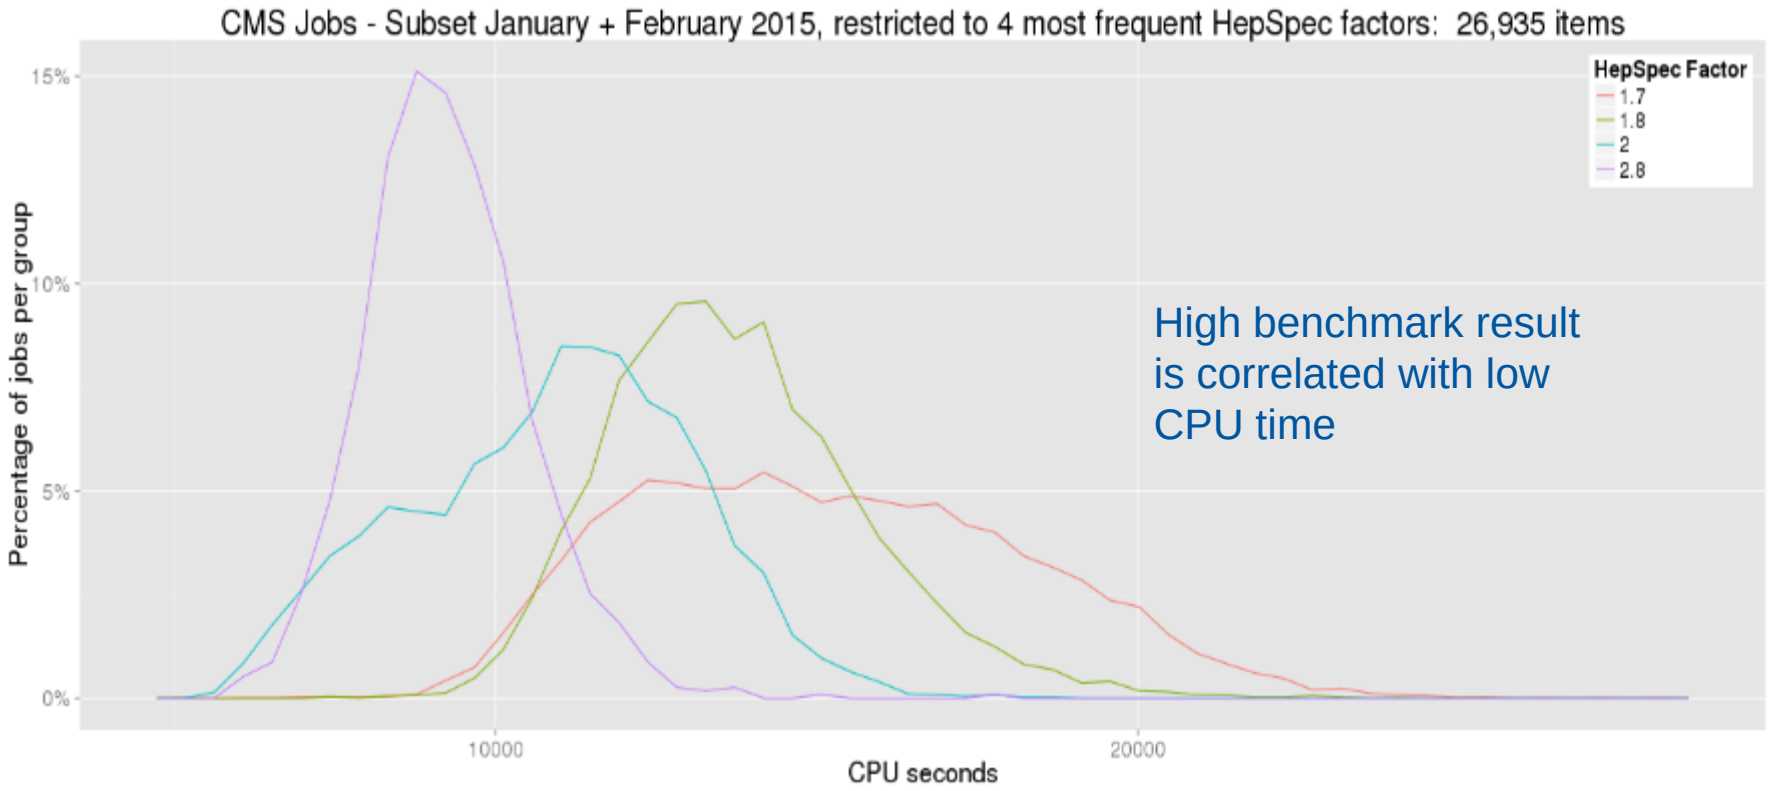

Scaling by CPU Factor

- Removes “expected” deviation

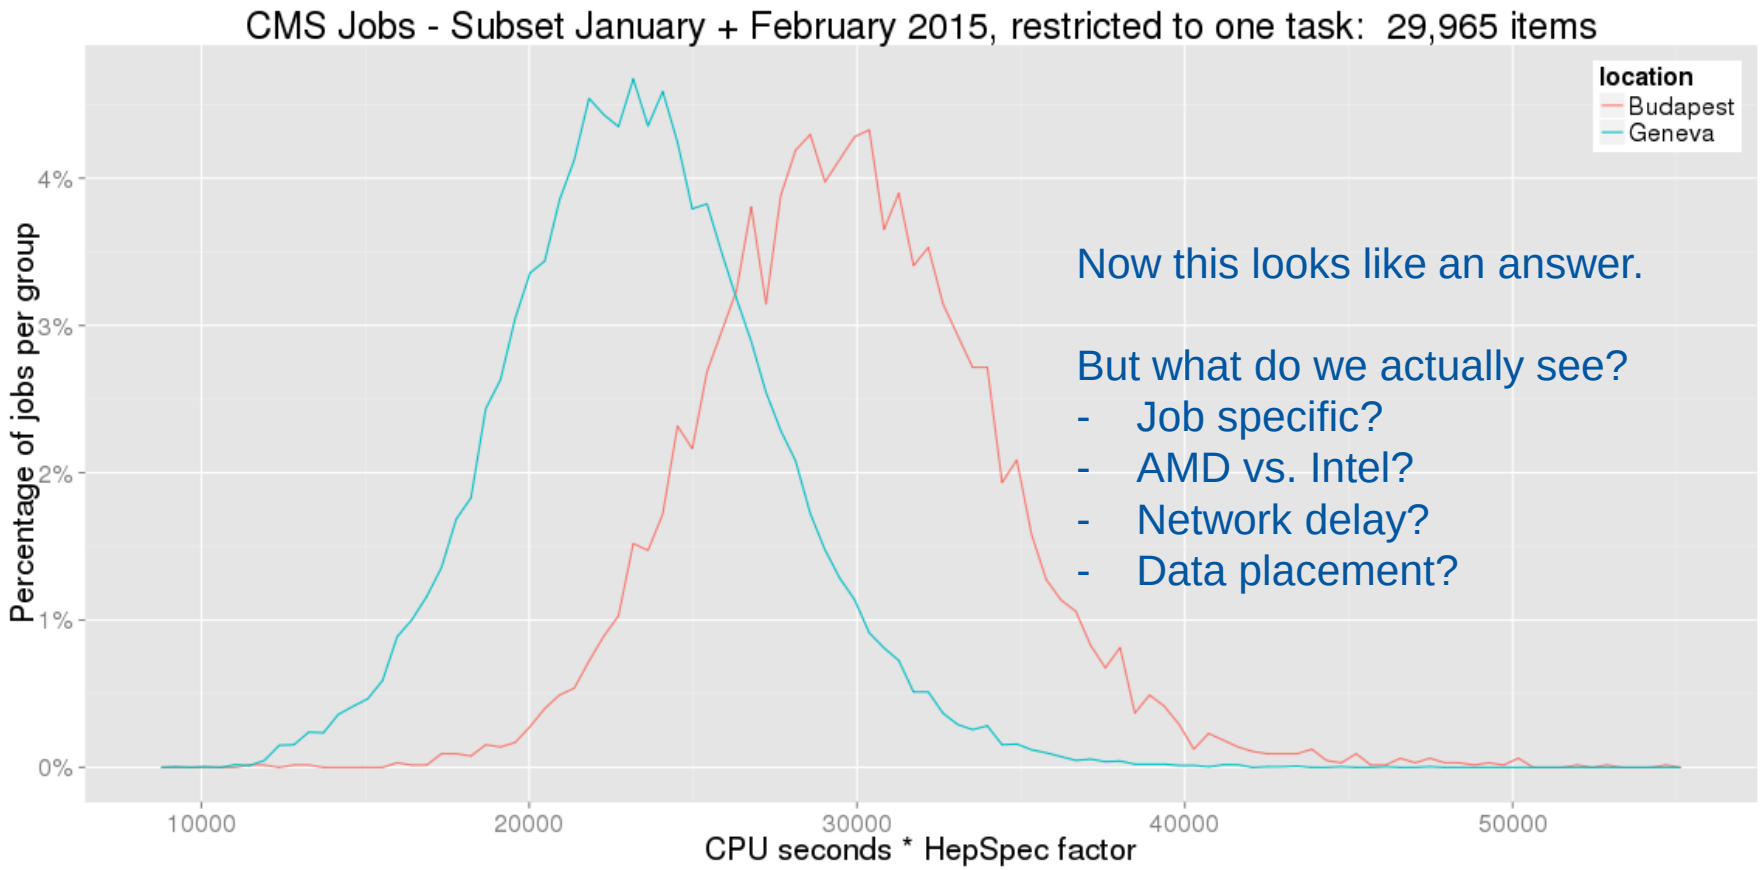

Conclusion

- Combined Effort
  - CERN IT and Experiments
  - Federated data repository for uniform access
  - Understanding the system as a whole
- Examples for Actions Taken
  - Rebalancing batch slots per machine to avoid swapping
  - User notification in case of inefficient jobs
  - Activated TTreeCache for ROOT in ATLAS

Resources

- Twiki
  - https://twiki.cern.ch/twiki/bin/view/ITAnalyticsWorkingGroup/WebHome

- Contact:
  - Dirk Duellmann, CERN IT (Working Group Chair)
  - or myself