

DUBNA

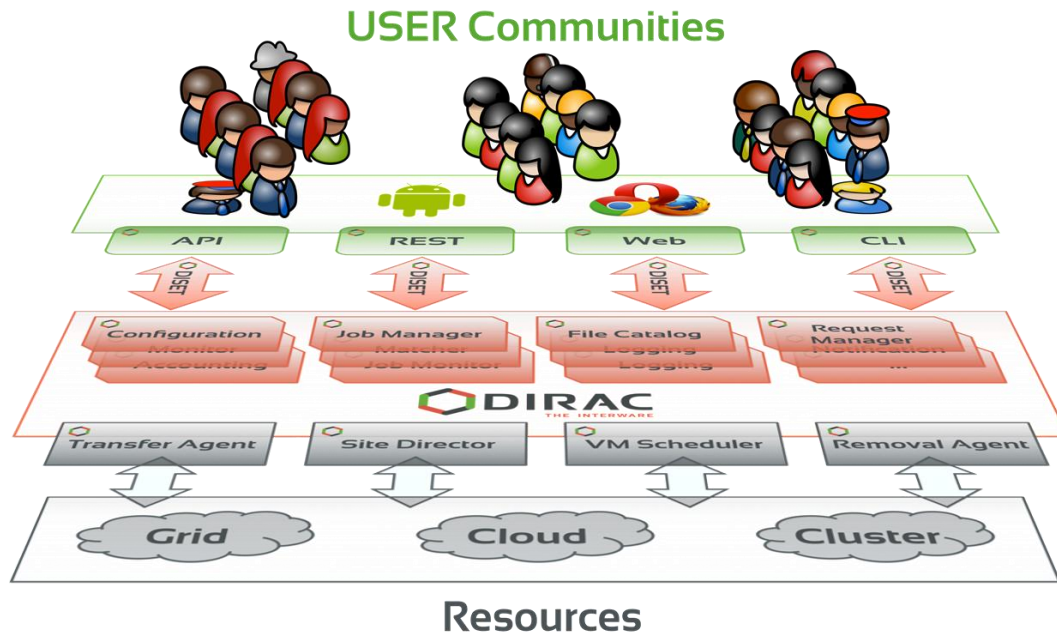
Evaluation
of
JINR computing resources
performance
with
DIRAC Interware

Speaker: Igor Pelevanyuk

Supported by the RFBR special grant ("Megascience – NICA"), №18-02-40101

What is DIRAC?

DIRAC provides all the necessary components to build ad-hoc grid infrastructures **interconnecting** computing resources of different types, allowing **interoperability** and simplifying **interfaces**. This allows to speak about the DIRAC *interware*.



Web

CLI

API

REST

What is DIRAC?



User Interface

API

Central configuration

Workload management

Data management

Integration tools

File Catalog

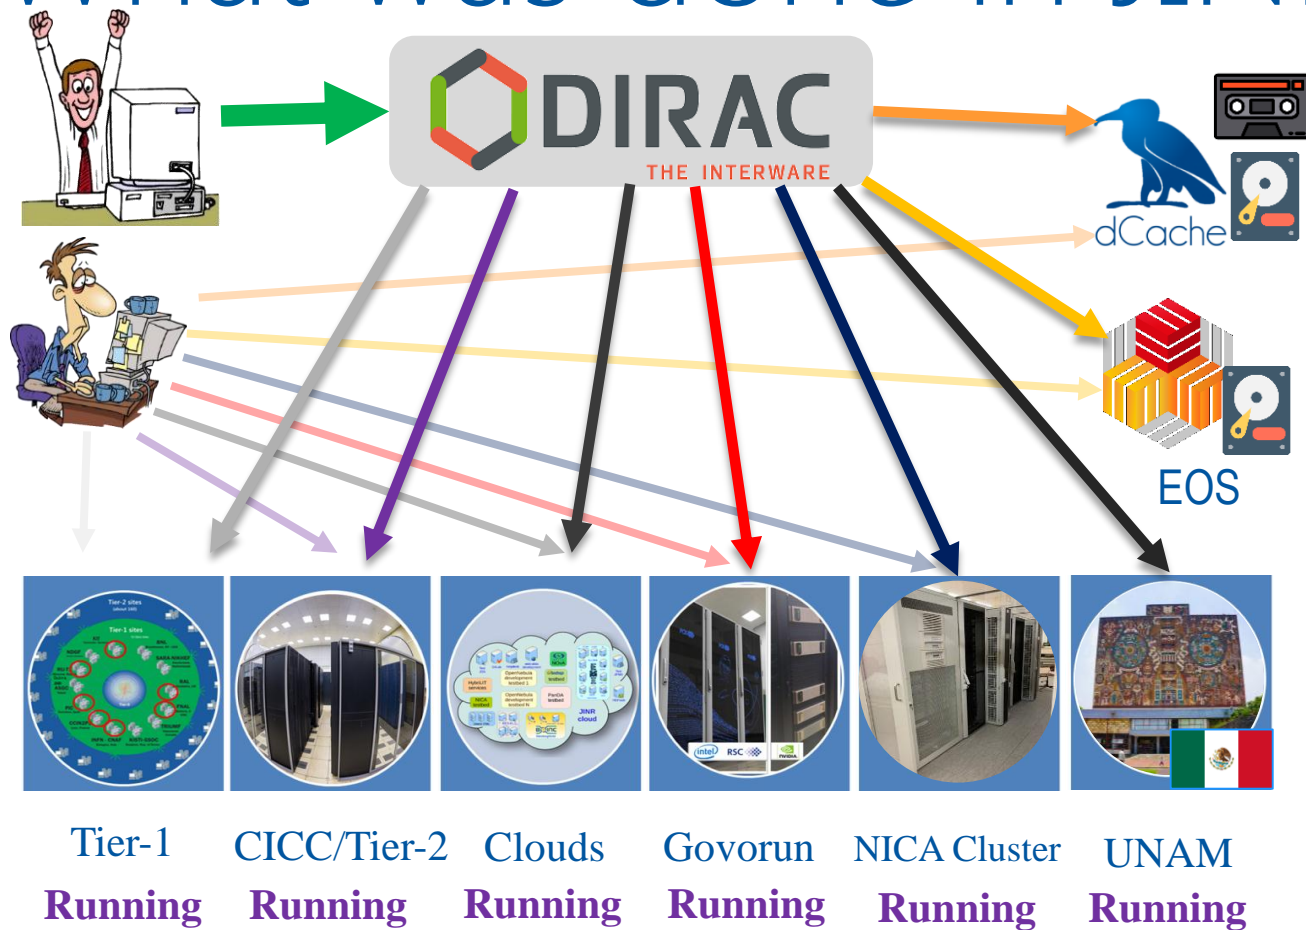
Workflow management

Metadata management

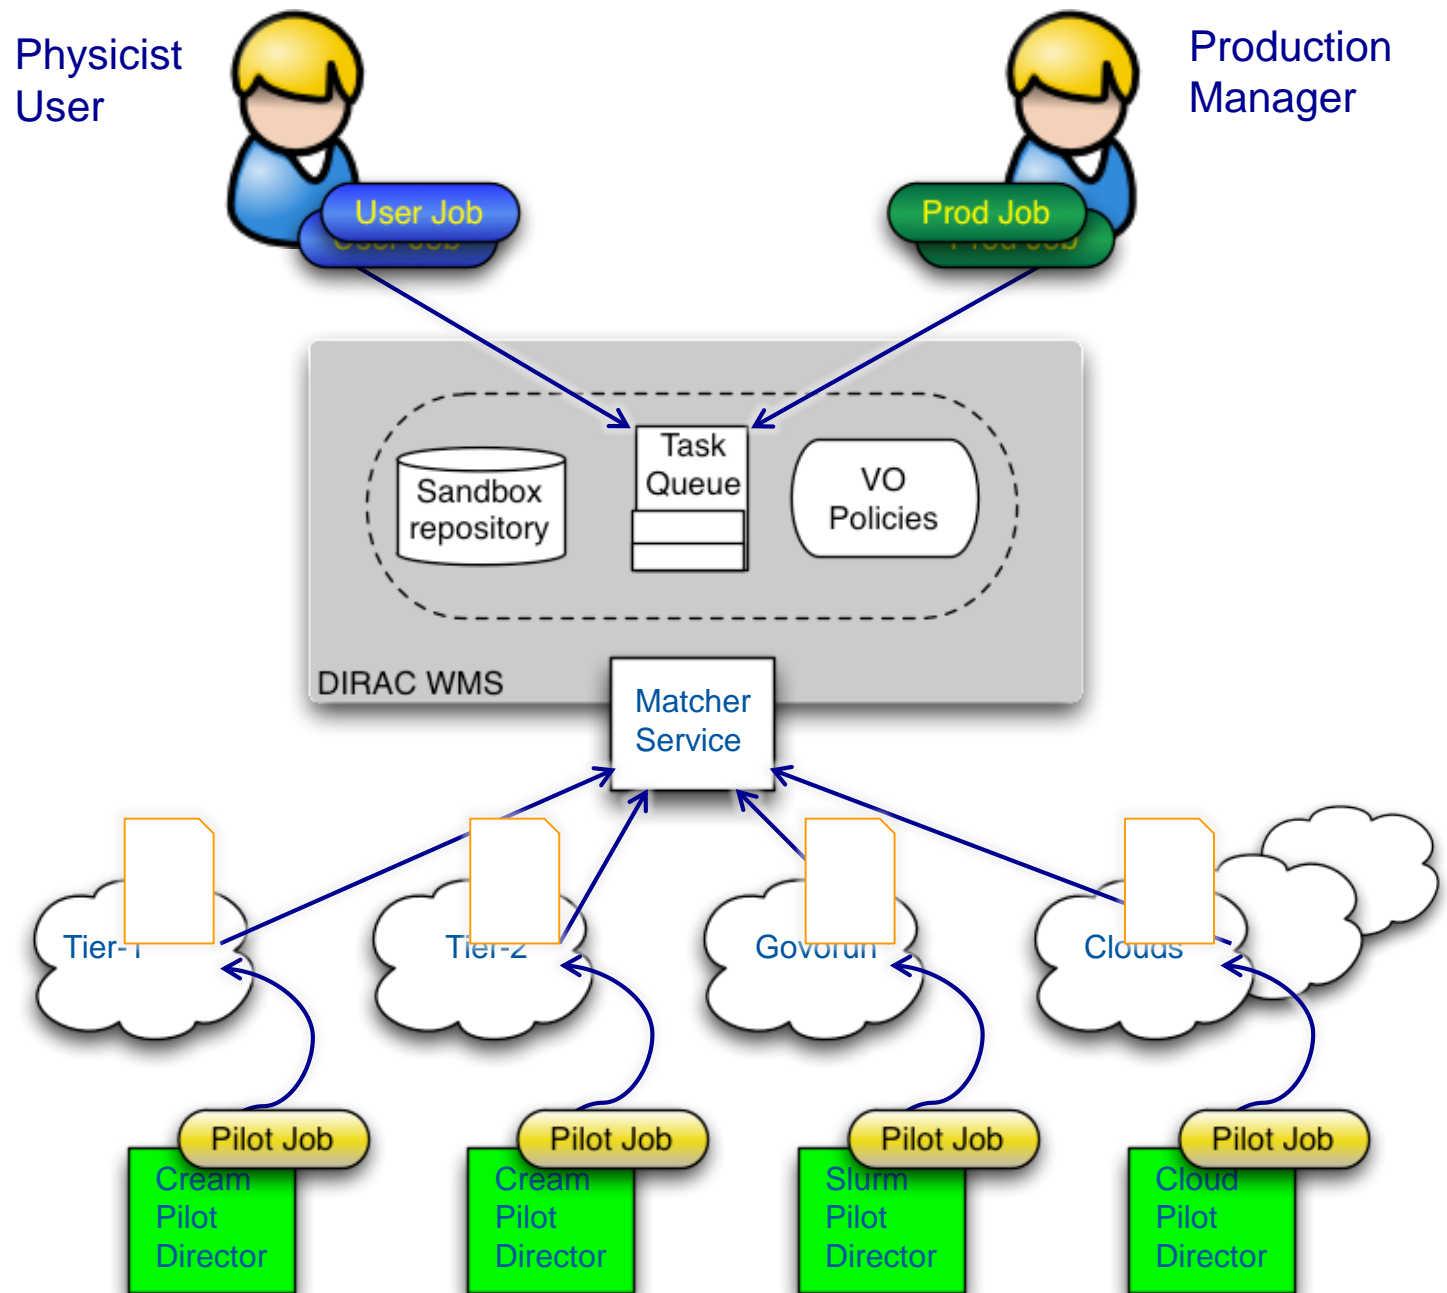
Accounting

Management

What was done in JINR



The computing resources of the JINR Multifunctional Information and Computing Complex, clouds in JINR Member-States, cluster from Mexico University were combined using the DIRAC Interware.



Individual CPU core performance study

- Centralized job management gives possibility for centralized and unified performance study of different computing resources.
- Before running user jobs DIRAC Pilots execute benchmark for CPU core they are running on.
- Benchmark is DiracBenchmark2012 or DB12. It evaluate just CPU core performance. Disk I/O, RAM speed, Network, CPU caches and other highly important aspects of performance are **neglected by DB12**.
- So, we are talking about *passive* estimation of computing performance

Benchmarks

- Benchmark is the act of running a computer program, a set of programs, or other operations, in order to assess the **relative performance** of an object, normally by running a number of standard tests and trials against it.



Benchmark of one aspect



Synthetic test



Benchmarks in grid

- HEP-SPEC2006 is a standard benchmark in computing for high energy physics.

“We propose a new name for the benchmark: HEP-SPEC06. This acknowledges the fact that we use a benchmark derived from the SPEC CPU2006 benchmark suite, but with a clearly defined way of running it (tuned for HEP: on Linux, with gcc compiler, with the optimization switches we defined and in Multiple speed mode), and underlines the difference between it and the other benchmarks of the SPEC family. “[1]

- The DIRAC Benchmark 2012 (DB12) is a good alternative to HEP-SPEC2006 [2].

This benchmark was originally created for prediction of the duration of LHCb Monte-Carlo tasks. It is fast (takes around 60 seconds), and it runs every time the DIRAC pilot job agent starts execution.

[1] Michele Michelotto et al. A Comparison of HEP code with SPEC benchmarks on multi-core worker nodes, Journal of Physics: Conference Series 219 (2010) 052009 doi:10.1088/1742-6596/219/5/052009

[2] P Charpentier, Benchmarking worker nodes using LHCb productions and comparing with HEPspec06, IOP Conf. Series: Journal of Physics: Conf. Series 898 (2017) 082011 doi:10.1088/1742-6596/898/8/082011

Analogy example

Piece of road from point A to B ← A Monte-Carlo task

Speed of the car ← Performance of the computer

Time to complete ← Time to complete

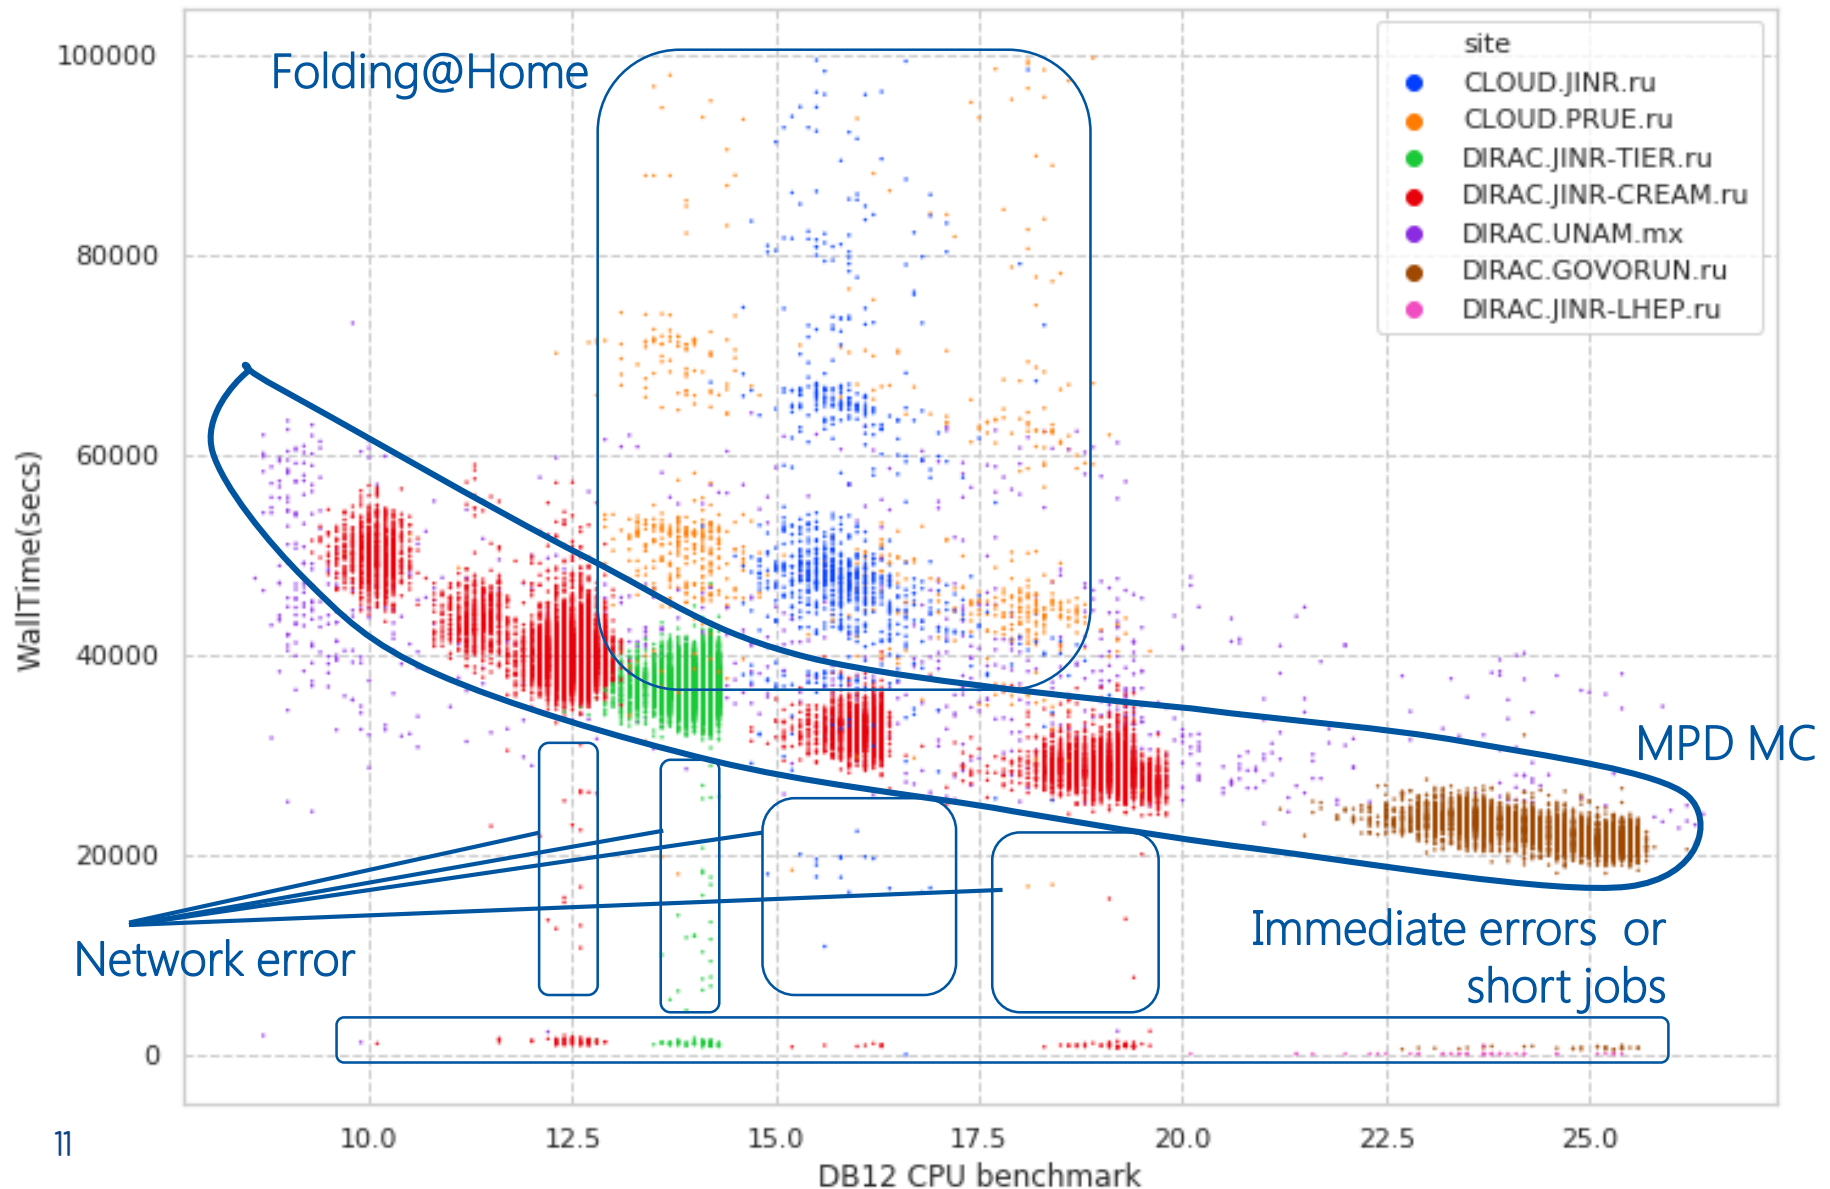
$$Time = \frac{Amount\ of\ work}{Speed\ of\ computer}$$

DB12 gives results like: 10(old slow core), 17 (standard server core), 27 (high performance core)

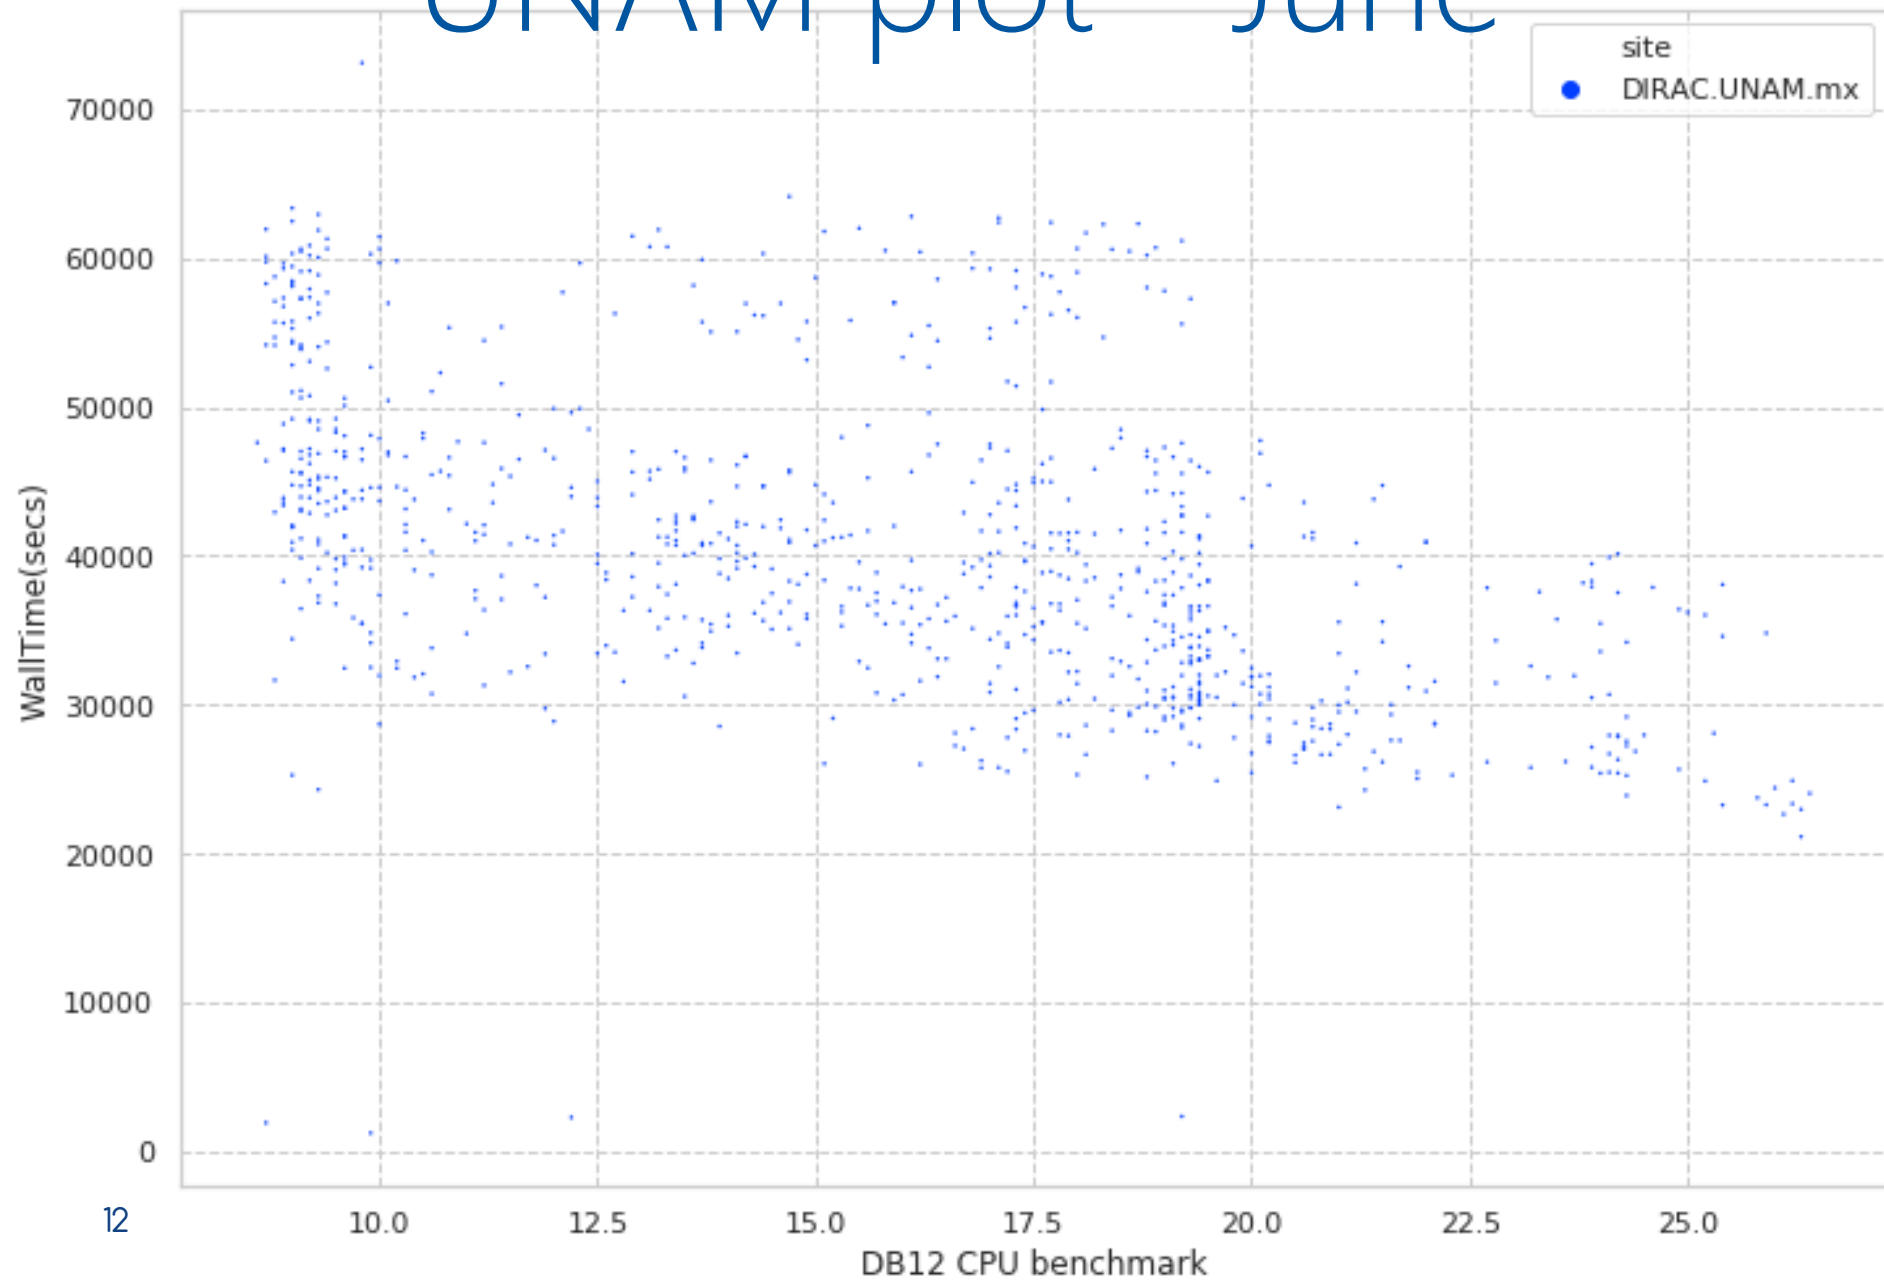
What if we build a plot, where X is DB12 result, Y is time in seconds. Then, every point on the plot represent one job.

It will be useless if all jobs are unique and different. But, in real life there are usually many similar jobs.

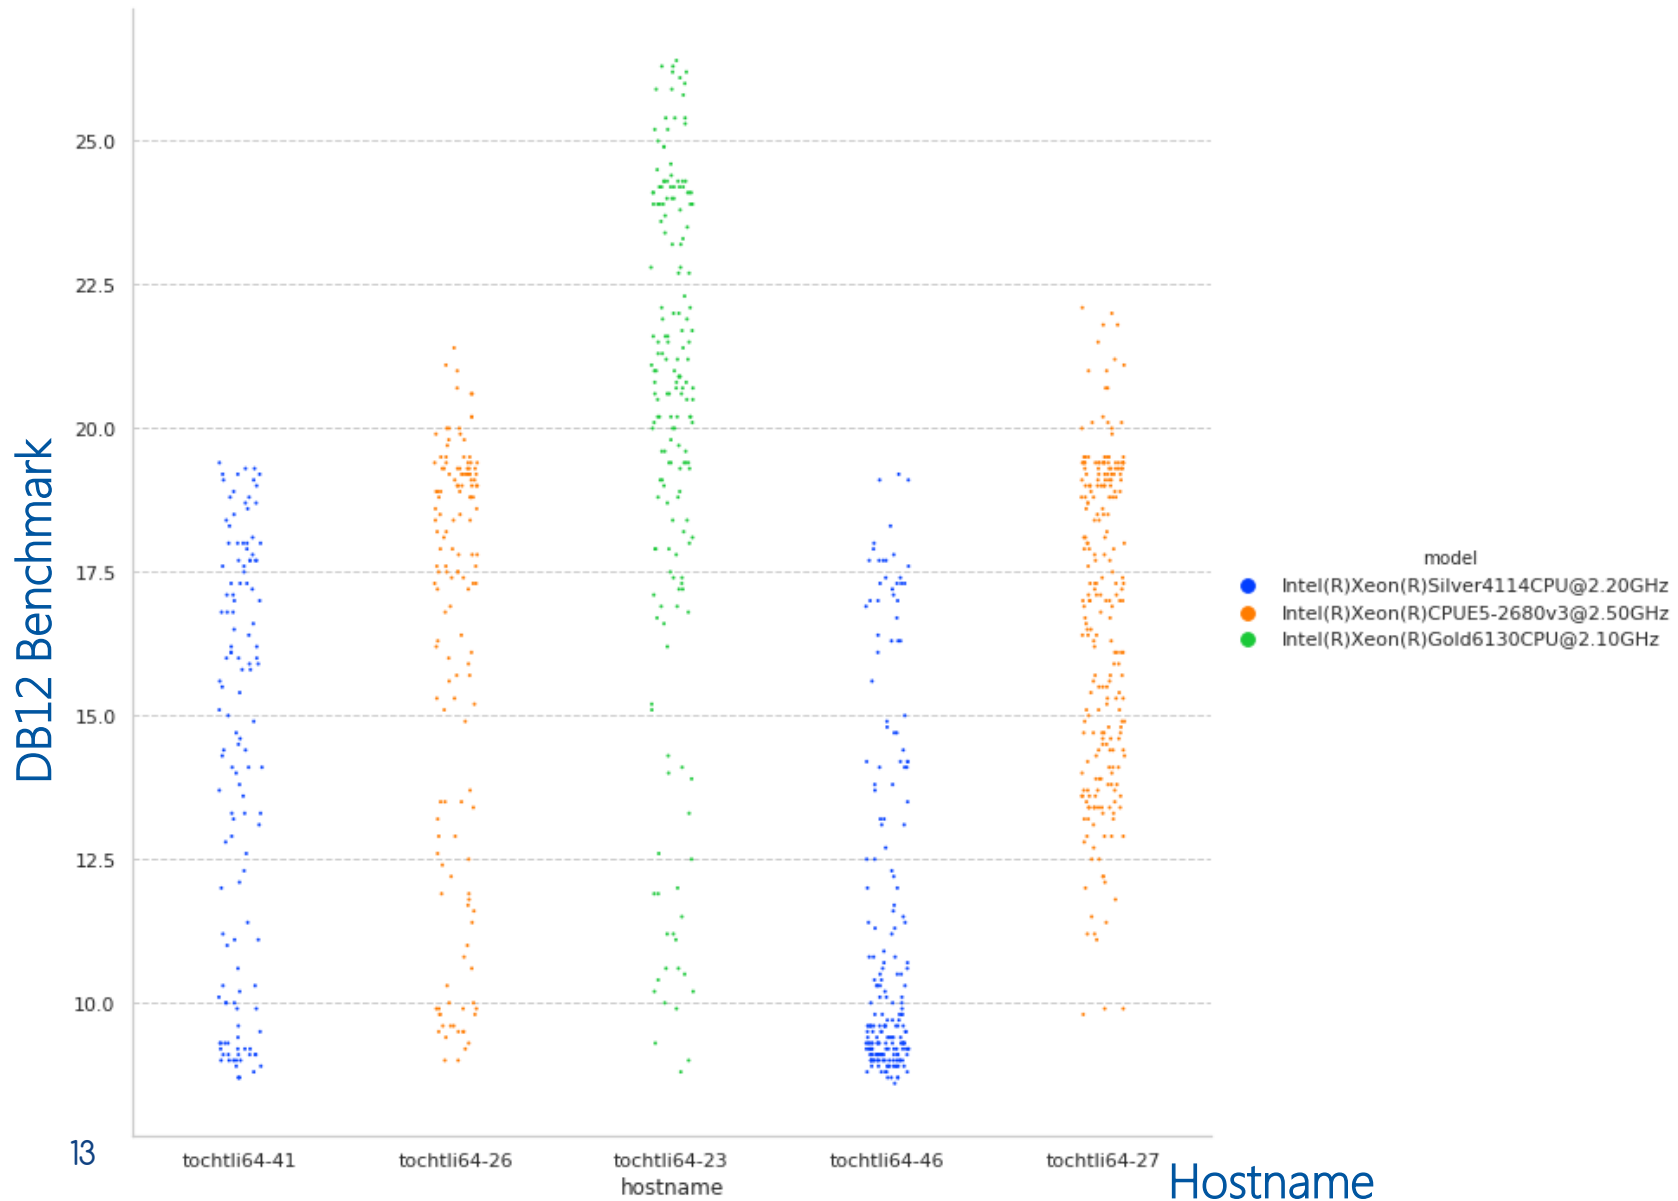
All jobs plot - June



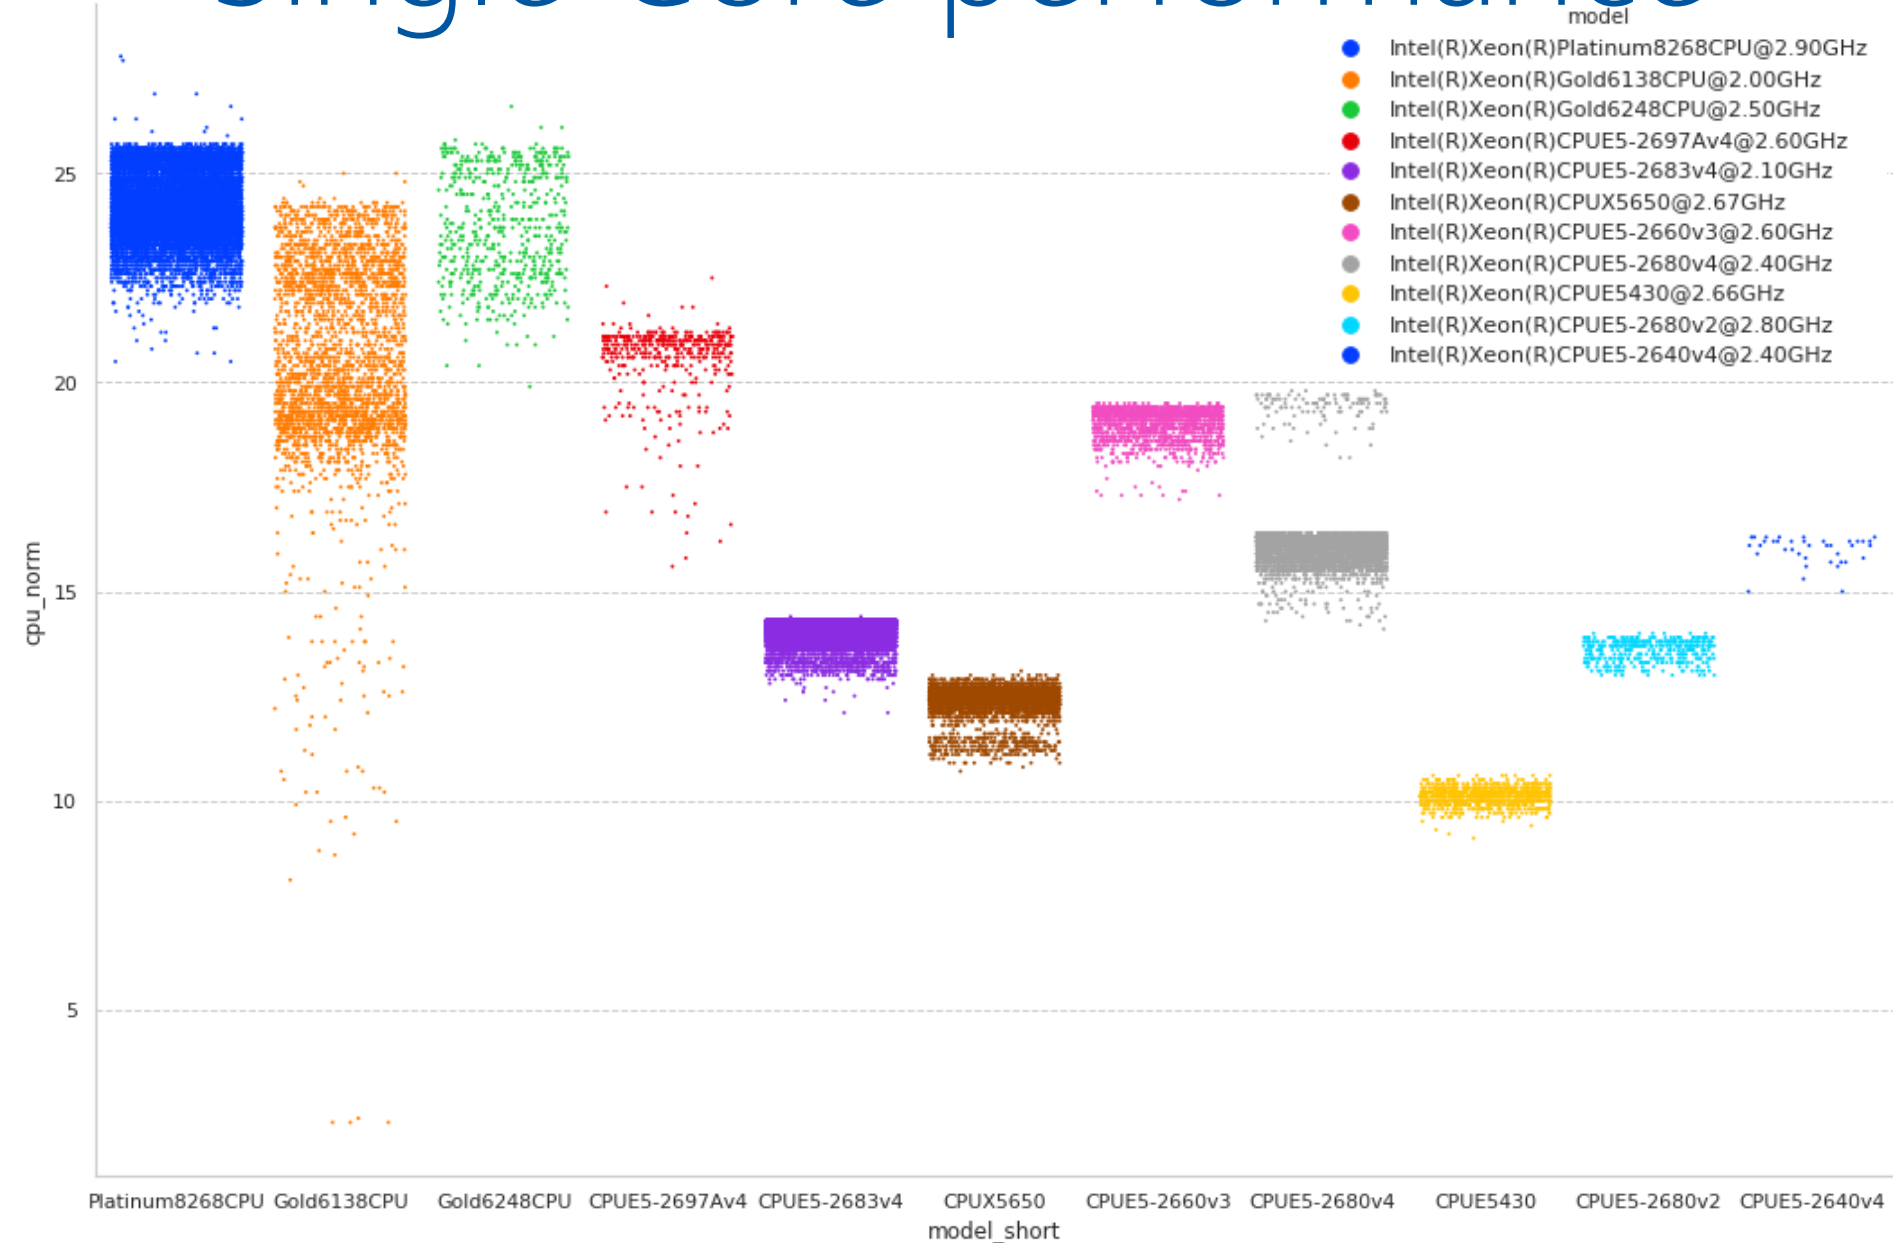
UNAM plot - June



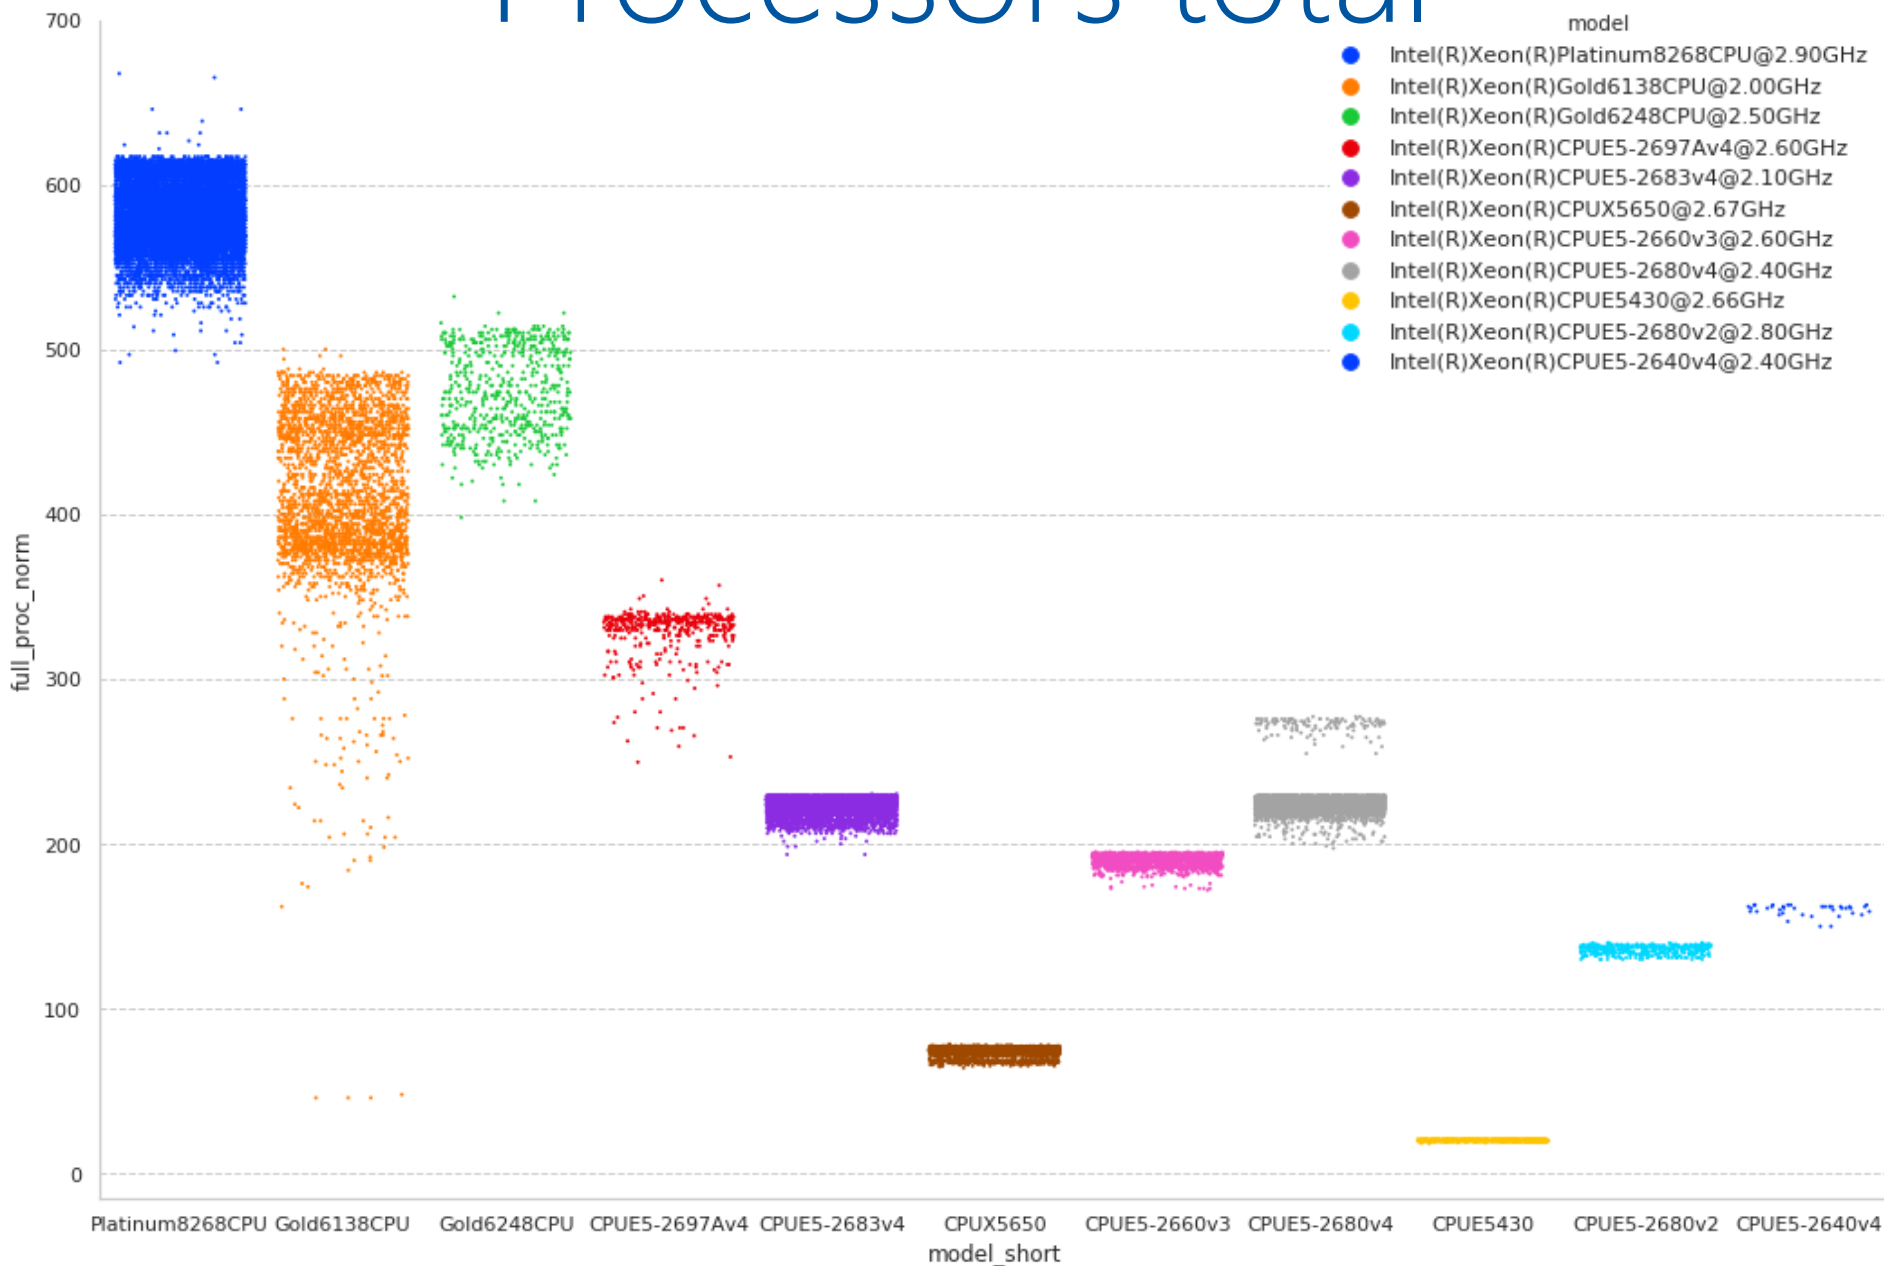
UNAM processors - June



Single Core performance



Processors total




Resources total

Resources	Slots	Average DB12
Govorun	375	24,2
LHEP Cluster	250	21,2
Tier1	440	14,05
Tier2	500	14,75
JINR Cloud	100	18,5

Conclusions

- DB12 benchmark correlates with performance of resources. Average DB12 are good for estimation of total performance of a cluster.
- Data transfer time is included in wall time increasing dispersion. More precise results may be received if we can subtract that time from walltime.
- When the system is operational, and users submit jobs, we may get intelligence about the performance, structure, components of heterogeneous resources "almost" "for free".



DUBNA

