Exascale Applications and Software Conference

Software System Stack for Efficiency of Exascale Supercomputer Centers

Vladimir Voevodin
Moscow State University

Prof. Vladimir V. Voevodin
Deputy Director, Research Computing Center
Head of the Department "Supercomputers and Quantum Informatics", CMC

voevodin@parallel.ru

April 21, 2015, Edinburgh

Computing Facilities of Moscow State University

(from 1956 up to now)

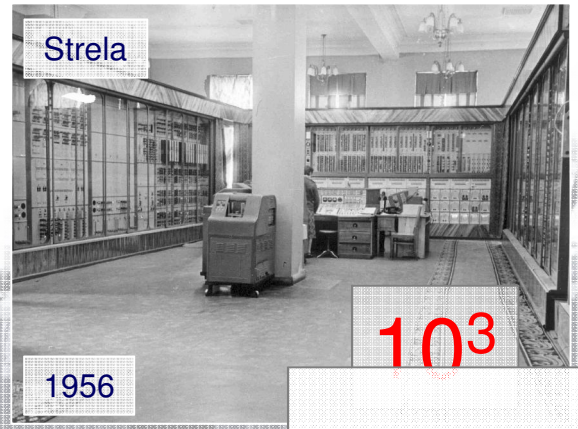

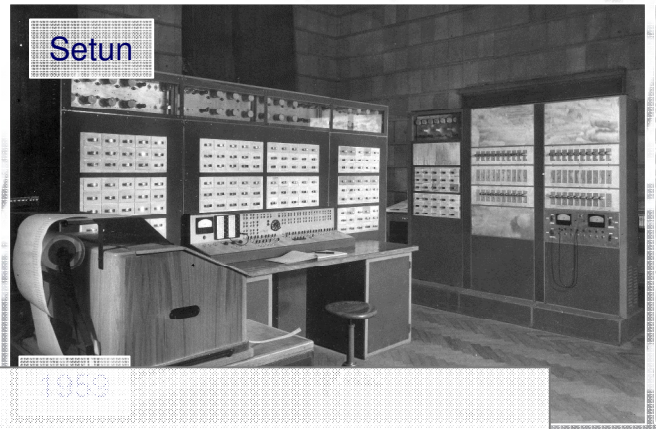

Brief stats on MSU Supercomputer Center

Users: 2511
Projects: 1607
Organizations: 302

MSU Faculties / Institutes: 20+

Computational science is everywhere…

Challenges of Large Supercomputer Centers : Efficiency, Efficiency, Efficiency…

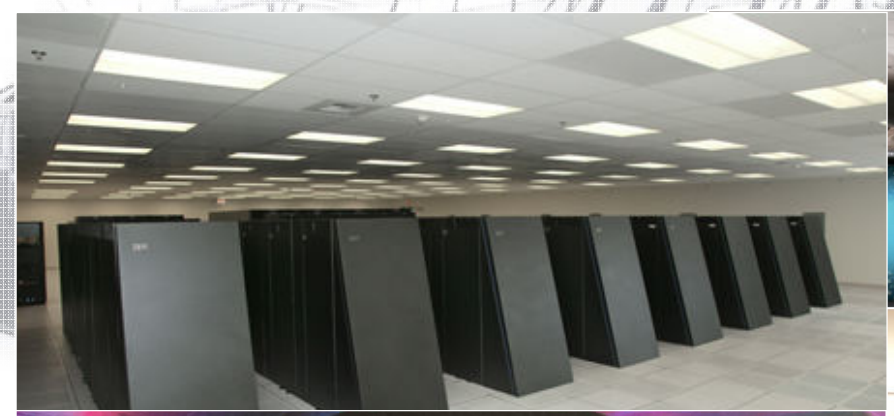

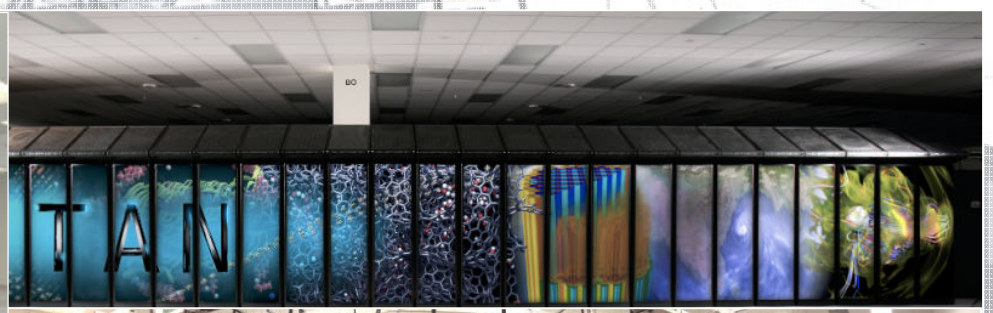

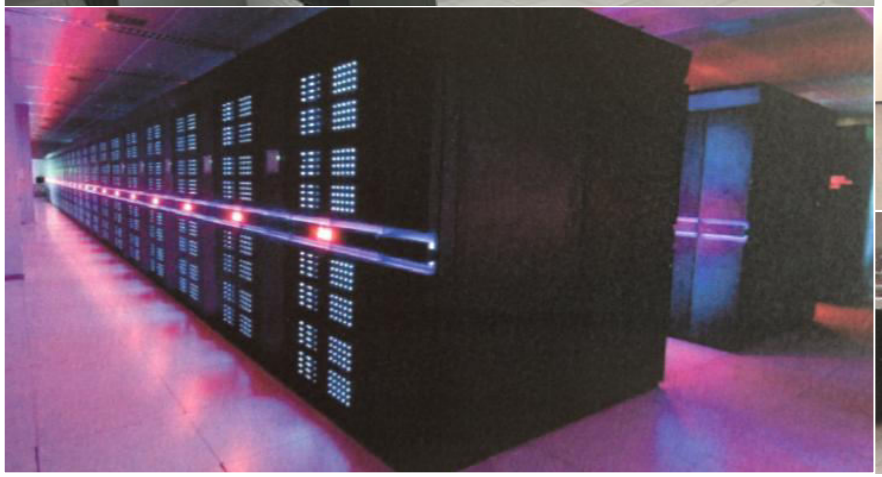

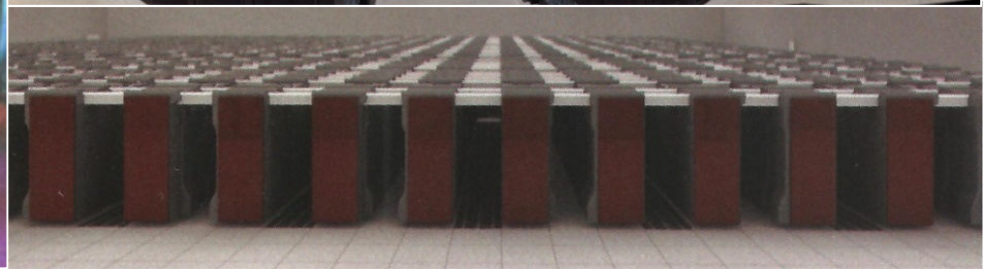

Efficiency of Supercomputing Centers

(practice, practice, practice)

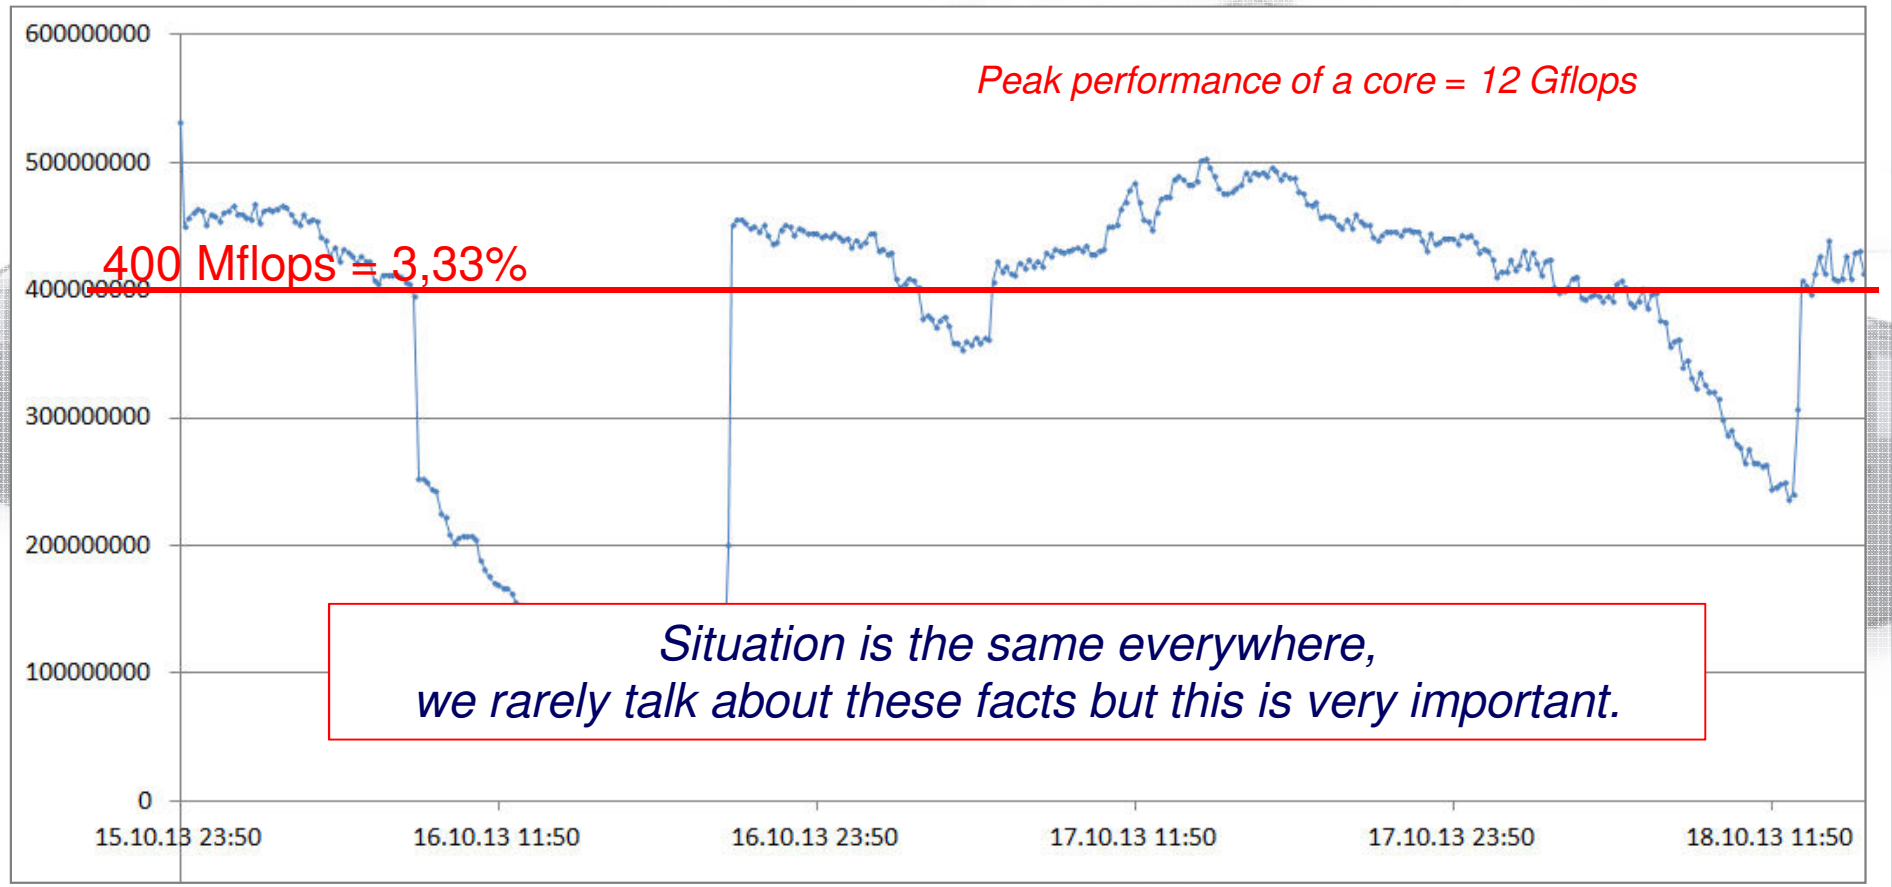

Average performance (one core) of "Chebyshev" supercomputer for 3 days

Efficiency of Supercomputing Centers
(typical questions)

- *What is a distribution of CPUhours between software packages for the last year?*
- *What is average intensity of InfiniBand and Ethernet networks usage for a particular user?*
- *How many nodes/cards/disks/cables do fail every month ?*
- *How often has Infiniband resent packages for the last week ?*
- *How often does LoadAVG exceed number of cores on computing nodes ?*
- *What is an min/max/average level of cache misses for applications of a particular user ?*
- *What is a distribution of waiting time in queues ?*
- *How does LoadAVG behave in my application during execution ?*
- *Who are 5% of the most inefficient applications/users ? (regarding CPULoad or LoadAVG or cache misses or …)*
- *…*

What do we need to know to control efficiency of supercomputer centers ?

Is it difficult to control few components ? A few ?..

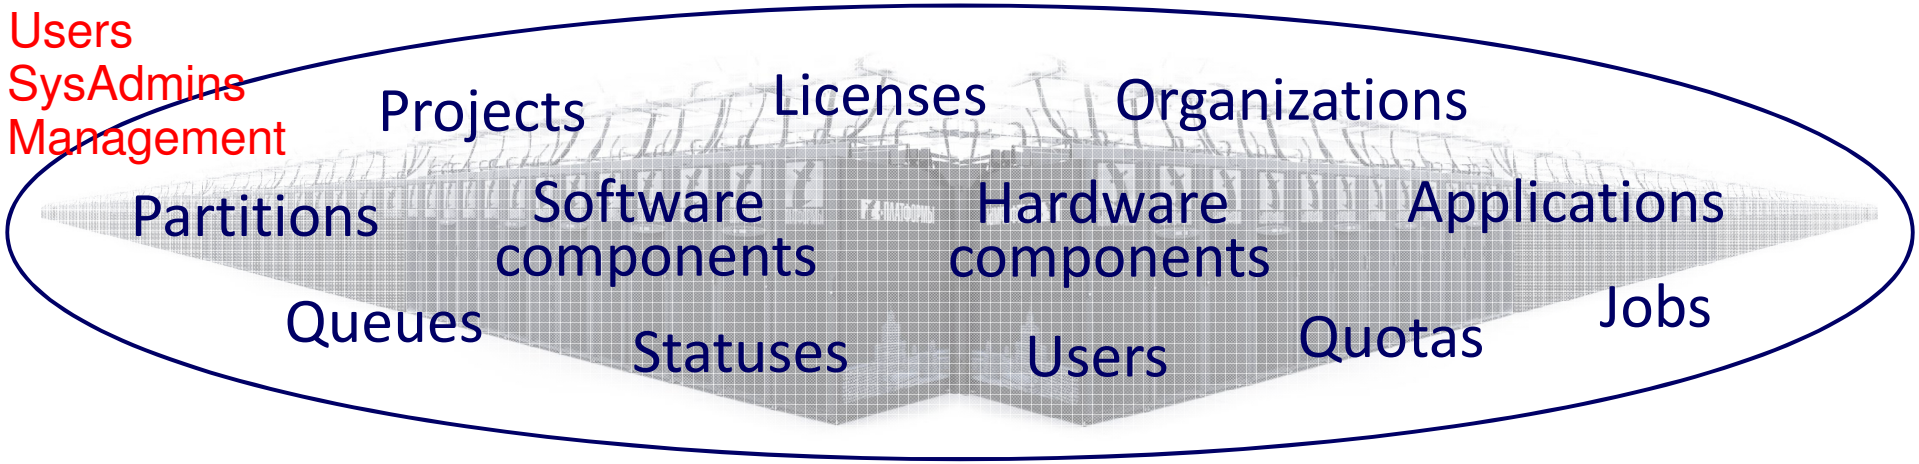

A few? Info on MSU "Lomonosov" Supercomputer :

(1.7 Pflops, 6000 computing nodes, 12K CPUs, 2K GPUs…)

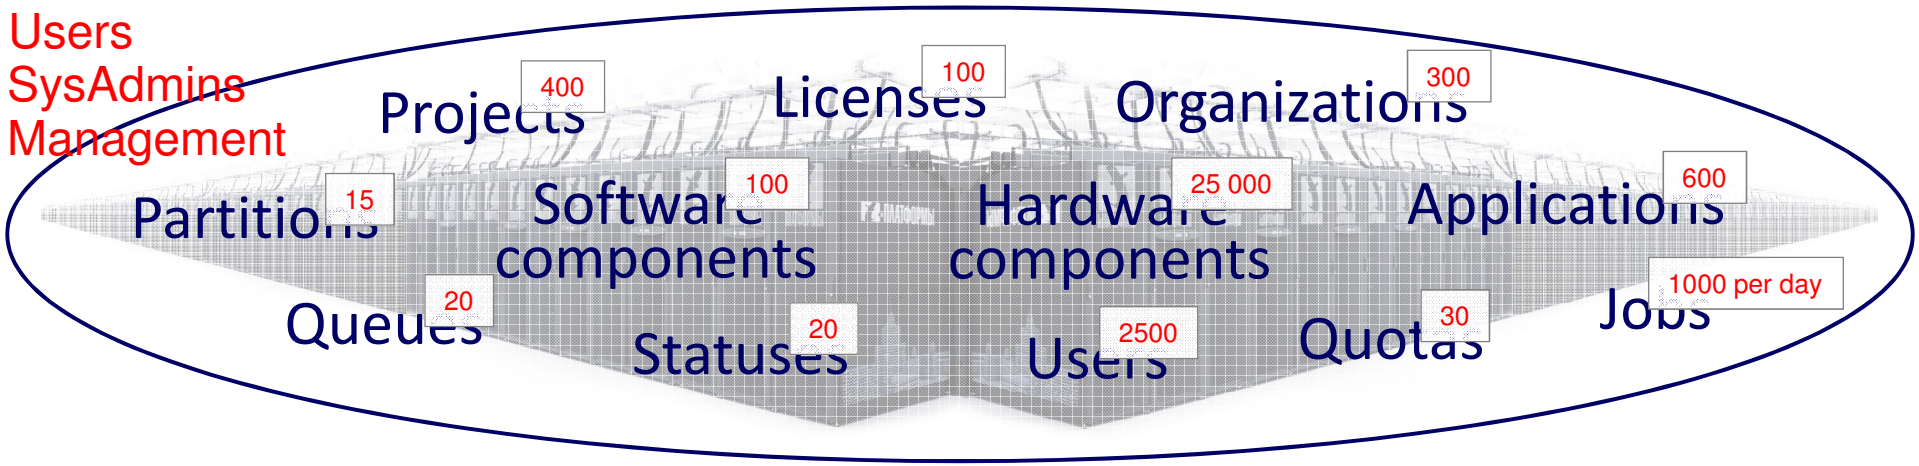

Current trend: all these numbers grow extremely fast !

⇩

It's impossible to predict/describe state of a supercomputer…

We have almost lost control over supercomputers…

Total control: cost of delay…

We need to keep control over supercomputers!

Current trend: the cost of delay with a proper reaction grows permanently !

Some facts:
One day of Lomonosov supercomputer (MSU) costs $25 000
One day of Titan supercomputer (ORNL) costs $300 000

Situation is the same everywhere,
we rarely talk about these facts but this is very important.

Total control: cost of delay…

We need to keep control over supercomputers!

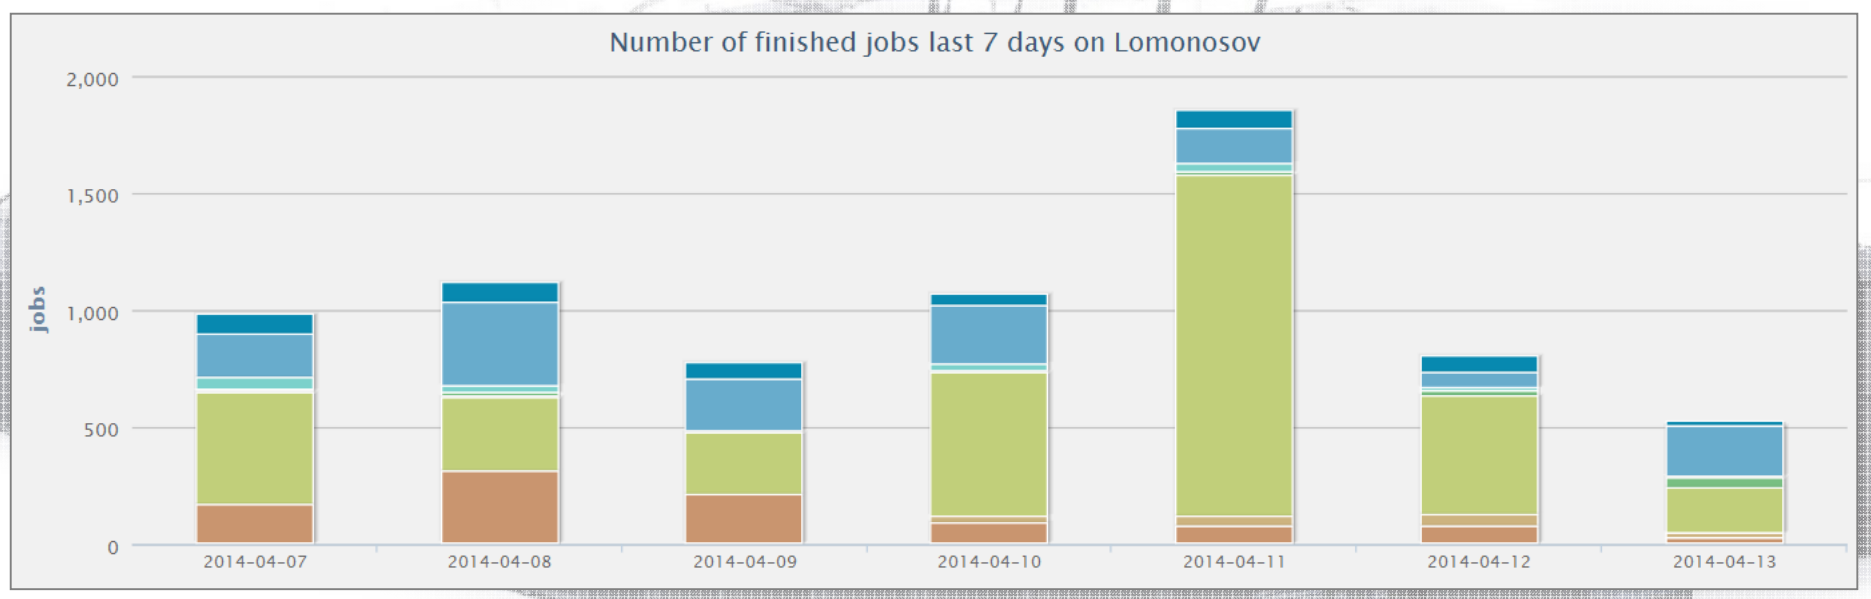

Supercomputer "Lomonosov":

- *about 1000 completed jobs per day,*
- *approx. 200 running jobs all the time,*

If the job scheduler hangs/dies, a half of the supercomputer will be idle in 2-3 hours.

What is our approach to control efficiency of supercomputer centers ?

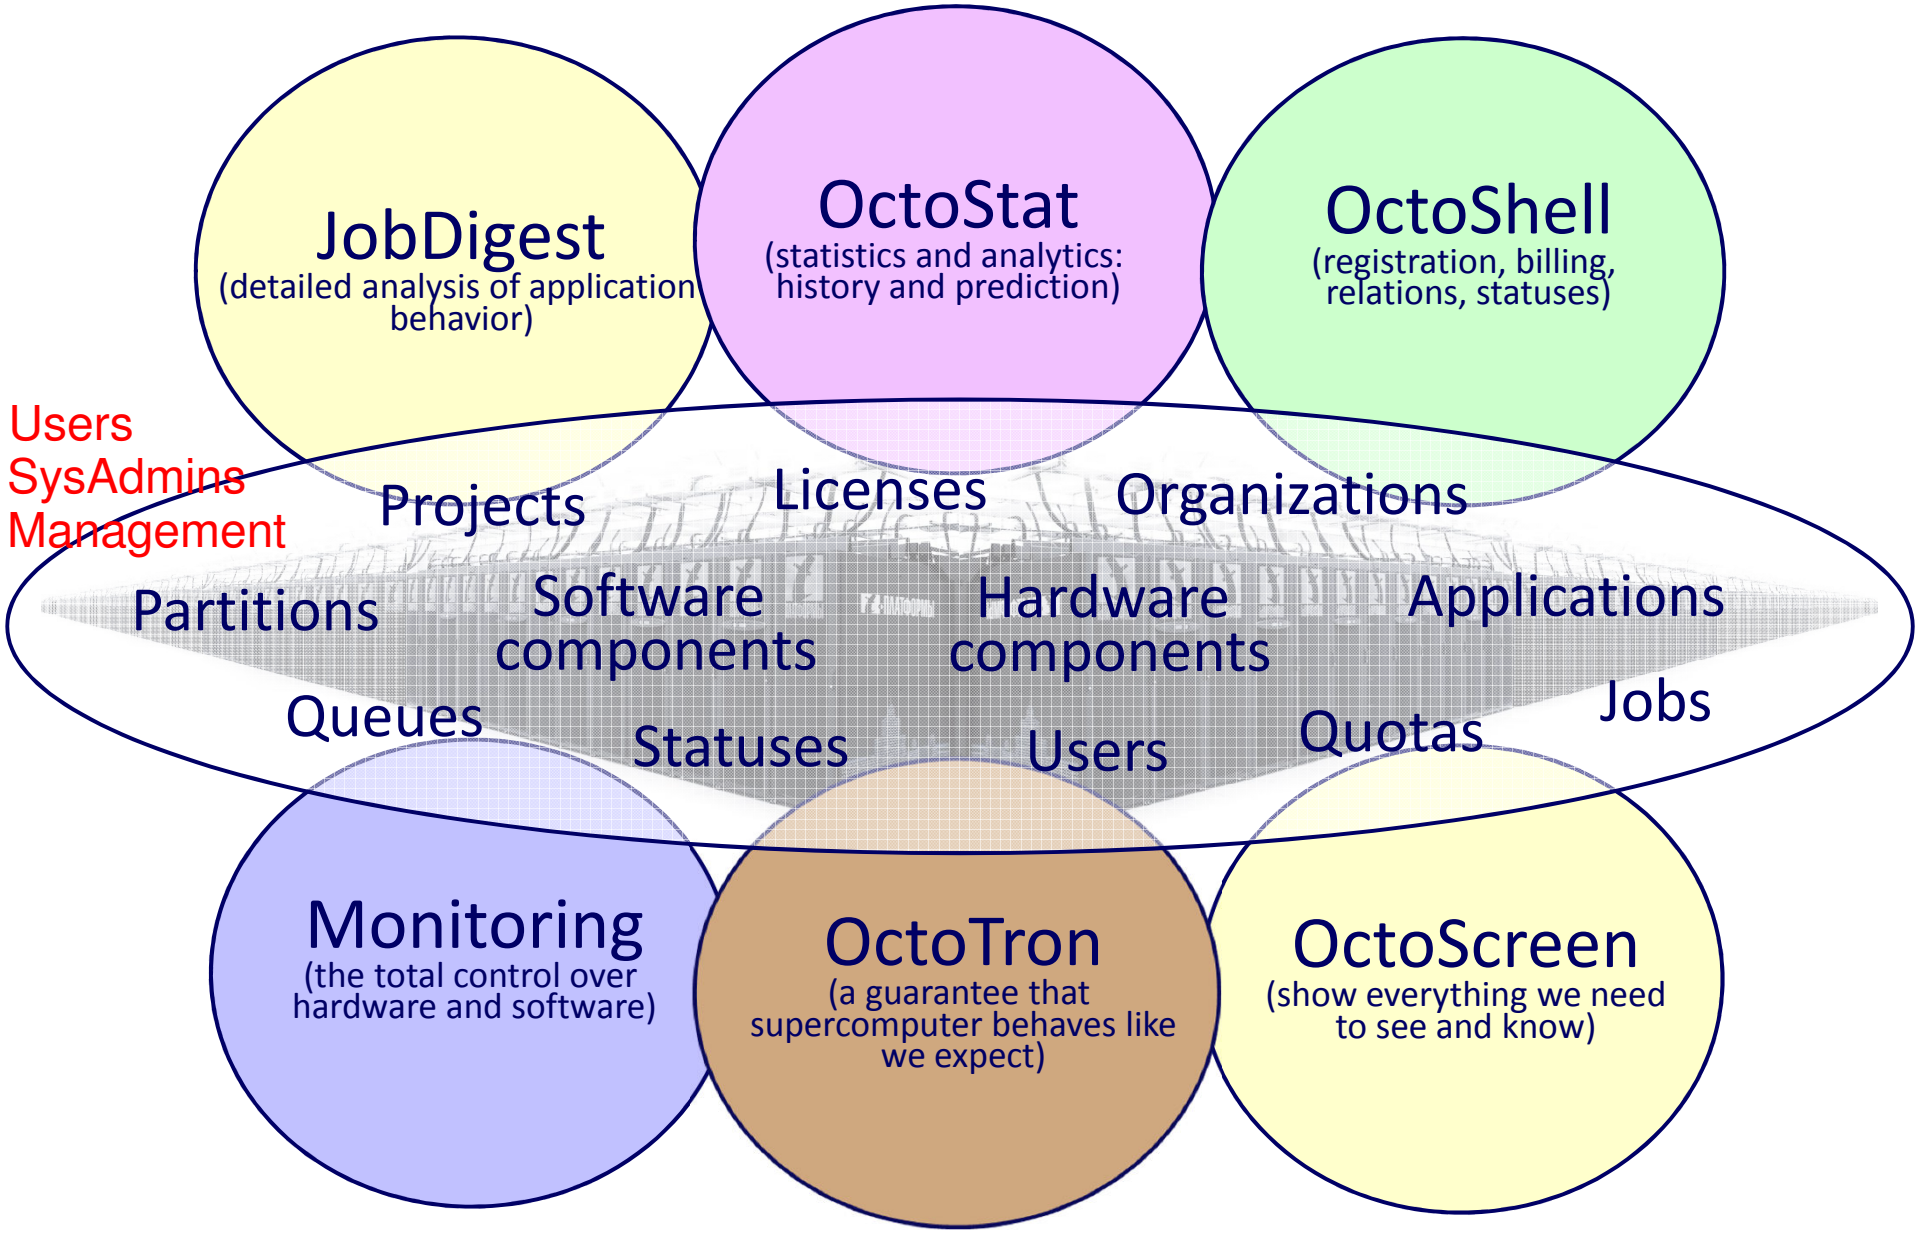

What are these systems designed for?

(in one word)

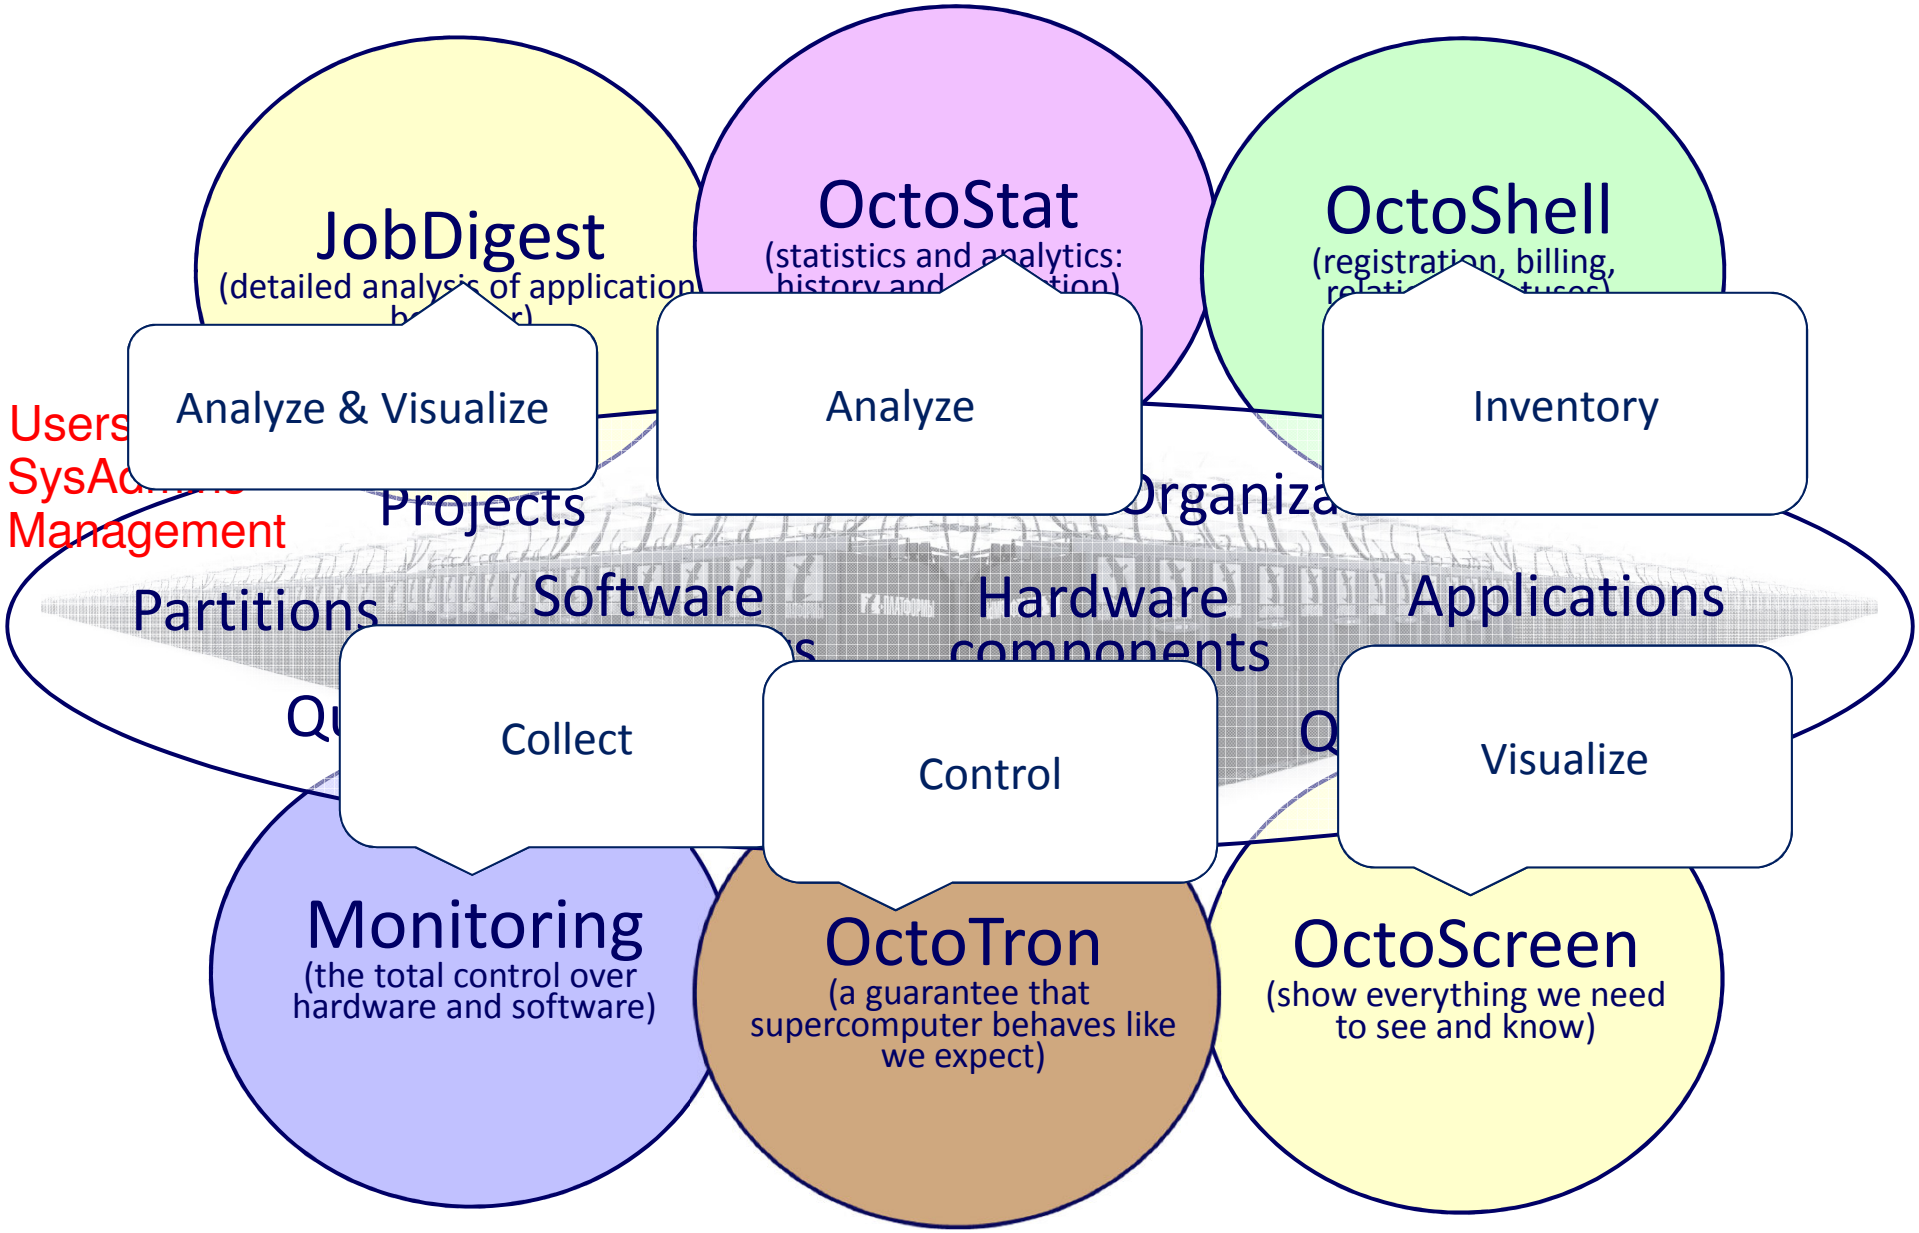

Some typical questions these systems address…

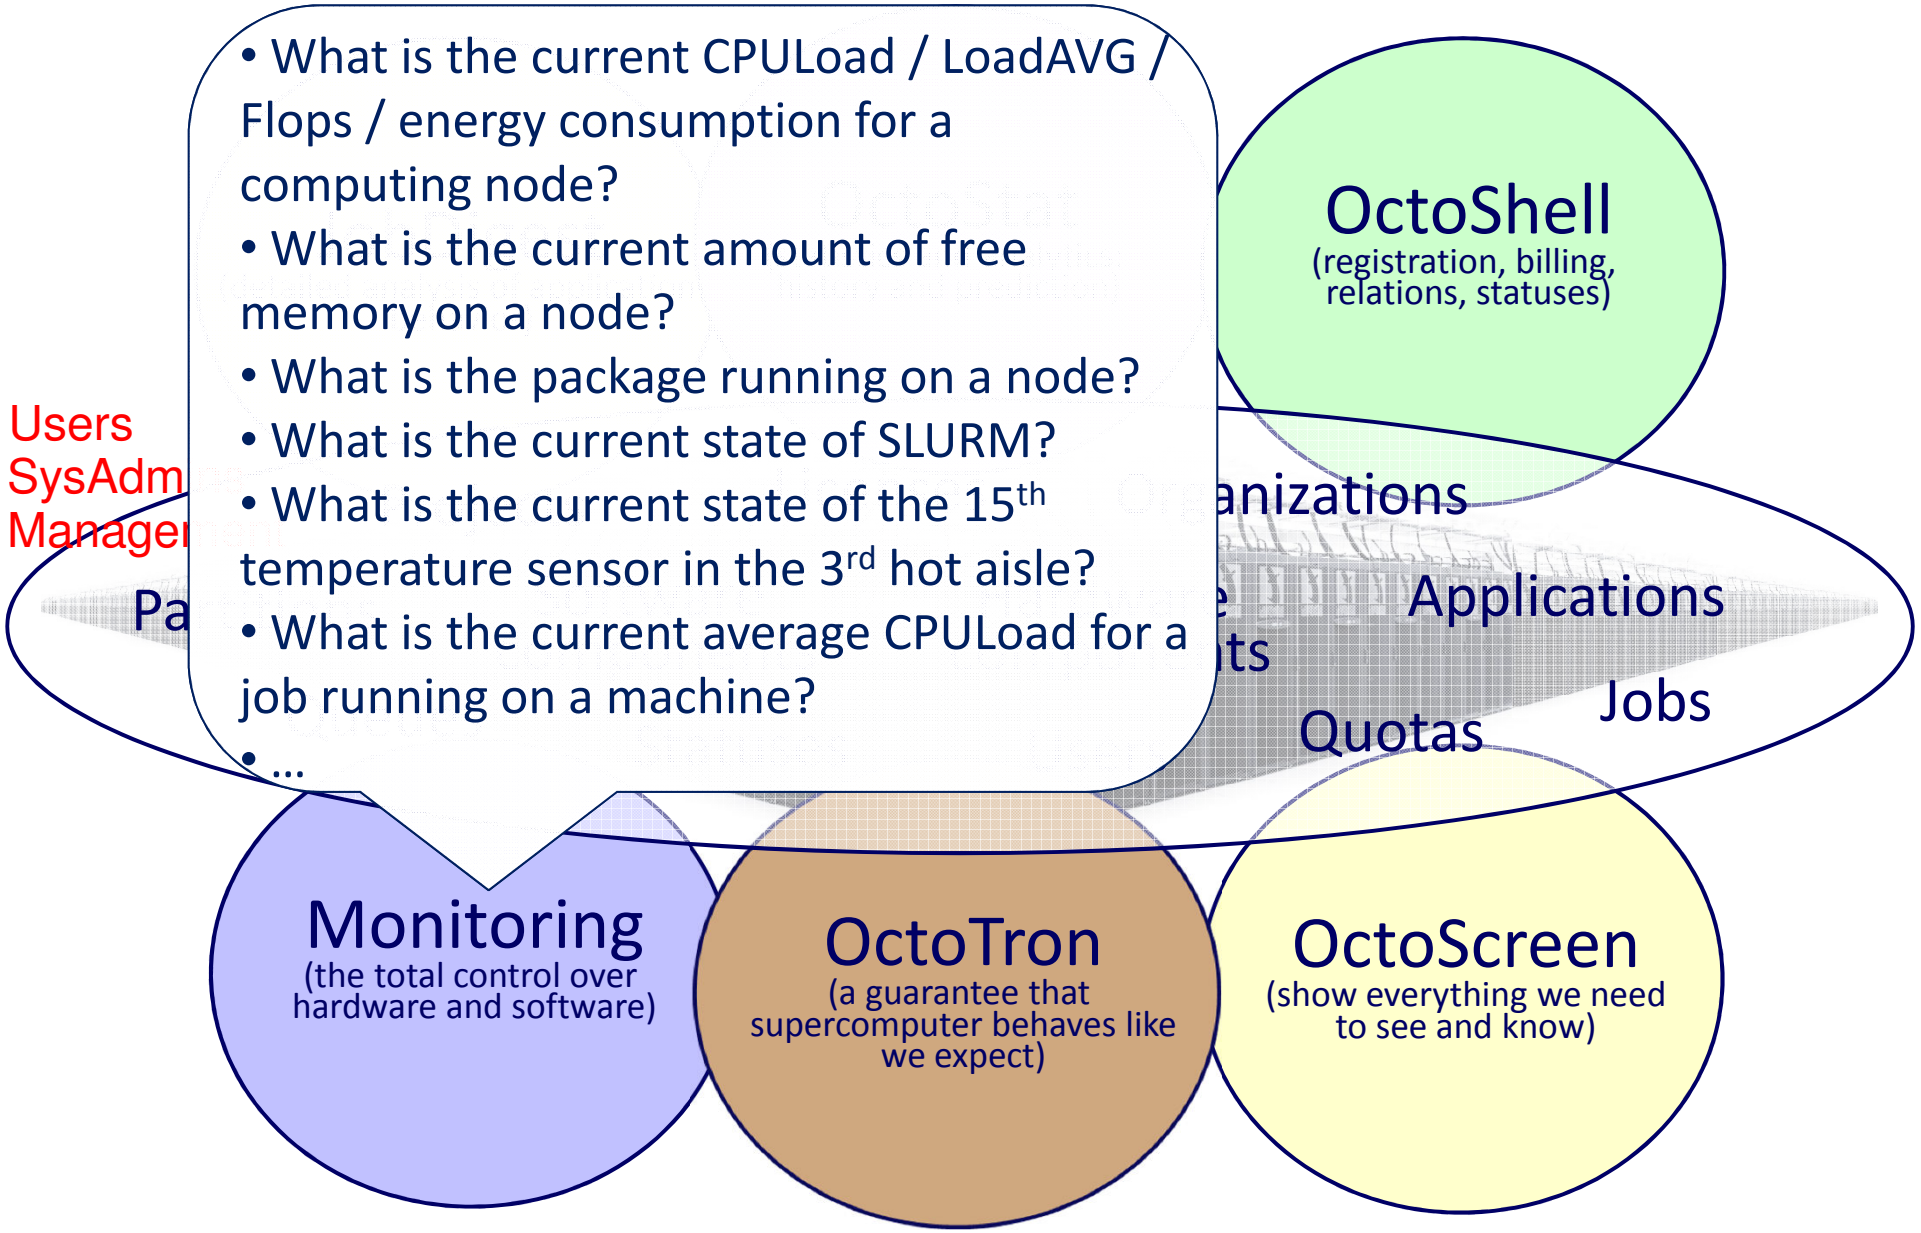

What is an average CPULoad for … ?

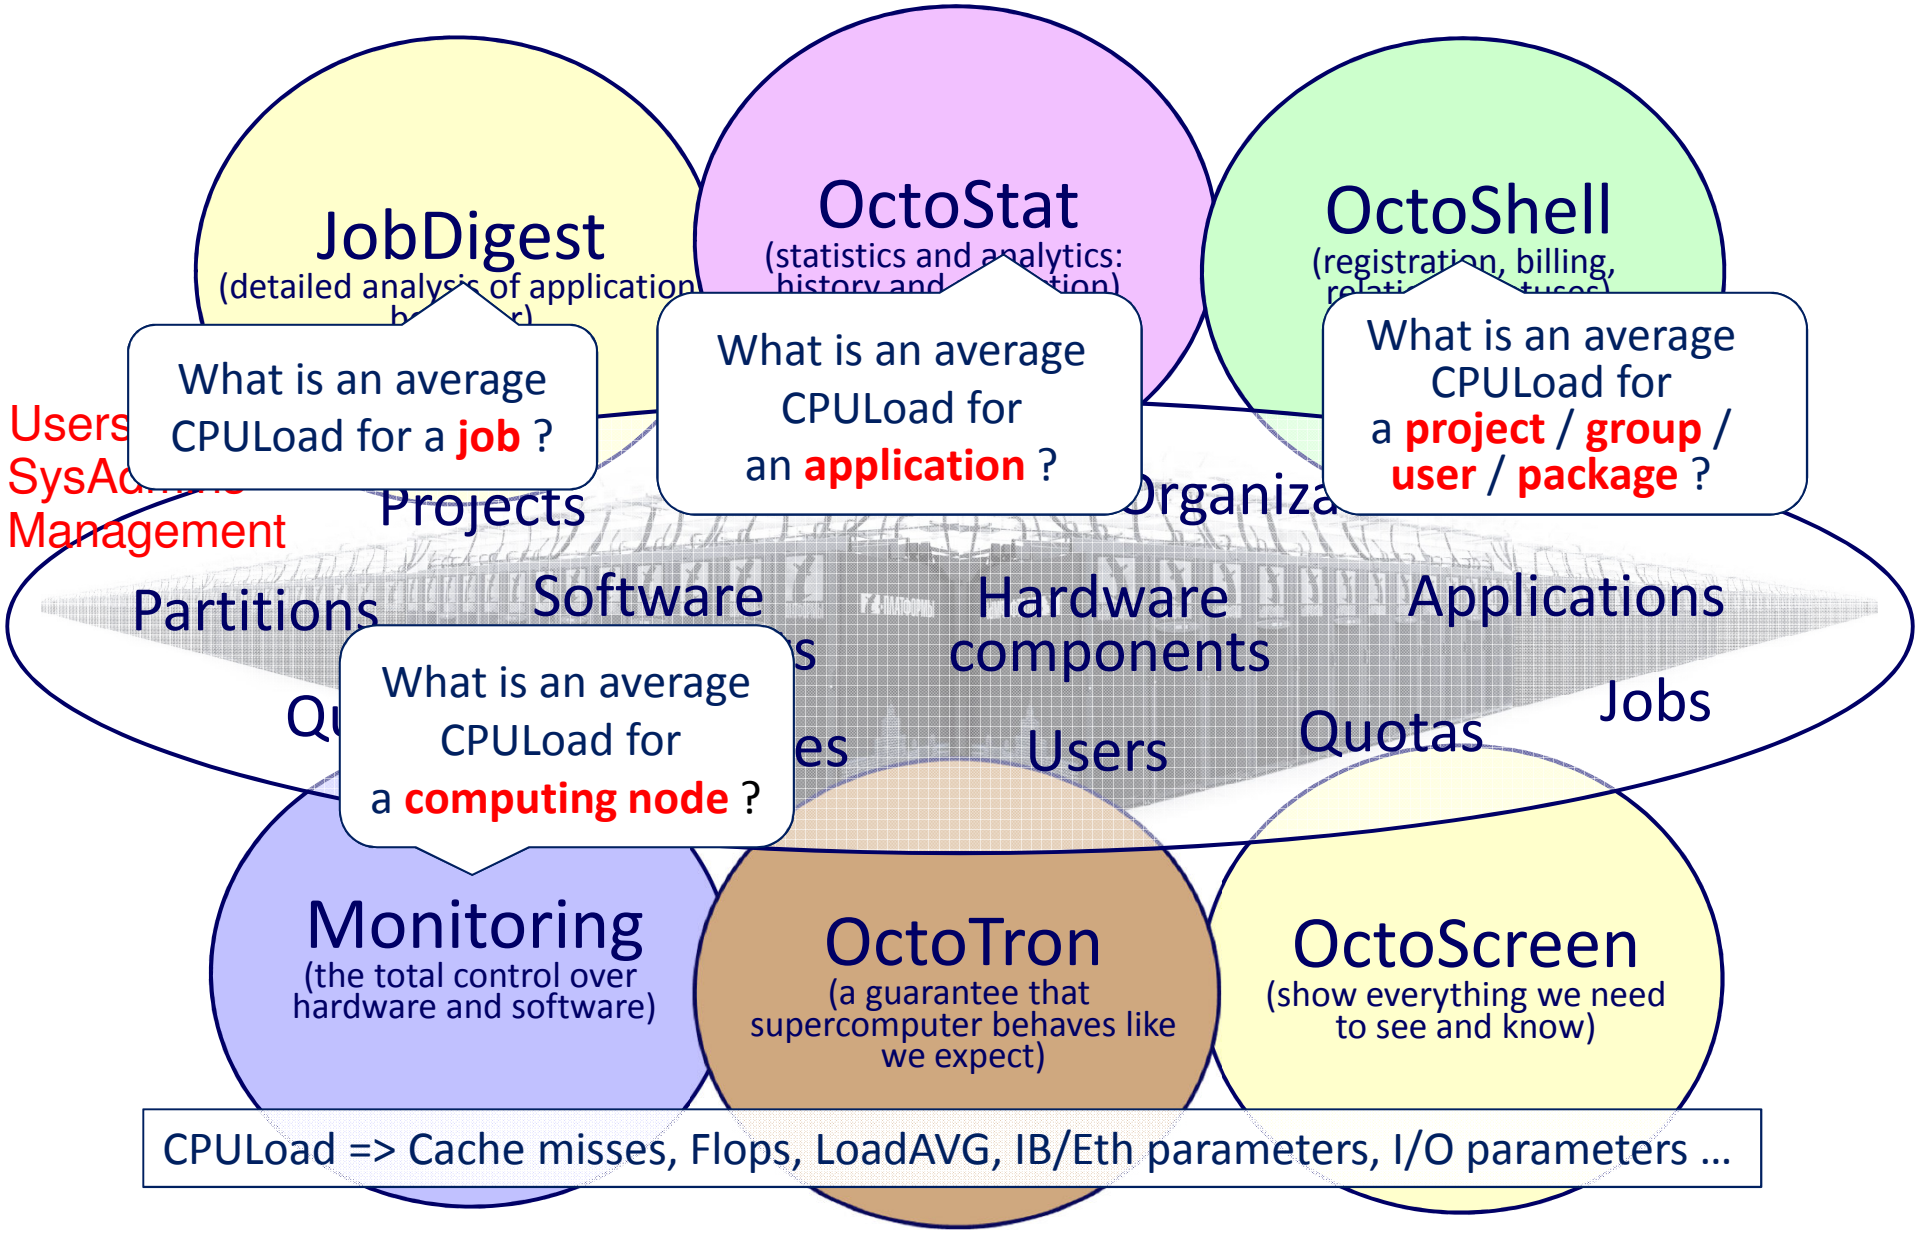

Systems and Challenges

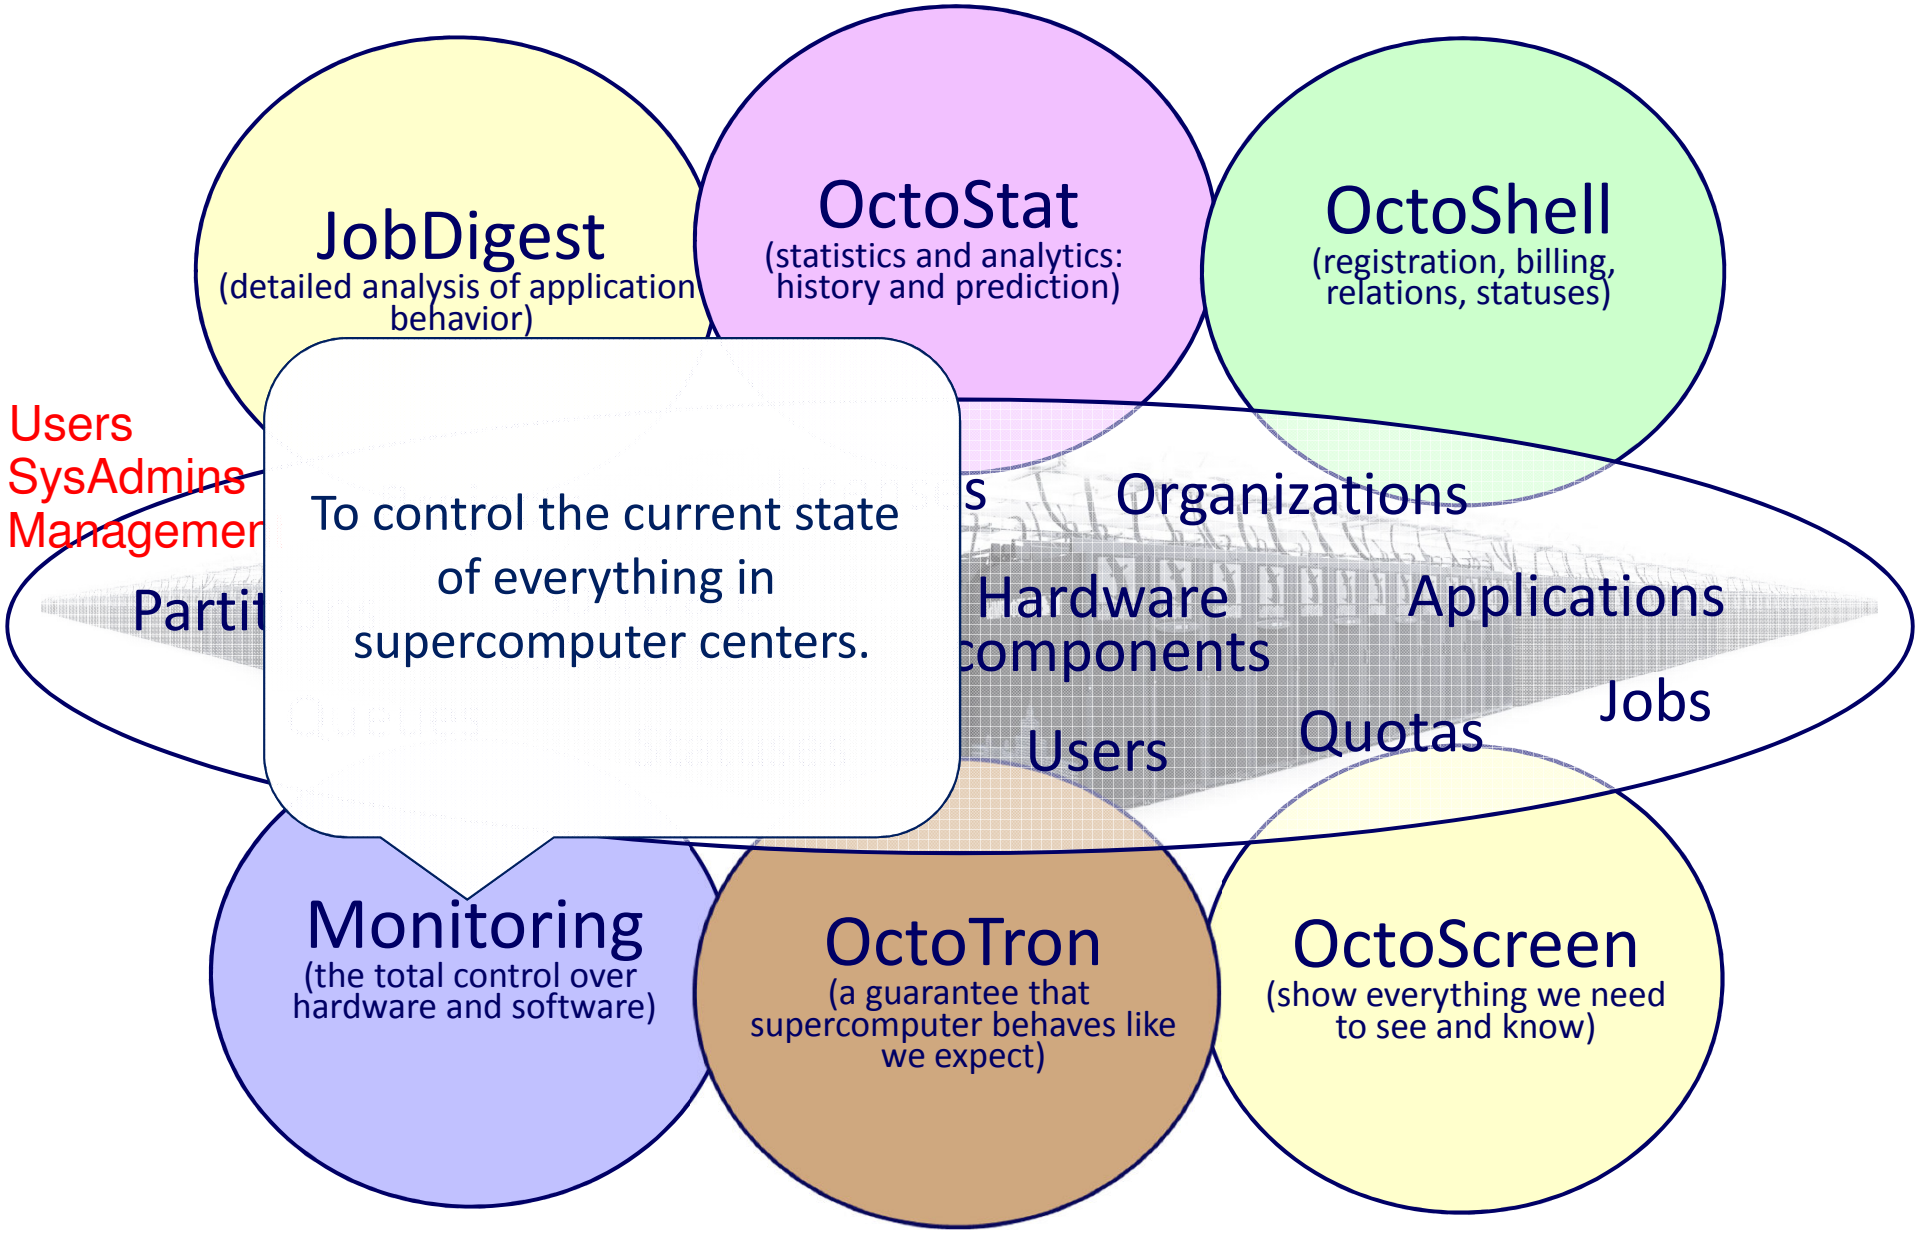

DiMMon Monitoring System

Monitoring system, requirements:

- *we need to know: what, where, when.*
- *scalability: millions of compute nodes, dozens sensors per node,*
- *low overheads: CPU, disks, interconnects (1% and less),*
- *frequency: a few seconds and less,*
- *easily reconfigurable and expandable,*
- *portable across platforms,*
- *active and passive modes*
- *…*
- *be careful: need to avoid BigData problems…*

Current trend: monitoring will be an integral part of all future complex HW&SW systems.

DiMMon Monitoring System

A smart approach to monitoring:

- *on-the-fly analysis: all relevant information should be extracted from the monitoring data before to be stored in a database (if necessary);*
- *in-situ analysis: monitoring data must be processed where the data were obtained (process first, move data later (if necessary));*
- *dynamic reconfiguration of monitoring systems: the monitoring system must be capable to change dynamically its configuration, depending on the current load on the supercomputer and the specific analysis objectives.*

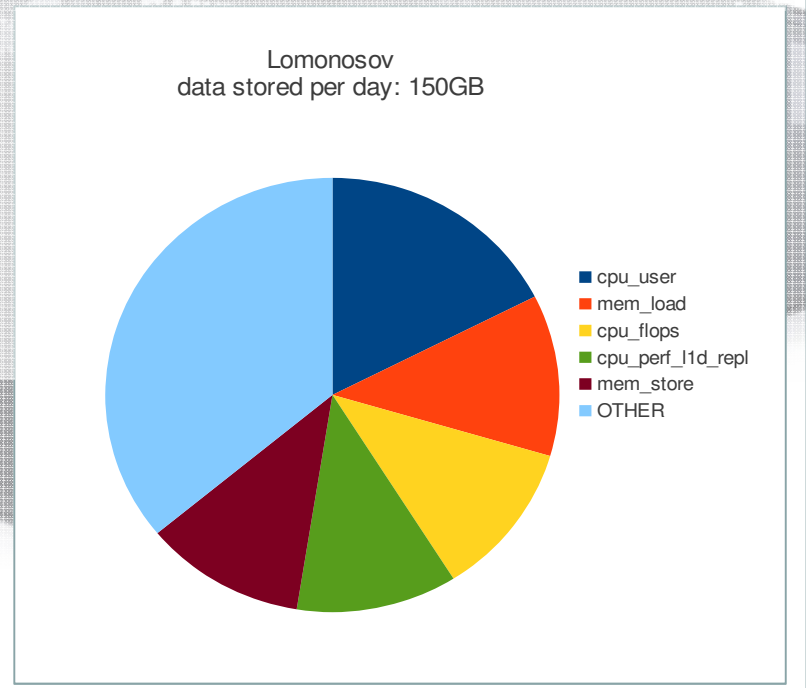

Aggressive filtering of data!

Current trend: total monitoring will be an integral part of all future complex HW&SW systems.

Efficiency of Supercomputing Centers

(practice, practice, practice)

Users

SysAdm

Manage

P

S

S

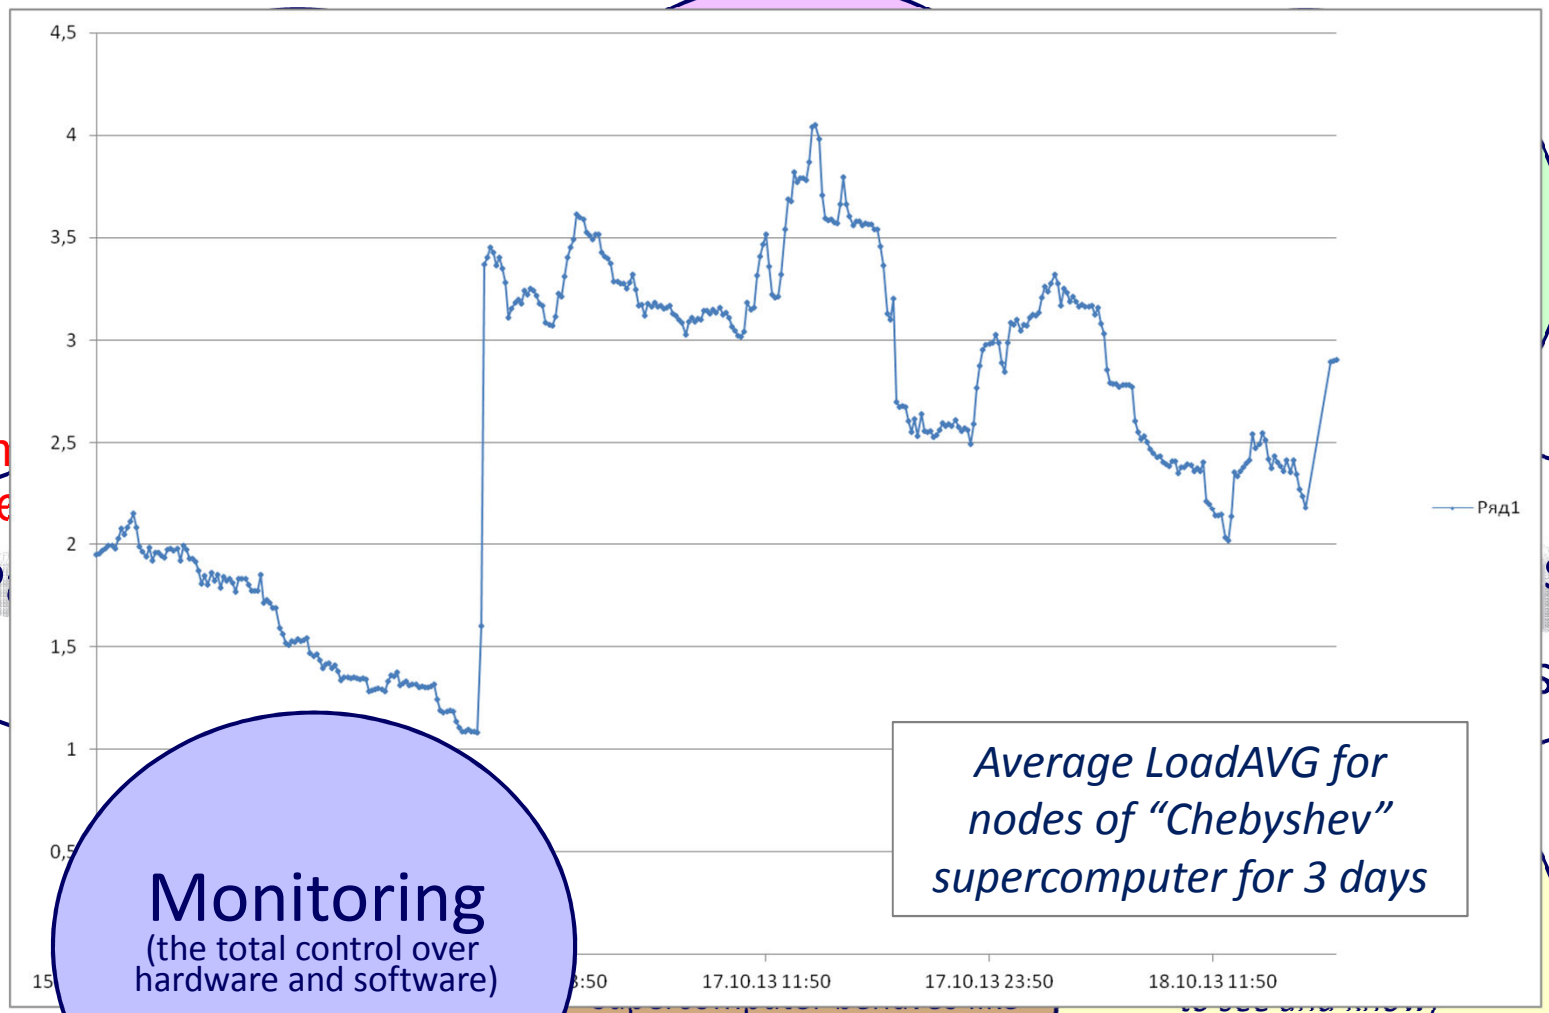

Average LoadAVG for nodes of "Chebyshev" supercomputer for 3 days

Monitoring
(the total control over hardware and software)

we expect)

Systems and Challenges

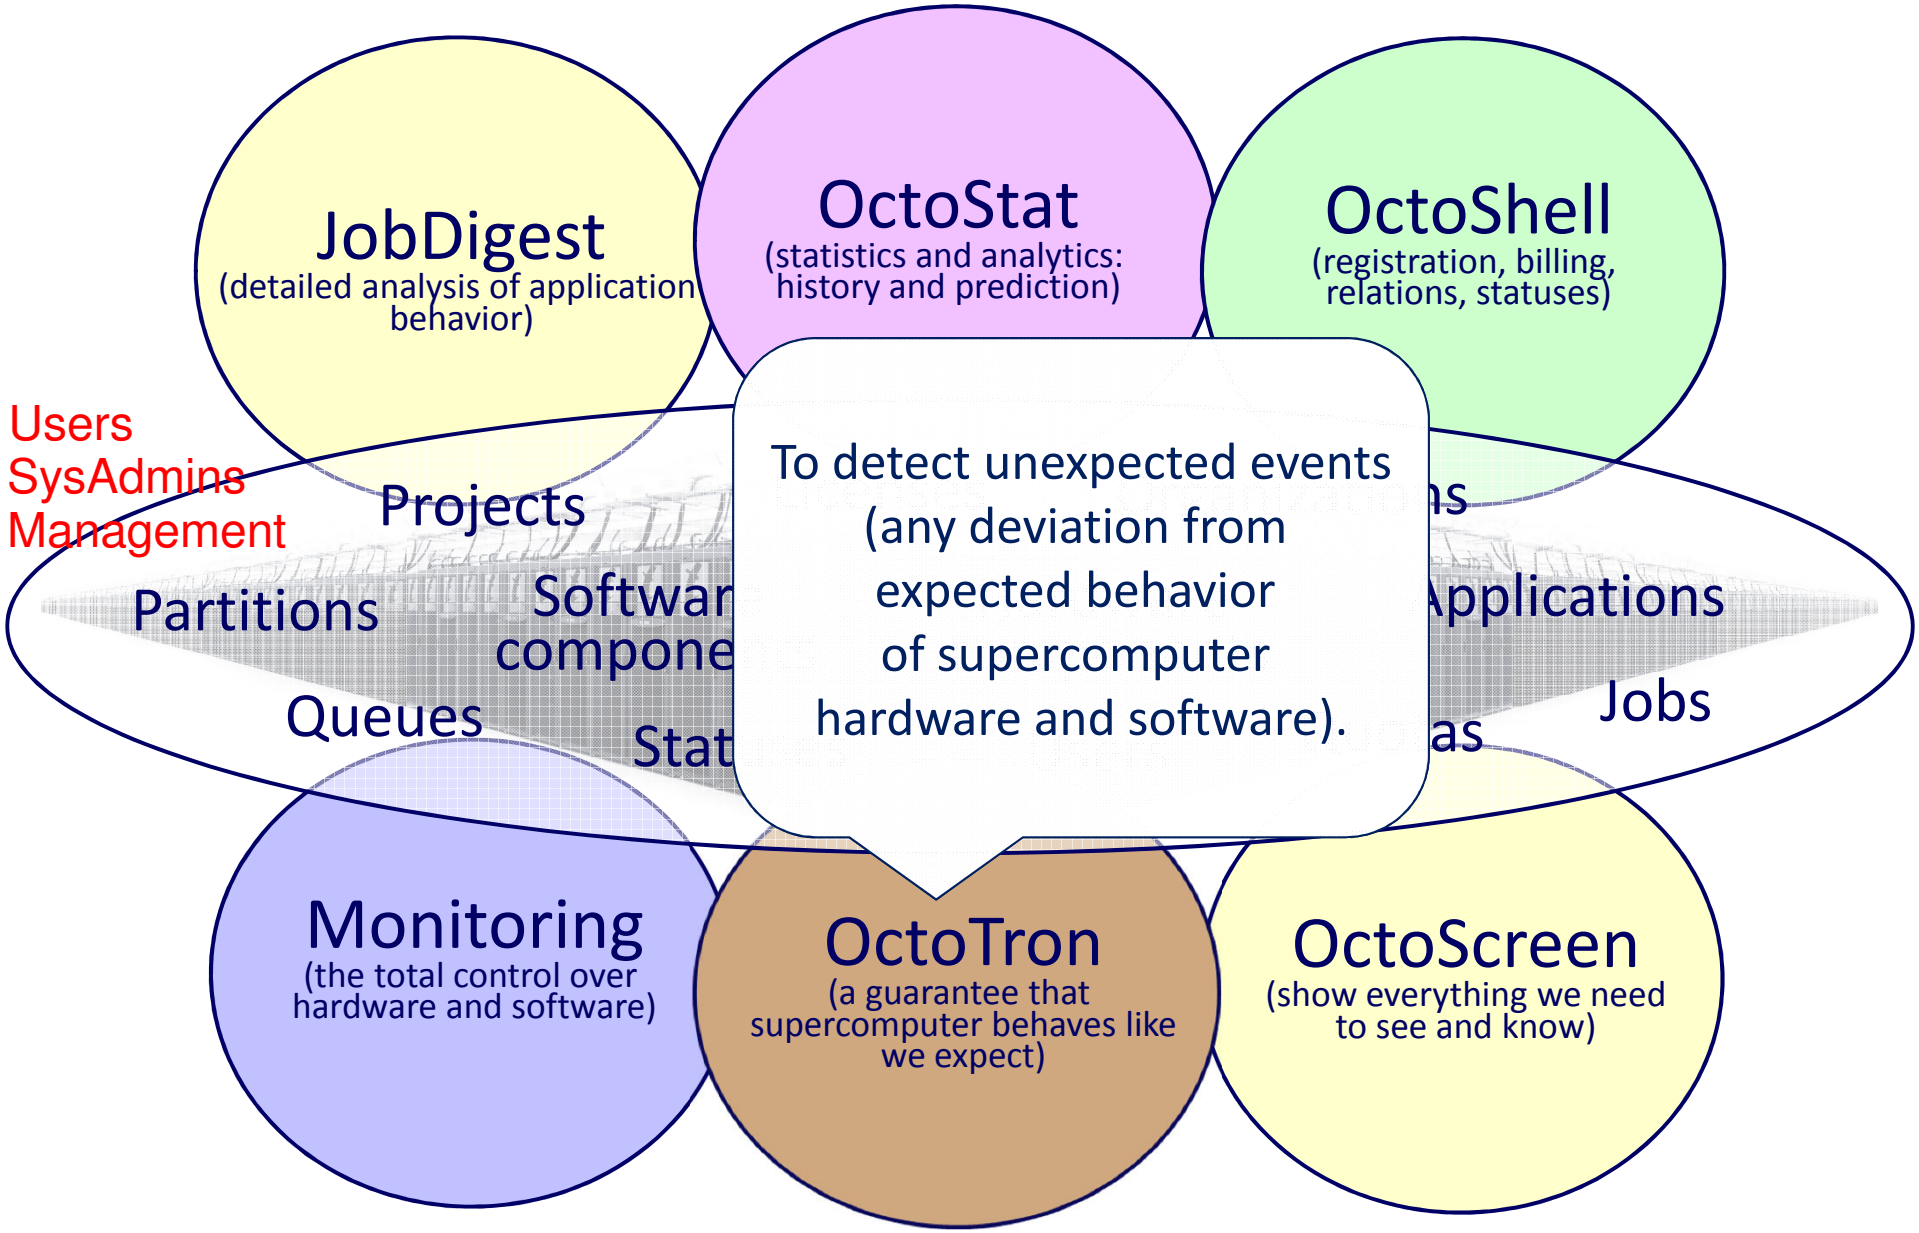

Guarantee, Predictability and Autonomous Life of Supercomputers

***Large numbers in supercomputers**: cores, processors, accelerators, nodes, HW&SW components, files, indexes, users, projects, processes, threads, running and queued jobs…*

We don't know for sure the current state of supercomputer's components …

Guarantee, Predictability and Autonomous Life of Supercomputers

What is now? *We hope a component works until we get an evidence that it has failed.*

What do we need?

We need a guarantee*:*
if something goes wrong inside a supercomputer we shall be notified immediately.

We want a system behaves in a way we expect it should behave.

Efficiency of Supercomputing Centers

(practice, practice, practice)

Users

SysAdmi

Managem

Par

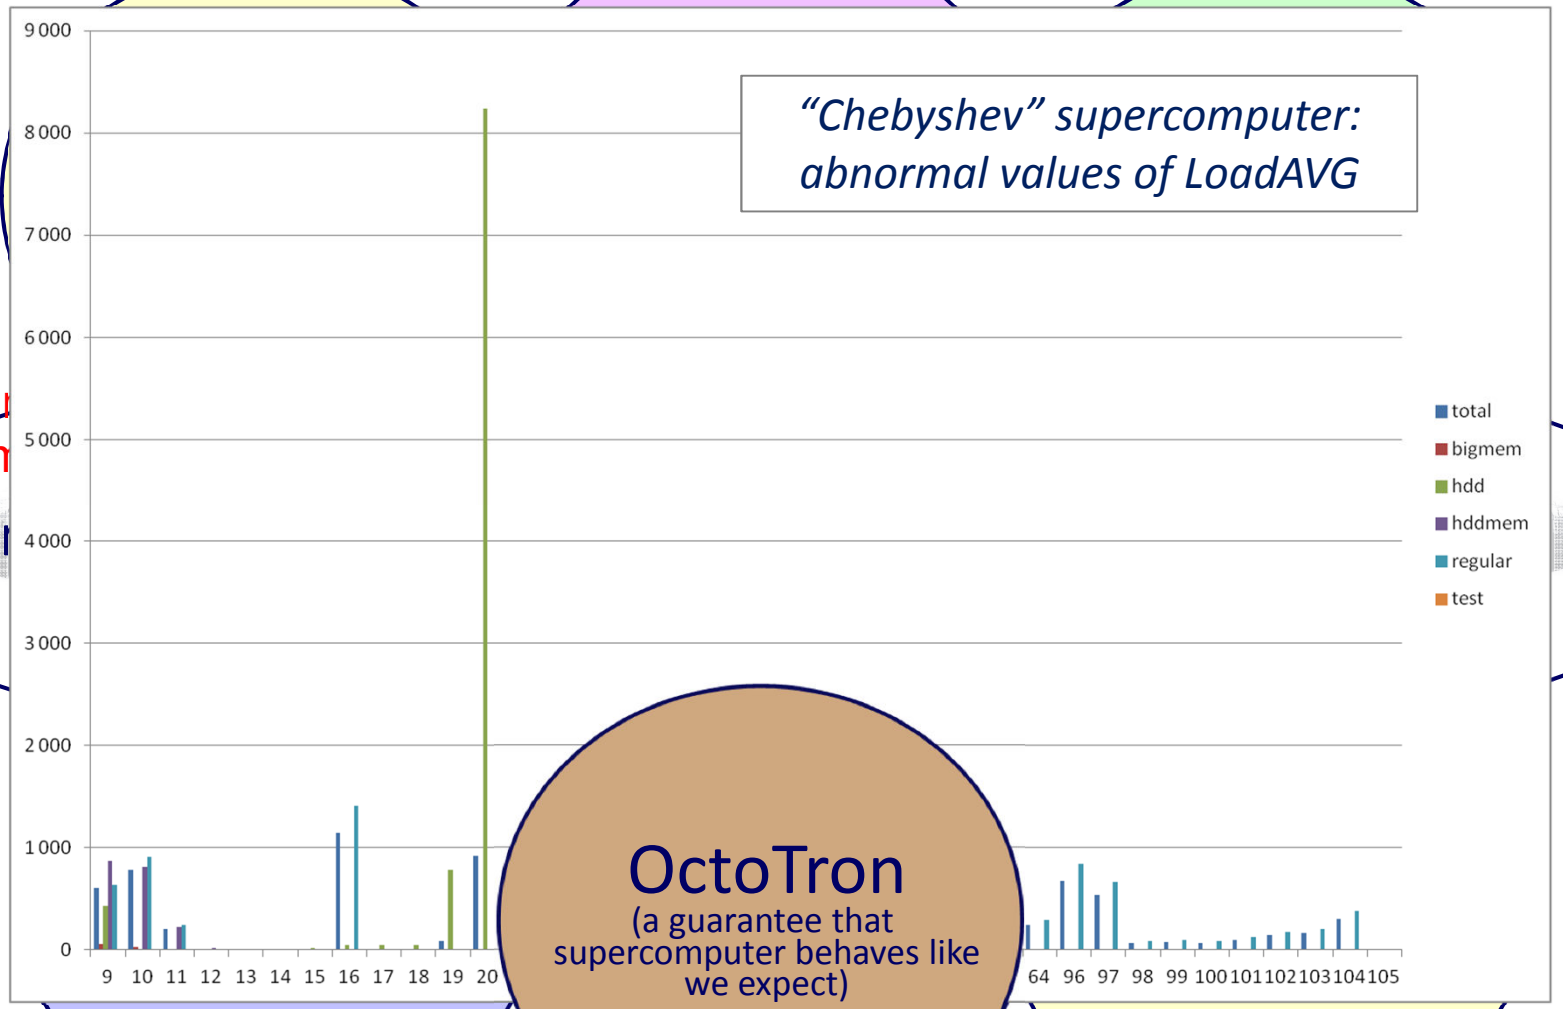

OctoTron

(a guarantee that supercomputer behaves like we expect)

Guarantee, Predictability and Autonomous Life of Supercomputers

If discrepancy occurs between our expectations and supercomputer behavior we need to know immediately about it.

How can we do that? Supercomputer is huge, we can't control it to a full extent any more.

But… supercomputer can do it itself (instead of us), if we explain what "our expectations" are.

Guarantee, Predictability and Autonomous Life of Supercomputers

Supercomputers should be autonomous in self-control.
("Systems become more dynamic." (Pete Beckman)

Moreover:
The larger supercomputers, the more autonomous they should be.

Exascale Applications and Software Conference

Software System Stack for Efficiency of Exascale Supercomputer Centers
(to be autonomous or not to be efficient)

Vladimir Voevodin
Moscow State University

Prof. Vladimir V. Voevodin
Deputy Director, Research Computing Center
Head of the Department "Supercomputers and Quantum Informatics", CMC

voevodin@parallel.ru

April 21, 2015, Edinburgh

Guarantee, Predictability and Autonomous Life of Supercomputers

A guarantee of "our expectations = reality", how this can be done?

- *a formal model of supercomputers (**model is a graph**),*
- *a set of formal rules,*
- *a set of reactions,*

Autonomous life and control of MSU supercomputers:

*-**"Chebyshev"** supercomputer, 60 Tflops, 625 CPUs:*
10 228 nodes, 24 698 edges, 205 044 attributes, 160 rules, 100 reactions;
*- **"Lomonosov"** supercomputer, 1.7 Pflops, 12K CPUs, 2K GPU:*
116 000 nodes, 332 000 edges, 2 400 000 attributes,...

Initial deployment, Detection of faults, critical and emergency situations, Turning off minimum amount of hardware, Self diagnostics, Previous accidents, etc. are done according to a model and rules.

Current trend: many decisions about control over HW&SW of supercomputers must be taken automatically.

Efficiency of Supercomputing Centers

(practice, practice, practice)

"Chebyshev" supercomputer: intensity of IB usage for different partitions

OctoStat

(statistics and analytics: history and prediction)

Users

SysAdmin

Managem

Par

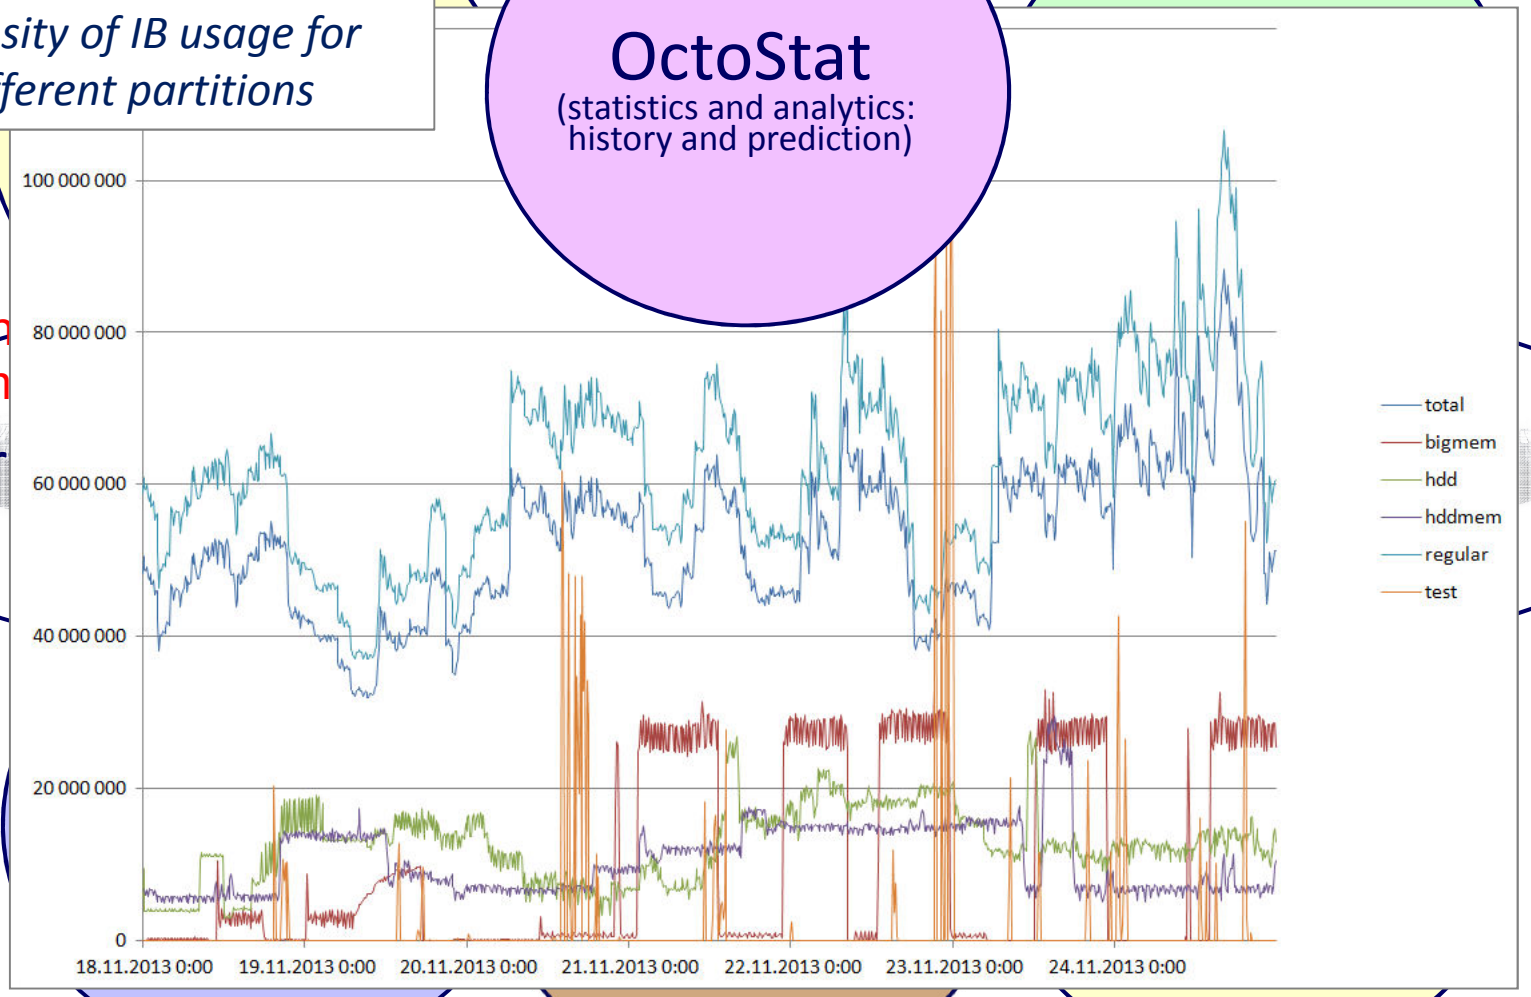

Efficiency of Supercomputing Centers

(practice, practice, practice)

JobDigest
(detailed analysis of application behavior)

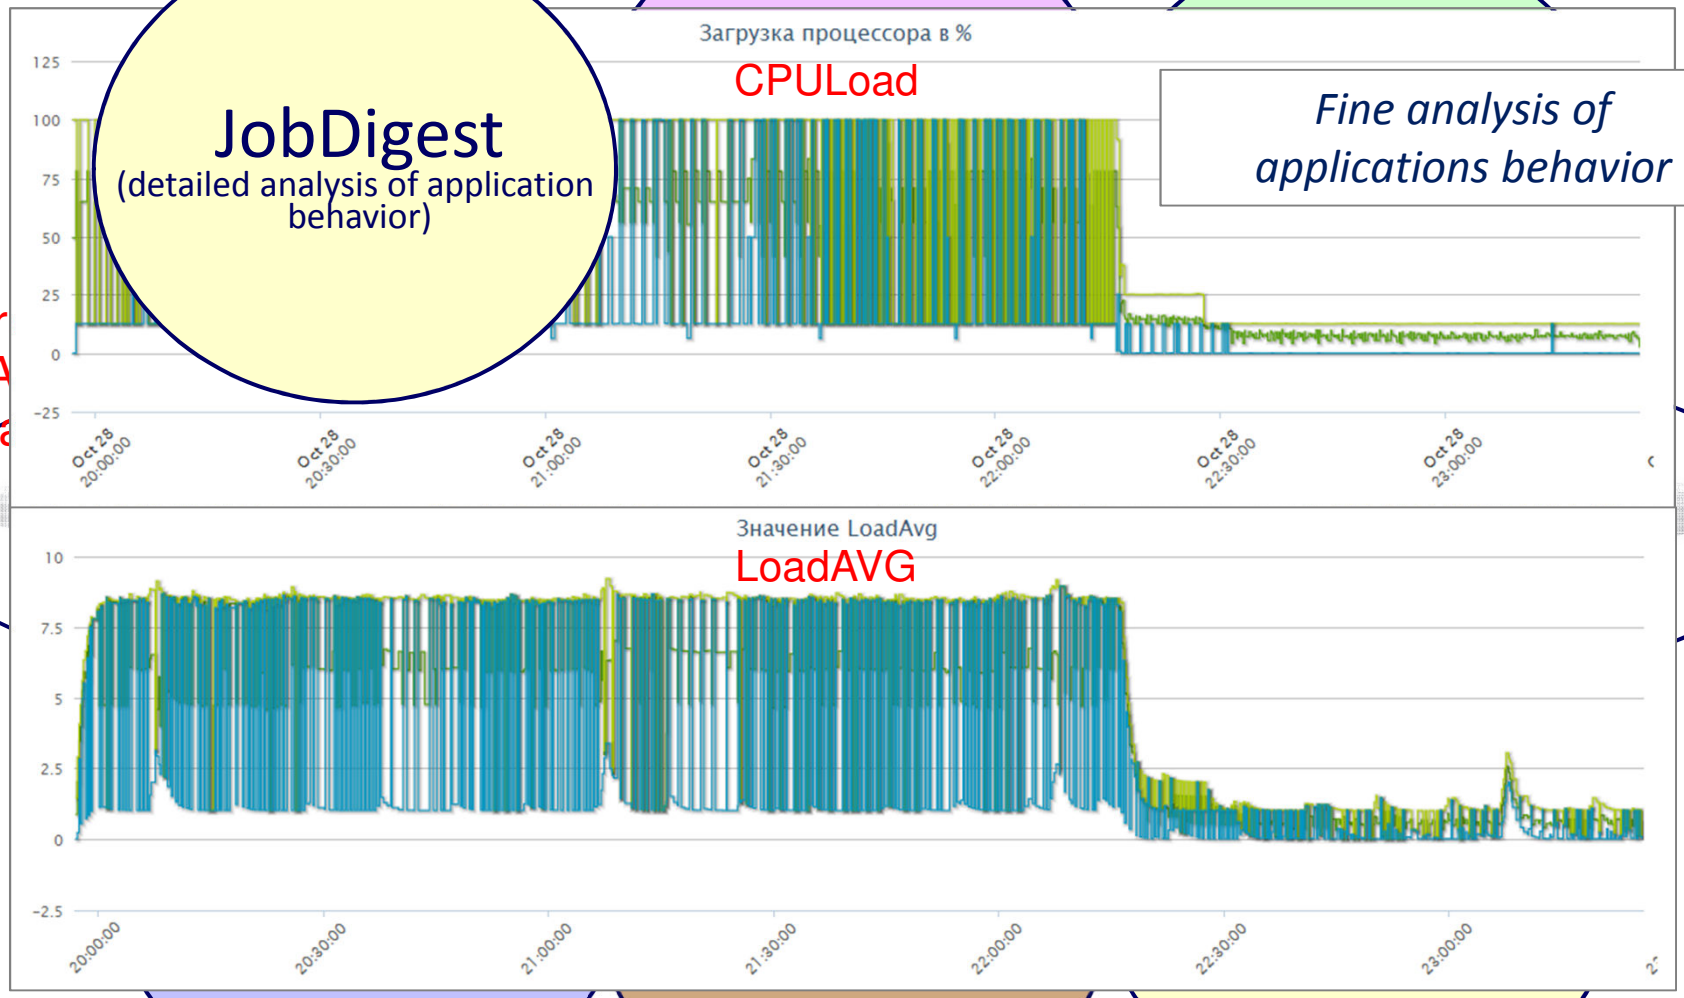

Fine analysis of applications behavior

User
SysA
Mana


Exascale Applications and Software Conference
epcc
THANK YOU
April 21, 2015, Edinburgh