Total Monitoring of Exascale Supercomputers

Why monitoring is necessary ?

- Very high degree of parallelism in HW&SW (Cores, processors, accelerators, nodes, hardware & Software components, files, indexes, buffers at data storage, traffic within interconnects, users, projects, processes, threads, running and queued jobs…).
- High energy consumption in small volume (one cabinet on the new “Lomonosov-2” MSU supercomputer: 420 Tflops and 120 kWt).
- Not only to detect but also to predict failures of supercomputers components.
- Many decisions about control over HW&SW of supercomputers must be taken automatically. Moreover, the larger system, the more autonomous supercomputer should be.

A new facet: monitoring of applications efficiency.
- very high degree of parallelism in HW&SW + sophisticated architecture of modern supercomputers = very low efficiency of supercomputing applications.

Fine analysis of different partitions
(control over everything!)

Fine analysis of supercomputing applications efficiency
(control over everything!)

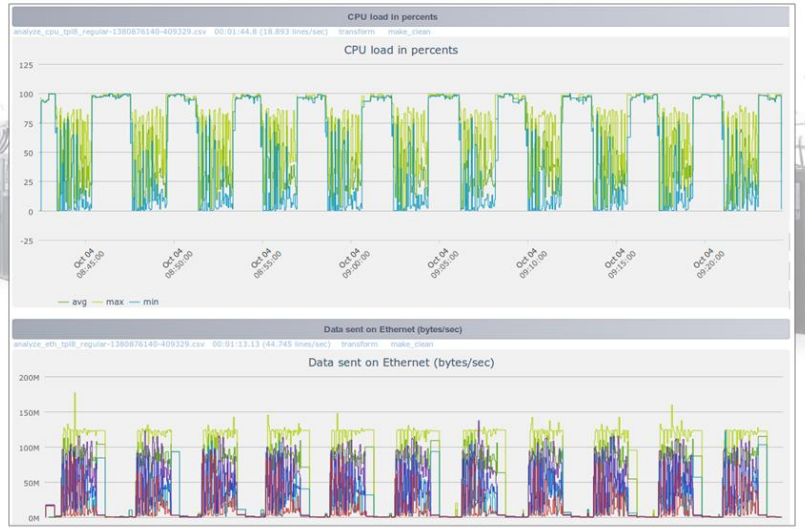

Efficiency of Supercomputing Centers
(straightforward approach)

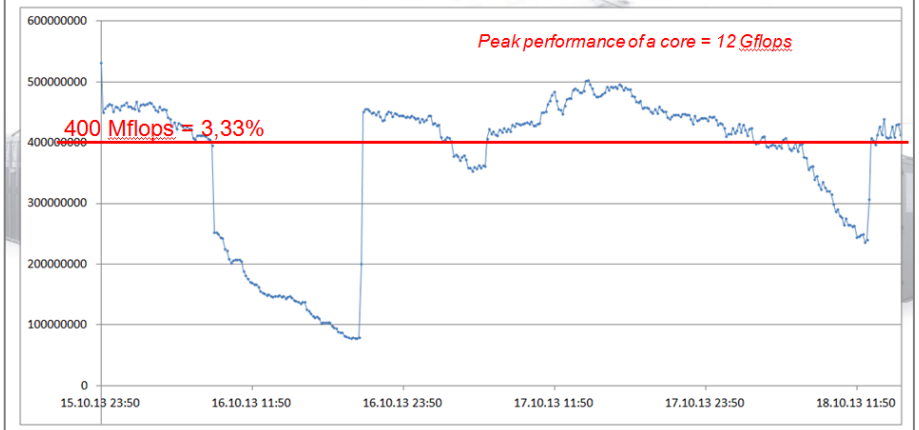

Average performance (one core) of "Chebyshev" supercomputer for 3 days

Total control: cost of delay…

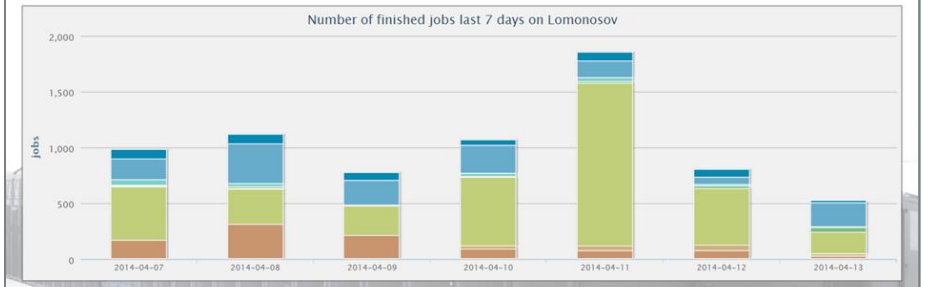

Supercomputer "Lomonosov":
- *about 1000 completed jobs per day,*
- *approx. 200 running jobs all the time,*

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

We need to keep control over supercomputers!

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

Total Monitoring of Exascale Supercomputers

==> We need the total control over HW&SW components of supercomputers, where total monitoring is a key issue. Why is the total monitoring really hard ?

Many different HW&SW components operating simultaneously
Many different applications running simultaneously,
Many different goals of monitoring…

Monitoring system, requirements:
- we need to know: what, where, when.
- scalability: thousands computing nodes,
  dozens sensors per node,
- high frequency: a few seconds and less,
- active and passive modes,
- …

A traditional approach (store all data first, process necessary data later) for the "Lomonosov" supercomputer doesn't work:
initial monitoring data rate – 120 MB/c
3+ Pbytes/year… BigData comes…
A simple analysis of the monitoring data required minutes, hours, days…

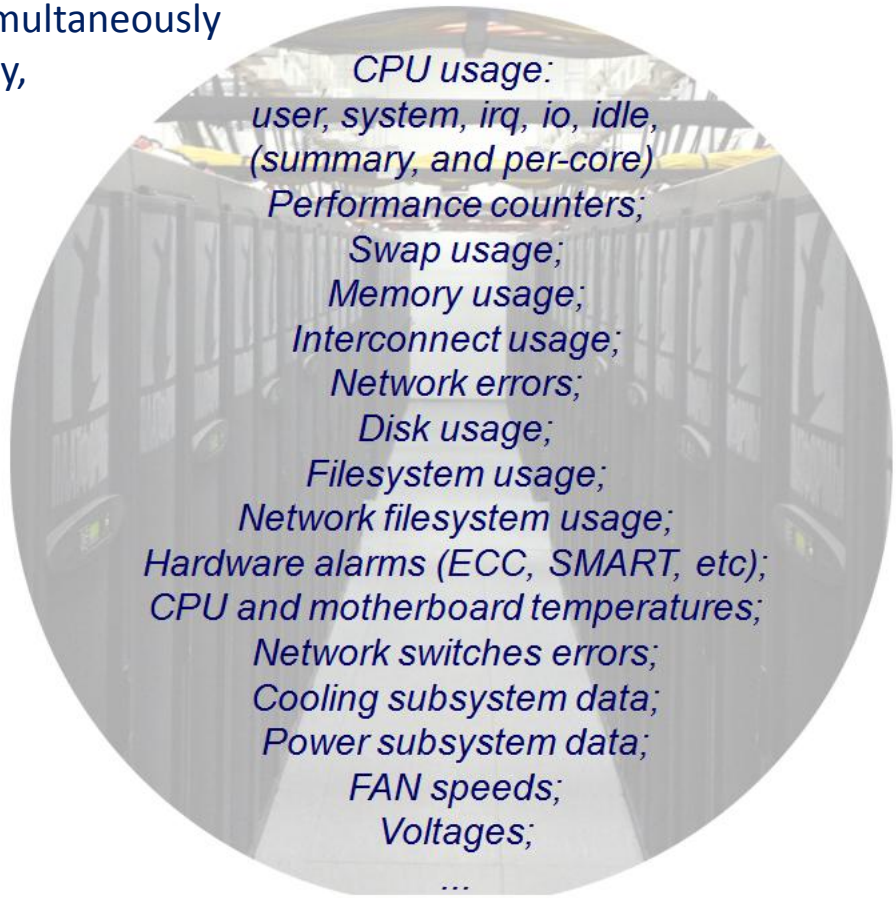

Too many possible reasons of performance degradation, but we need to collect and keep all this information

Total Monitoring of Exascale Supercomputers

Good questions:
Do we really need to keep all the data? Is “store first, process later” the best strategy?
Where is the best point to process data? Static/Dynamic? …

A smart approach to monitoring:
- on-the-fly analysis: all relevant information should be extracted from the monitoring data before it’s stored in a database;
- on-site analysis: monitoring data must be processed where the data were obtained (process first, move data (if necessary) later);
- 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.

No problem with monitoring of the whole “Lomonosov” supercomputer…
(estimations up to x100)

Two interesting statements inspired by the practice:
• What is BigData: characteristic of a certain problem or lack of our understanding (nature of data, structure of data, objectives of analysis…) ?
• If you have to deal with BigData, typically you don’t need the most of the data…