



Published on *CERN openlab* (<http://test-static-05.web.cern.ch>)

[Home](#) > Breakthroughs in security, efficiency, and performance with on-die hetero-processing

---

## **Breakthroughs in security, efficiency, and performance with on-die hetero-processing** <sup>[1]</sup>

**Date:**

Thursday, 28 April, 2016 - 10:00 to 12:00

**Location:**

[IT Amphitheatre](#) <sup>[2]</sup>

The Oracle database is the world's leading data management product. After the acquisition of Sun Microsystems in 2010, engineers from Oracle and Sun started work on a new category of microprocessor designed to process data several times faster, many times more efficiently, and qualitatively safer. This kind of goal cannot be reached by running software unchanged -- we needed to design new hardware and write new software at all levels of the system to utilize it. The approach we took, and are taking, exploits the following ideas:

- **Big is better:** Large scale computing systems give access to huge amounts of data without the costs of moving data between systems. This allows for larger tables, bigger sorts, fatter graphs, and more cloud tenants sharing the the same resource pool on SPARC systems that scale linearly in cost and performance from 8 to 512 cores, 64 to 4096 threads.
- **Secure is better:** Cache-line level memory access checking allows our instrumented memory allocators to manage memory at production speed while detecting bugs and reporting attacks in real time.
- **Information Density is better:** With hardware designed for scanning n-gram compressed, bit packed, dictionary and run-length encoded columnar data at full memory bandwidth, we make maximal use of every bit stored and every cache line transferred over the memory channels with no impact on performance.
- **Fast is better:** With hardware support for database operators running on specialized streaming processors, we can drive the memory channels at maximum rate, freeing up power and cores for running user computations on the result of these operators.
- **Connected is better:** Integrating EDR InfiniBand on-chip and on-board with low-latency, high-throughput, one-sided networking.
- **Portable is better:** By supporting platform independent acceleration APIs inside the database we can support a wide variety of acceleration techniques and give applications

and query planners the information to make the best use of the available hardware.

- Integrated is better: By supporting and accelerating multiple storage types (In-memory, NFS, NVMe, Exadata, HDFS, Fibre Channel), data formats (row major, column major, graph, JSON, spatial, MIME, Hive), algorithms, query languages, network protocols, and hardware platforms in a single product, we can share resources, increase usability and reduce the cost and the cognitive load in acquiring, storing, securing and understanding data.

In this talk, I will describe the experience that drives our acceleration priorities, the constraints and joys of the hardware-software co-design process, the HW features that resulted, and how software engineers have exploiting these features in ways we expected and ways we didn't. The industry has seen this approach used in the acceleration of linear algebra and computer graphics, and in this talk we'll see how we apply similar techniques to data processing but with changes to match the lower compute density of the problem space.

### **Indico or other event webpage:**

[For more information about the event](#) [3]

- [Visit Us](#)
- [RSS Feeds](#)

DISCLAIMER: This Web page contains pointers to material related to the management of CERN openlab in the Information Technology Department at the European Organization for Nuclear Research (CERN). Their use and distribution are regulated by the [CERN copyright notice](#).



---

**Source URL:** <http://test-static-05.web.cern.ch/news/breakthroughs-security-efficiency-and-performance-die-hetero-processing>

### **Links**

[1] <http://test-static-05.web.cern.ch/news/breakthroughs-security-efficiency-and-performance-die-hetero-processing>

[2] <https://maps.cern.ch/mapsearch/mapsearch.htm?n=%5B%2731/3-004%27%5D>

[3] <https://indico.cern.ch/event/515025/>