Evaluation of the LDC Computing Platform for Point 2

SuperMicro X6DHE-XB, X7DB8+

Andrey Shevel

CERN PH-AID

ALICE DAQ

CERN

10 October 2006

Purpose

- **Background:**
  - Test the machine (X6DHE-XB) as LDC with 6 D-RORCs
- **Machines:**
  - Elonex: Supermicro X7DB8+(2x Intel dual core Xeon)
  - Elonex: Supermicro X6DHE-XB (2x Intel Xeon)
- **Evaluation program:**
  - Linux installation
  - DATE V5 installation
  - D-RORC throughput with DATE
  - Perform a range of runs with internal D-RORC generator and receiving through DATE
  - Quick comparison of two machines based on mentioned motherboards.

Supermicro X6DHE-XB

CPU	2x Xeon 2.8 GHz, 2 MB
Chipset	E7520
FSB	800 MHz
Memory	4 GB (max 16GB), DDR 333
I/O	2x PCI-X (64/133) 4x PCI-X (64/100) 1x PCI-e (x4)
Network	2x GbE, BCM5721
Disk	74 GB, SATA
Graphics	Rage XL 8 MB
Periphery	2x USB front, 2x USB back Serial, PS/2
Chassis	4U, 1x 550W

Software Installation

- **Linux**
  - SLC4.3 with kernel 2.6.9-42.0.2.EL.cernsmp
    no specific problems but X7DB8+, USB boot successful, BIOS ok
  - To use D-RORC modules the BIOS setting needed to be changed: all PCI slots were set to 100 MHz
- **DATE 5.21 installation**
  - MySql
  - mysqltcl
  - BWidget
  - SMI
  - DIM
  - Drivers: rorc and physmem

  have been installed as rpm with additional scripts:

  DATE_basic_Install-SLC4x.bash, DATE_User_Install-SLC4x.bash, DATE_DriversInstall.bash
  - No specific problems

Common info about measurement

- **Total main memory volume = 4 GB**
  - 1 GB for Linux (in /etc/grub.conf)
  - other 3 GB for physmem
- **SLC 4.3, Kernel 2.6.9-42.0.2.EL.1.cernsmp**
- **6 cards D-RORC rev 4 (internal data generator).**
- **All measurement were done with DATE 5.21. No received event data were recorded to HD. The data from the database were analyzed with scripts. The pictures were built up with gnuplot.**
- **In total there were performed 390 measurement runs. Because the figures (rates, bandwidth) are fluctuating every run was continued at least 5 minutes.**
- **Two points on the further graphs: random event size with max=10**4 and max=10**6 Bytes were measured during two days each.**
- **All measurements with randomly distributed event size were performed with event checking in DATE.**
- **D-RORC Page Size was 1.5*10**5 Bytes.**

Example of editDb for the cards

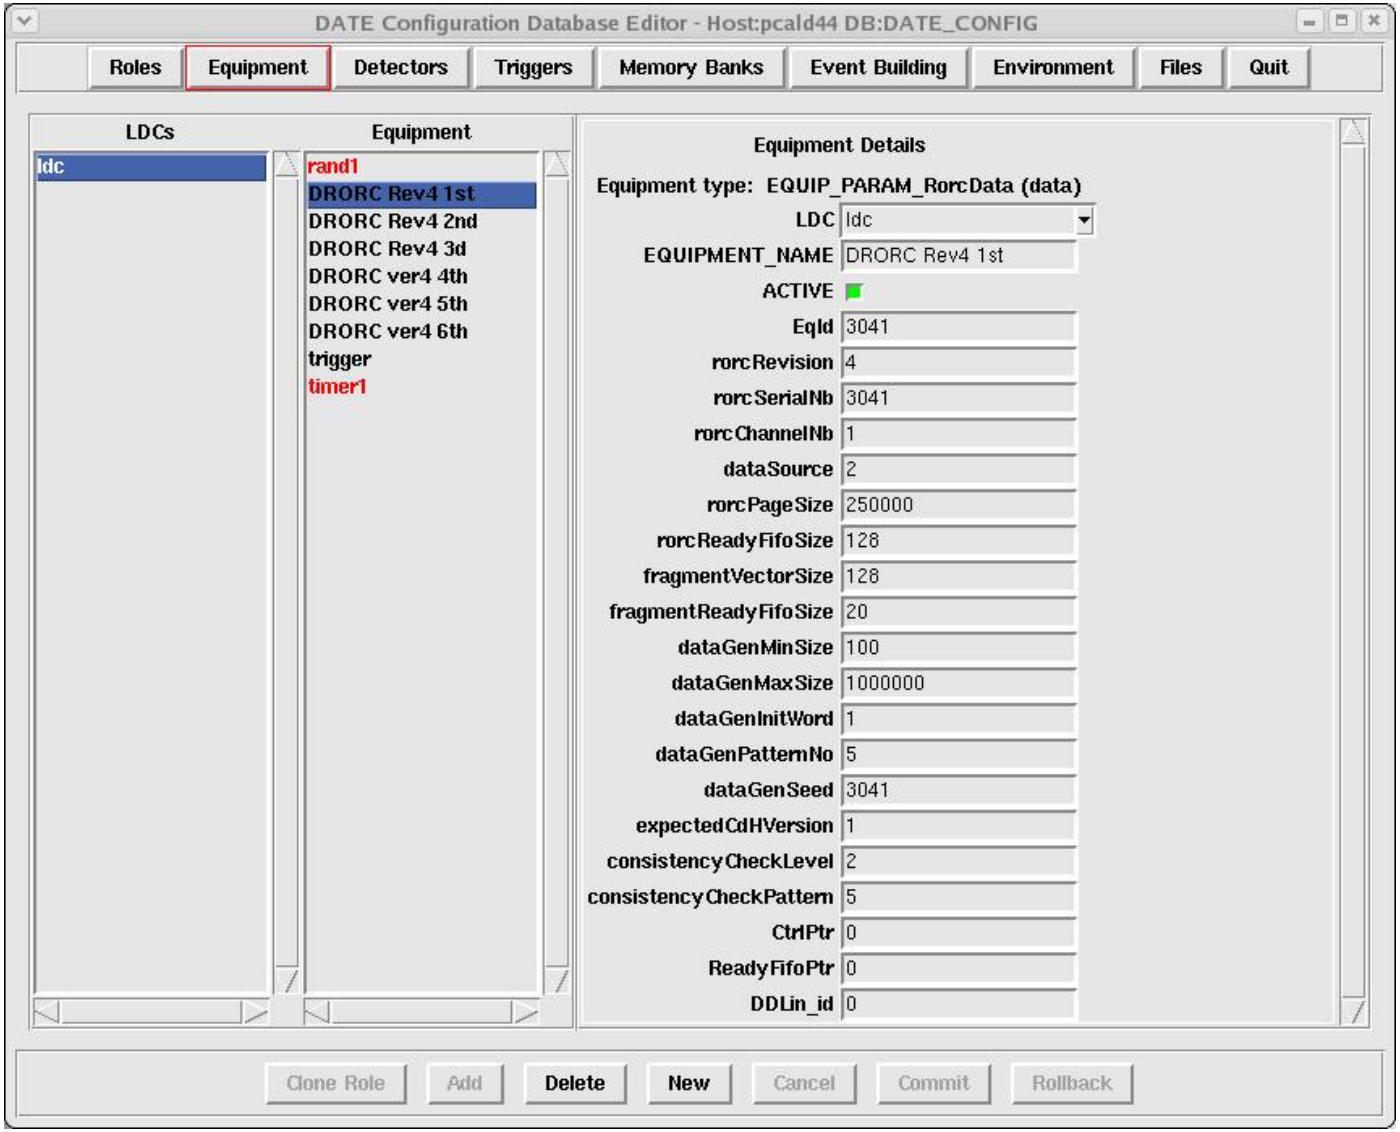

Architecture

6x D-RORCs

SN3042

SN3038

SN3041

SN3036

SN3040

SN3033

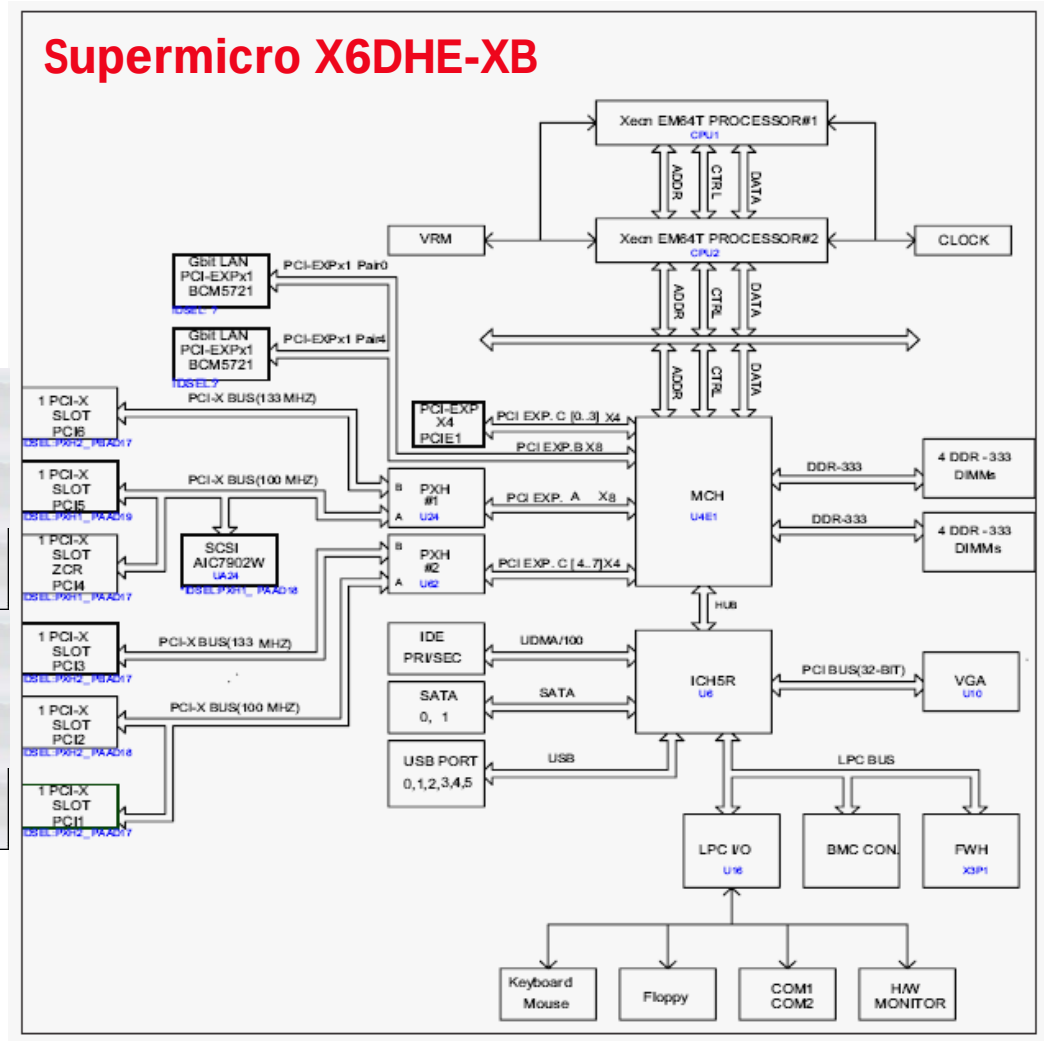

One D-RORC card: readout event rate

Fixed event fragment size

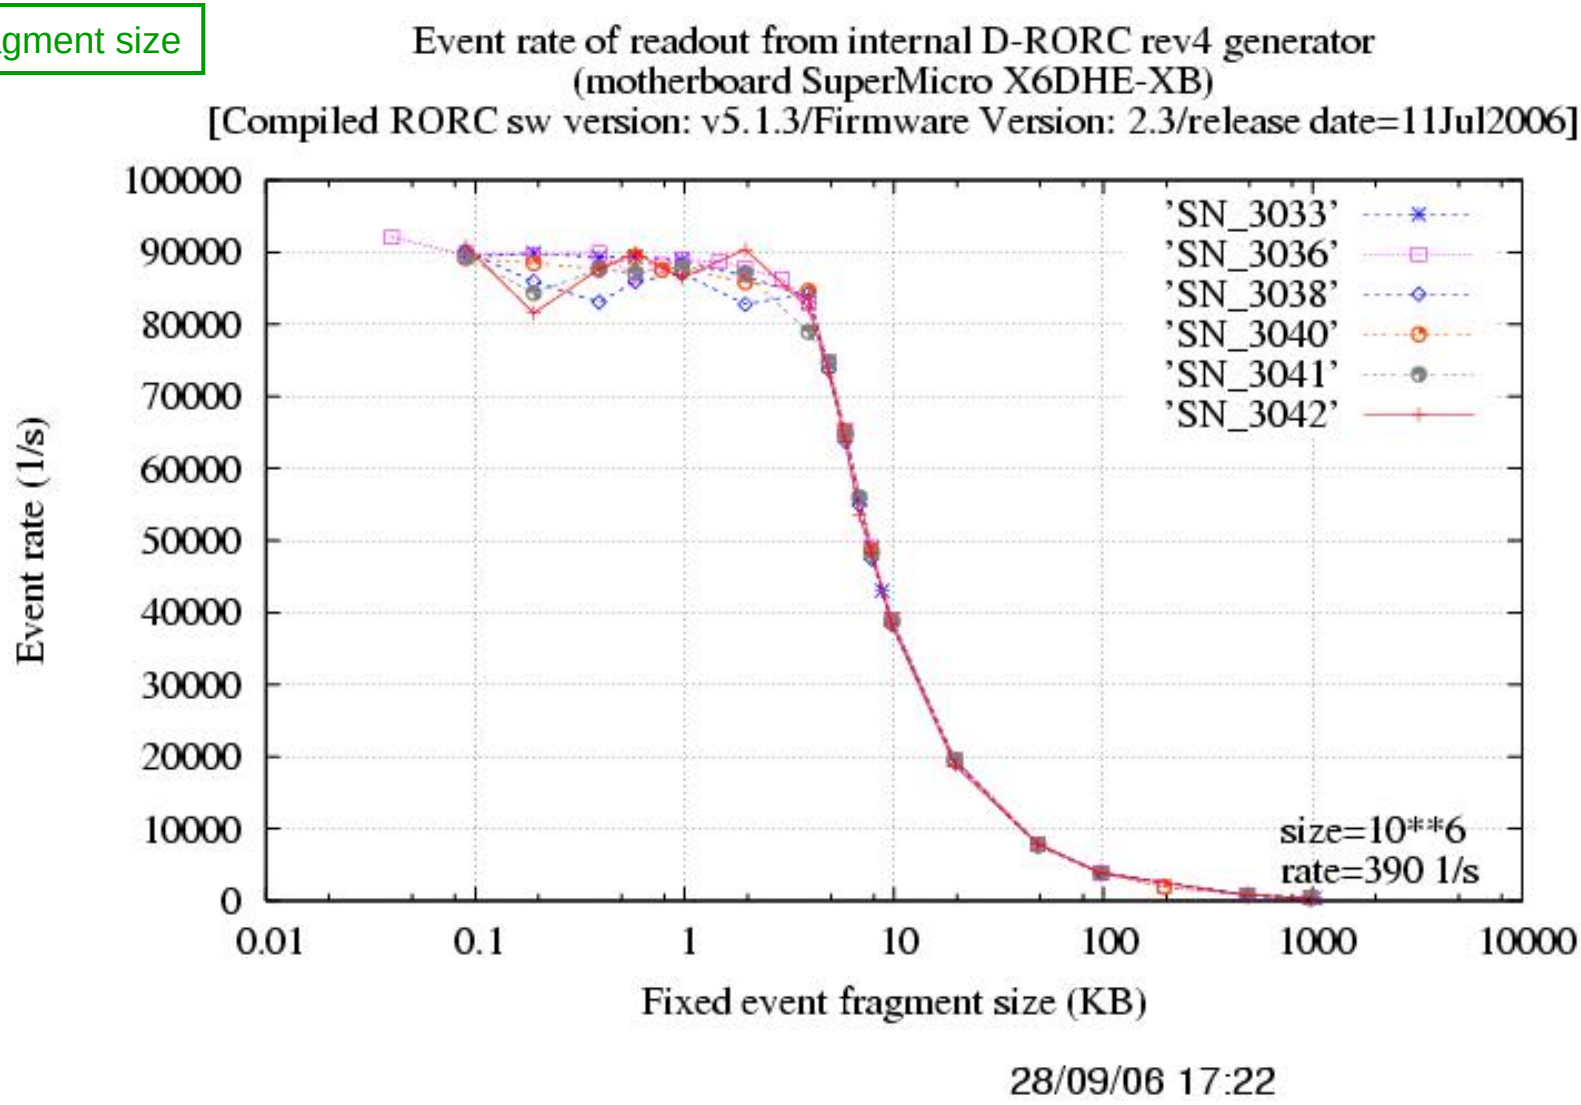

Testing the machine (X6DHE-XB) with 6 D-RORCs, X7DB8+

One D-RORC card: Bandwidth of readout

Fixed event fragment size

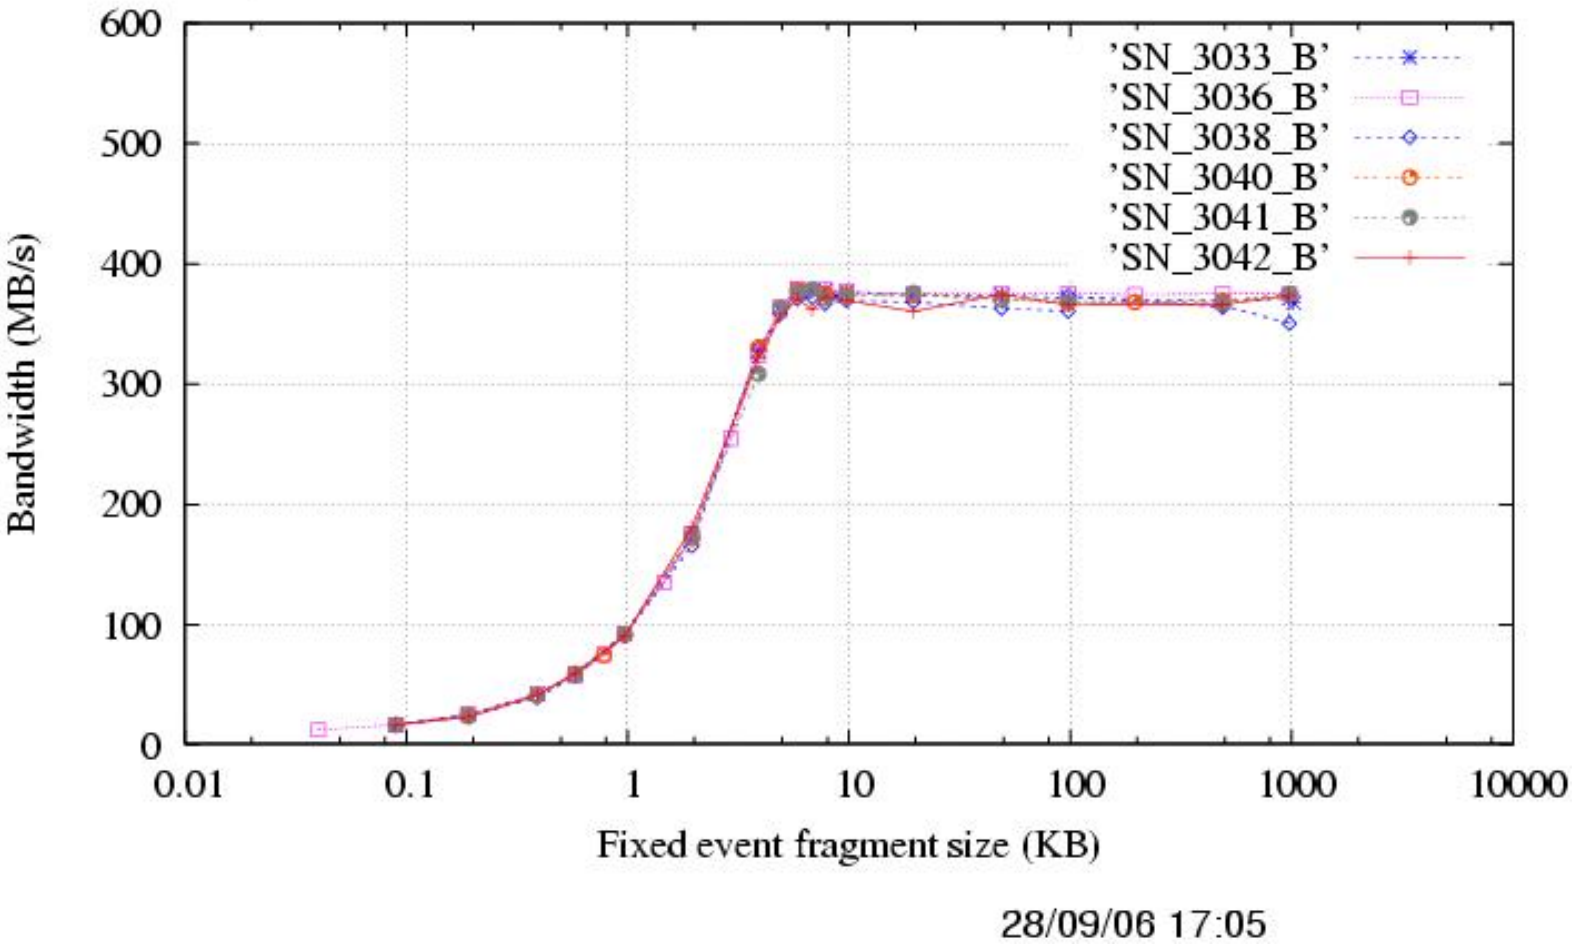

Testing the machine (X6DHE-XB) with 6 D-RORCs, X7DB8+

Multiple D-RORC cards: Readout Event rate

Fixed event fragment size

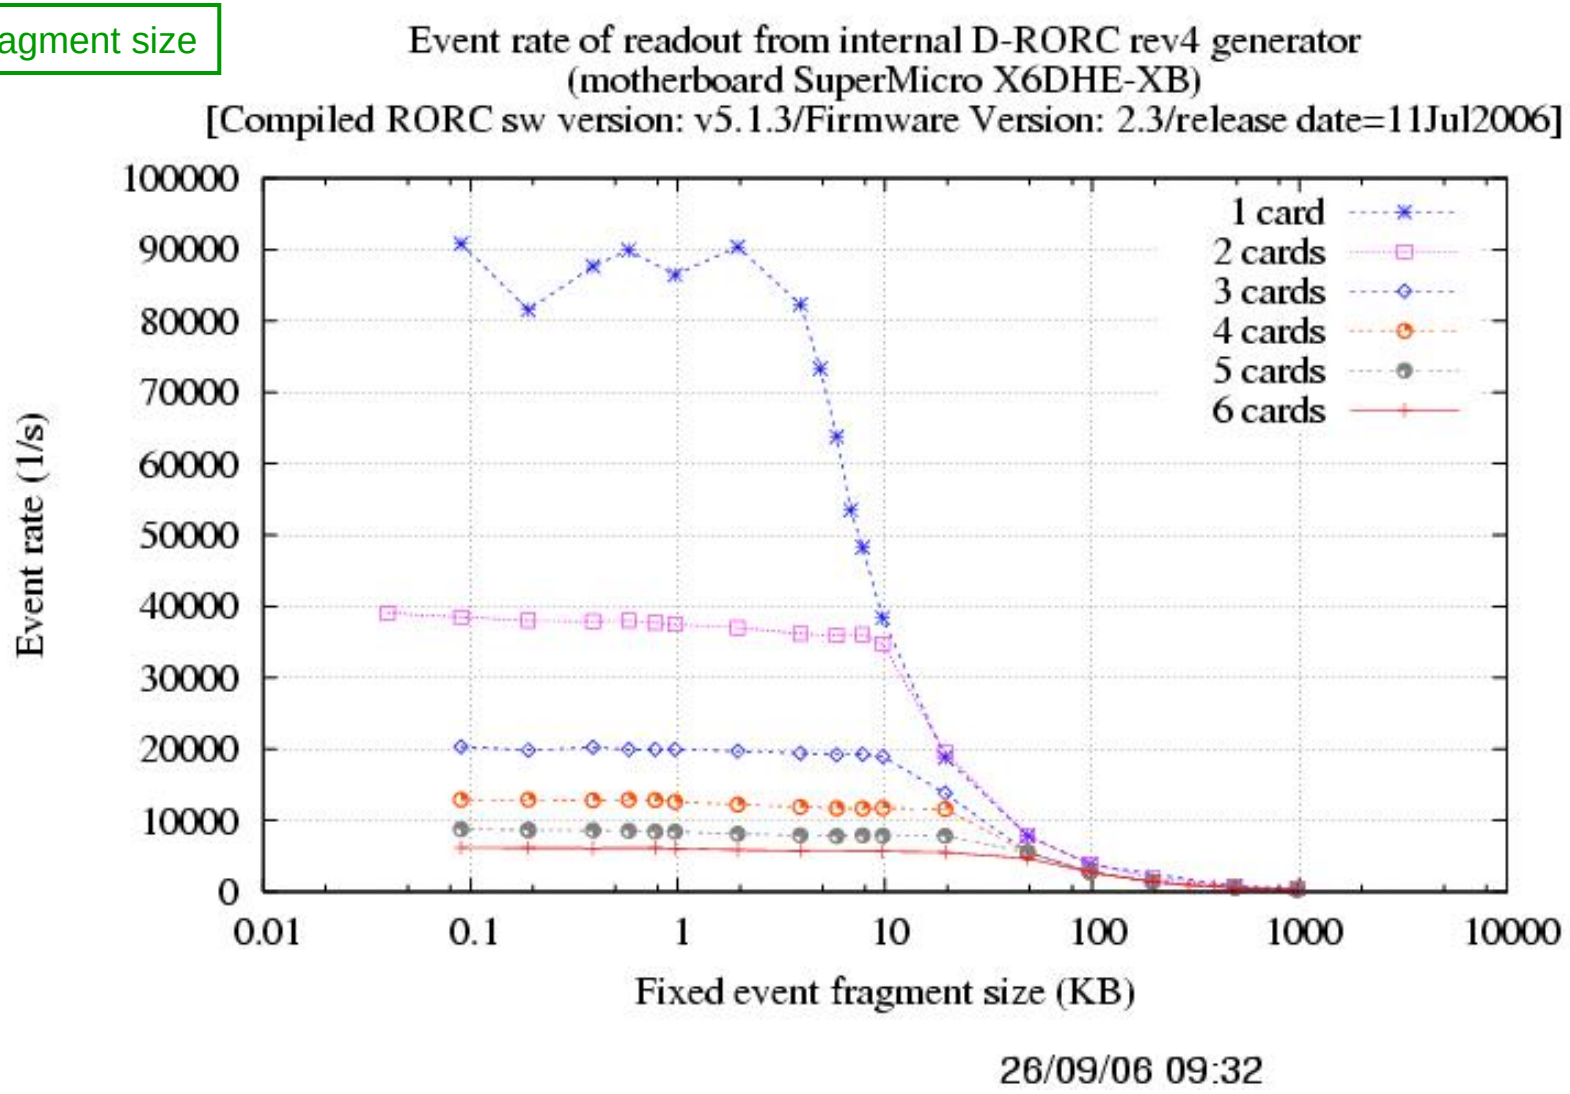

Testing the machine (X6DHE-XB) with 6 D-RORCs, X7DB8+

Multiple D-RORC cards: Bandwidth of readout

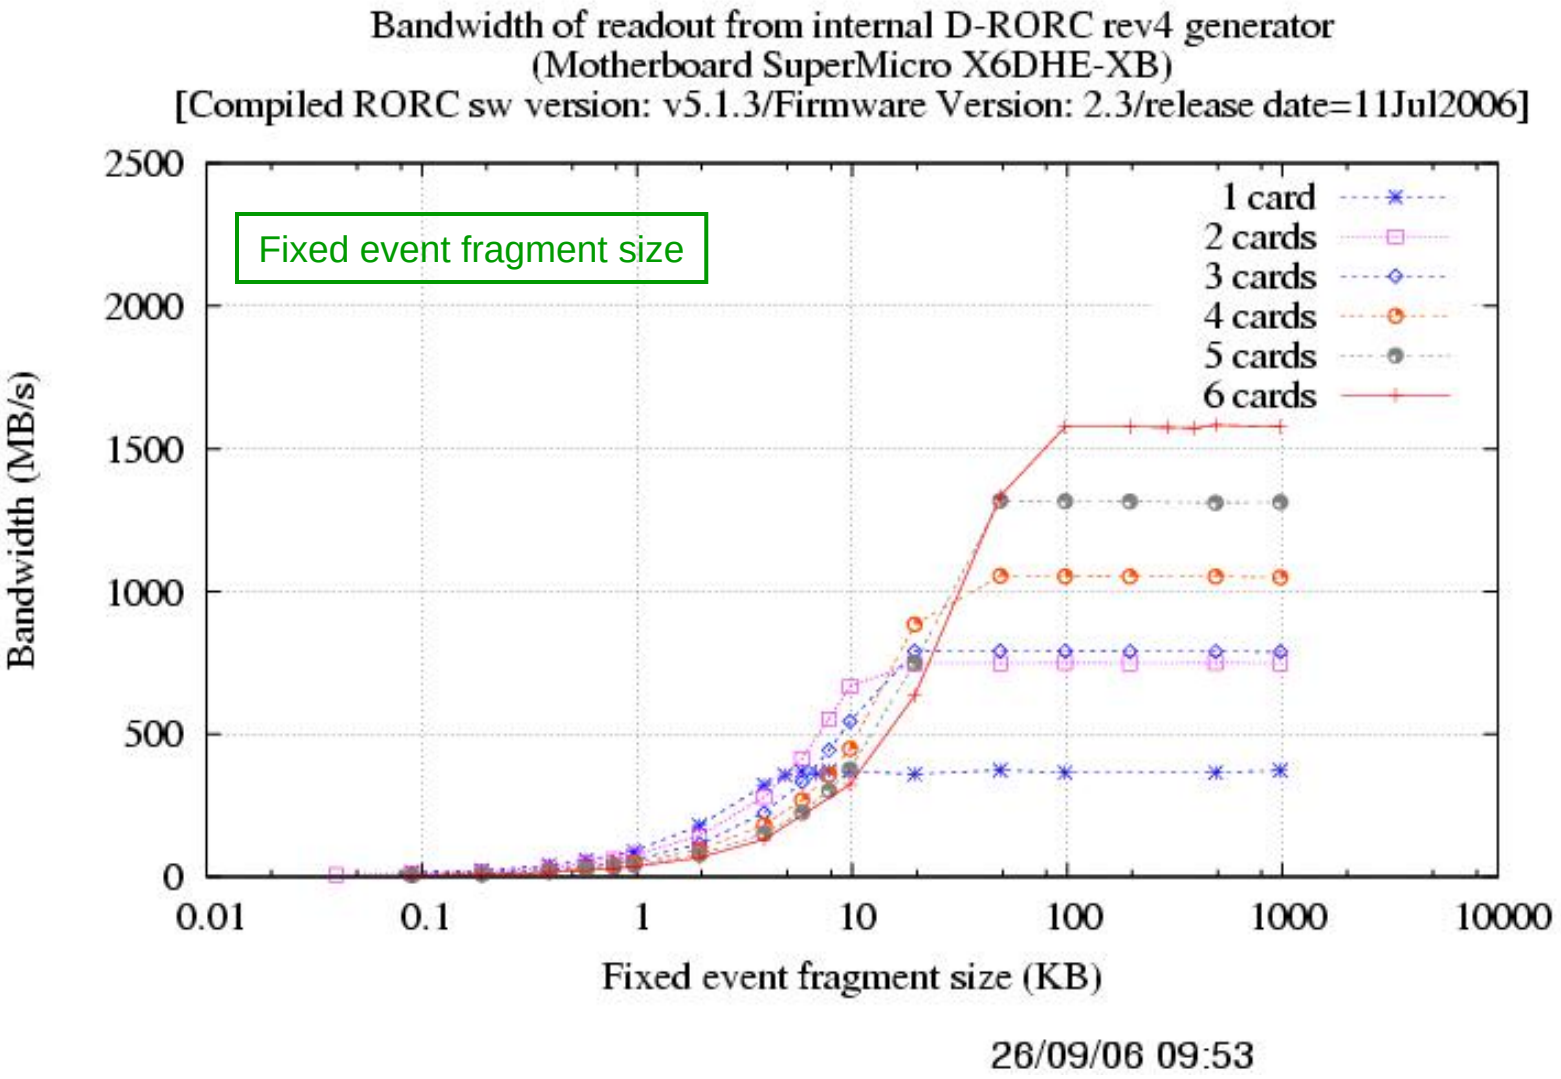

Testing the machine (X6DHE-XB) with 6 D-RORCs, X7DB8+

6xD-RORC cards: Readout Event rate

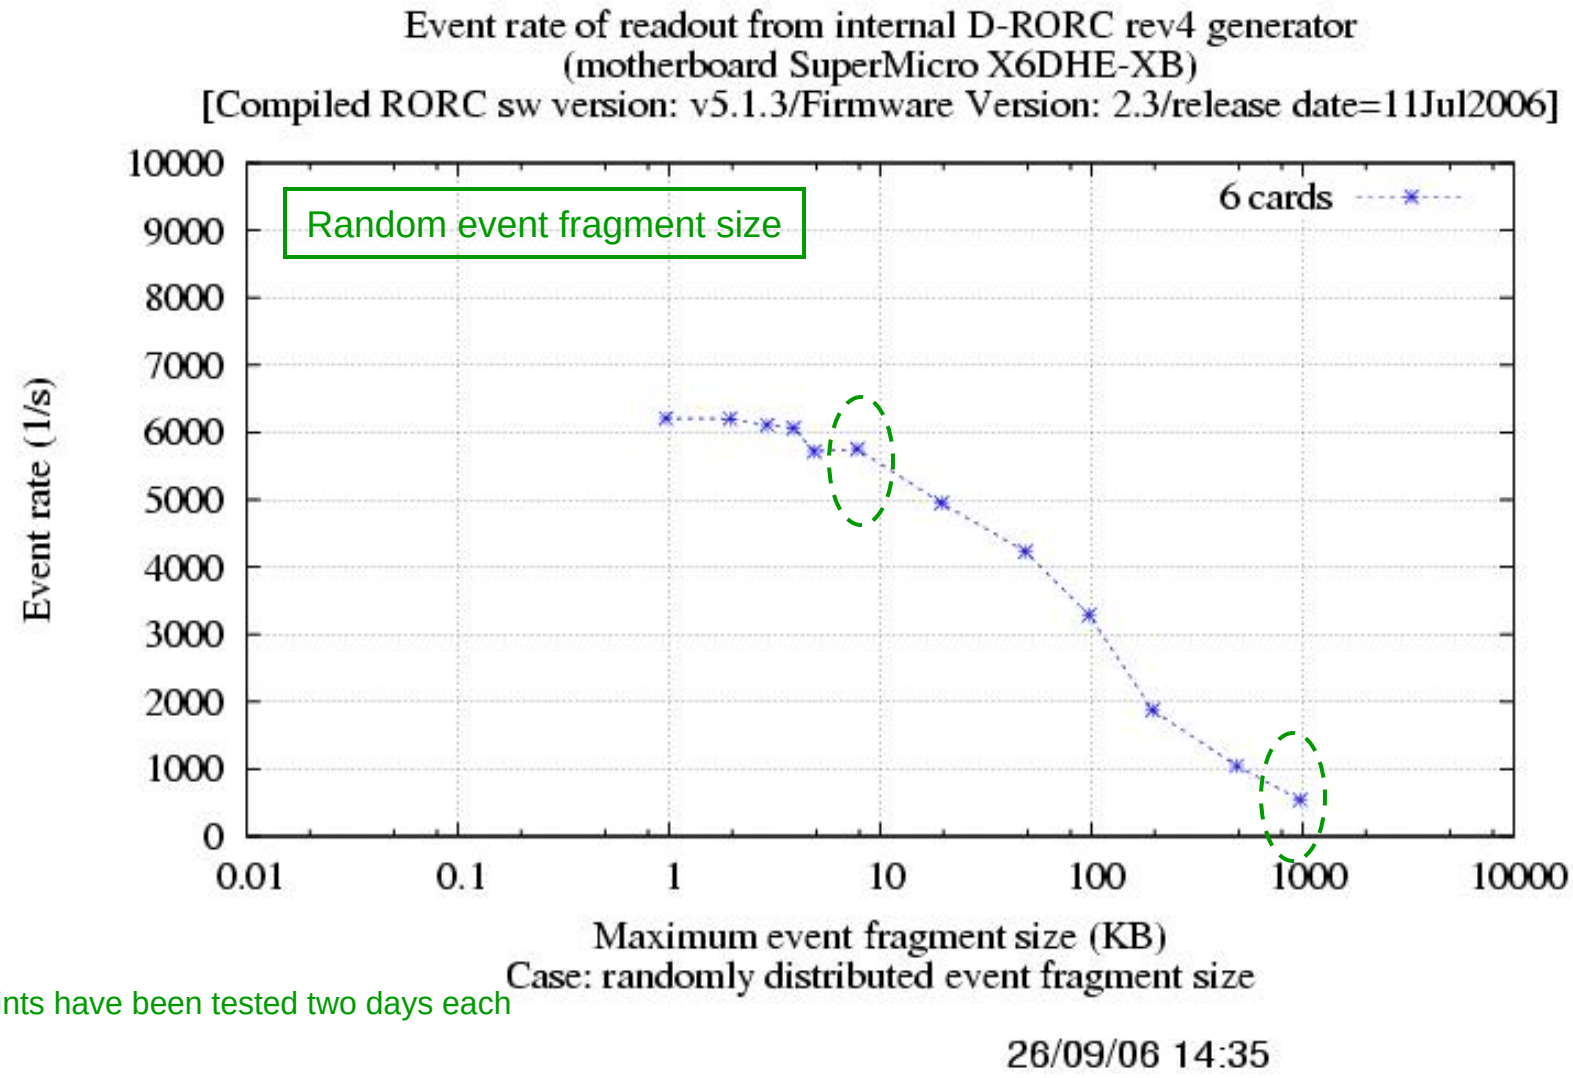

- Marked points have been tested two days each

Testing the machine (X6DHE-XB) with 6 D-RORCs, X7DB8+

6xD-RORC cards: Bandwidth readout

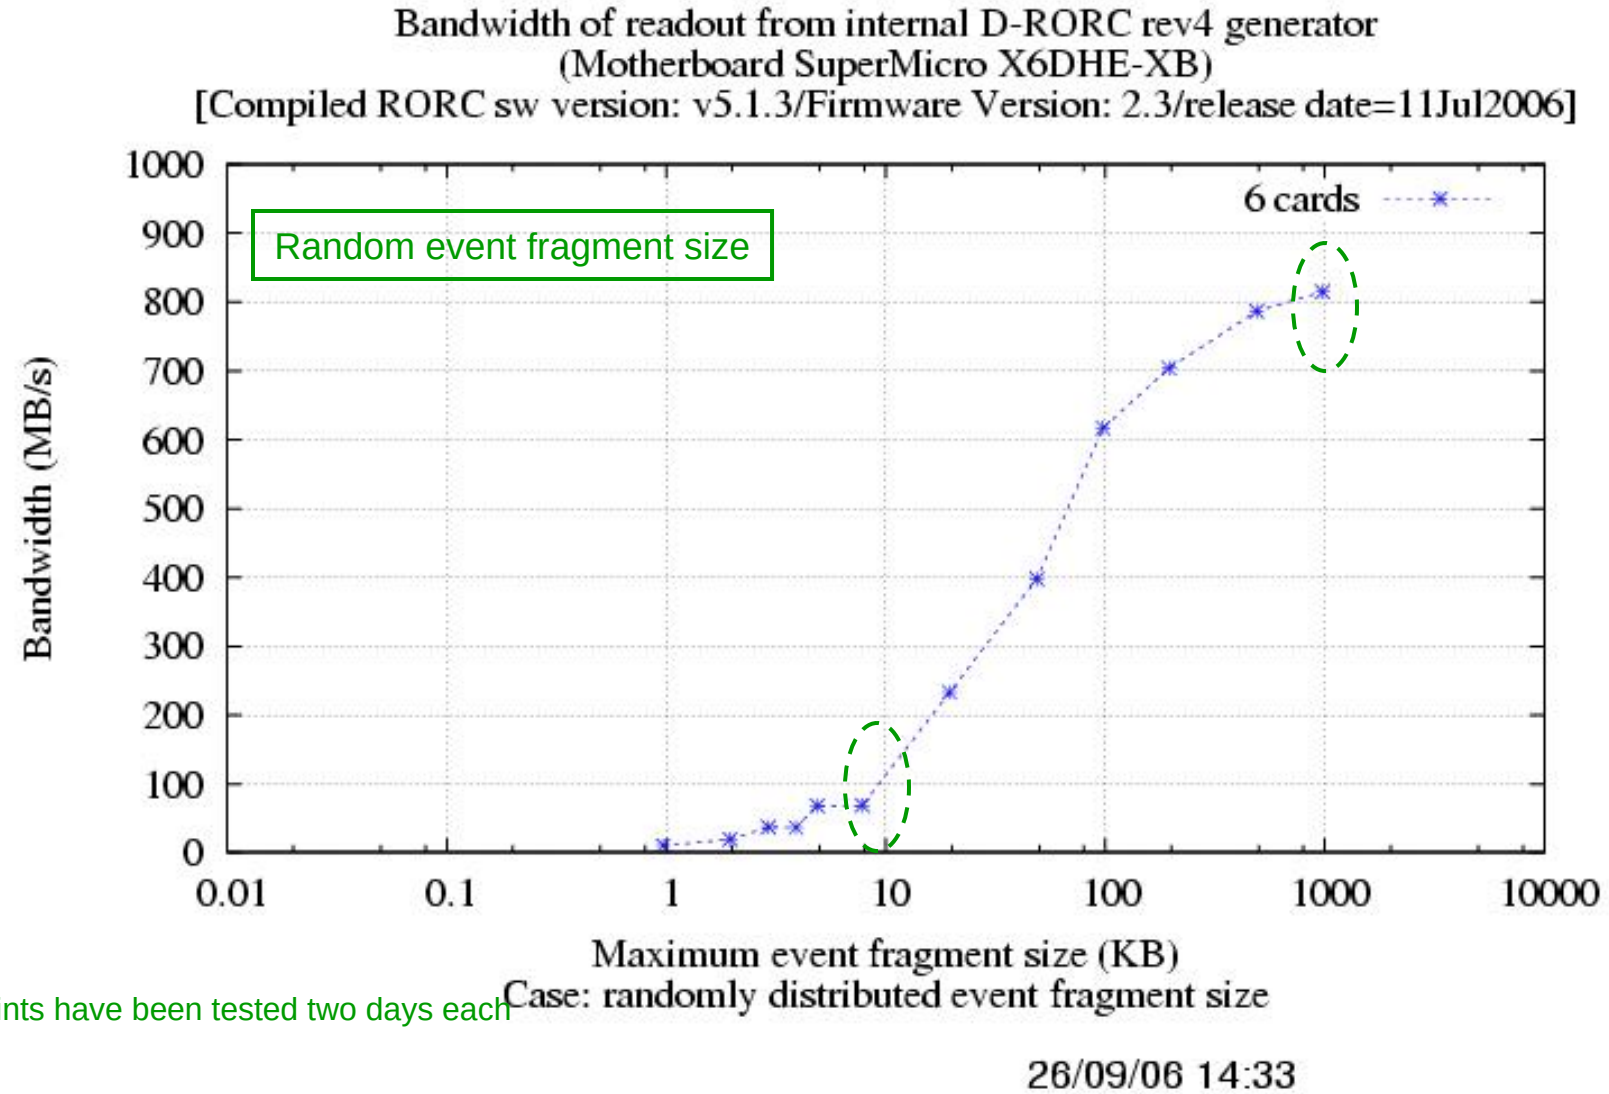

- Marked points have been tested two days each

Testing the machine (X6DHE-XB) with 6 D-RORCs, X7DB8+

Event rate: Influence of D-RORC page size

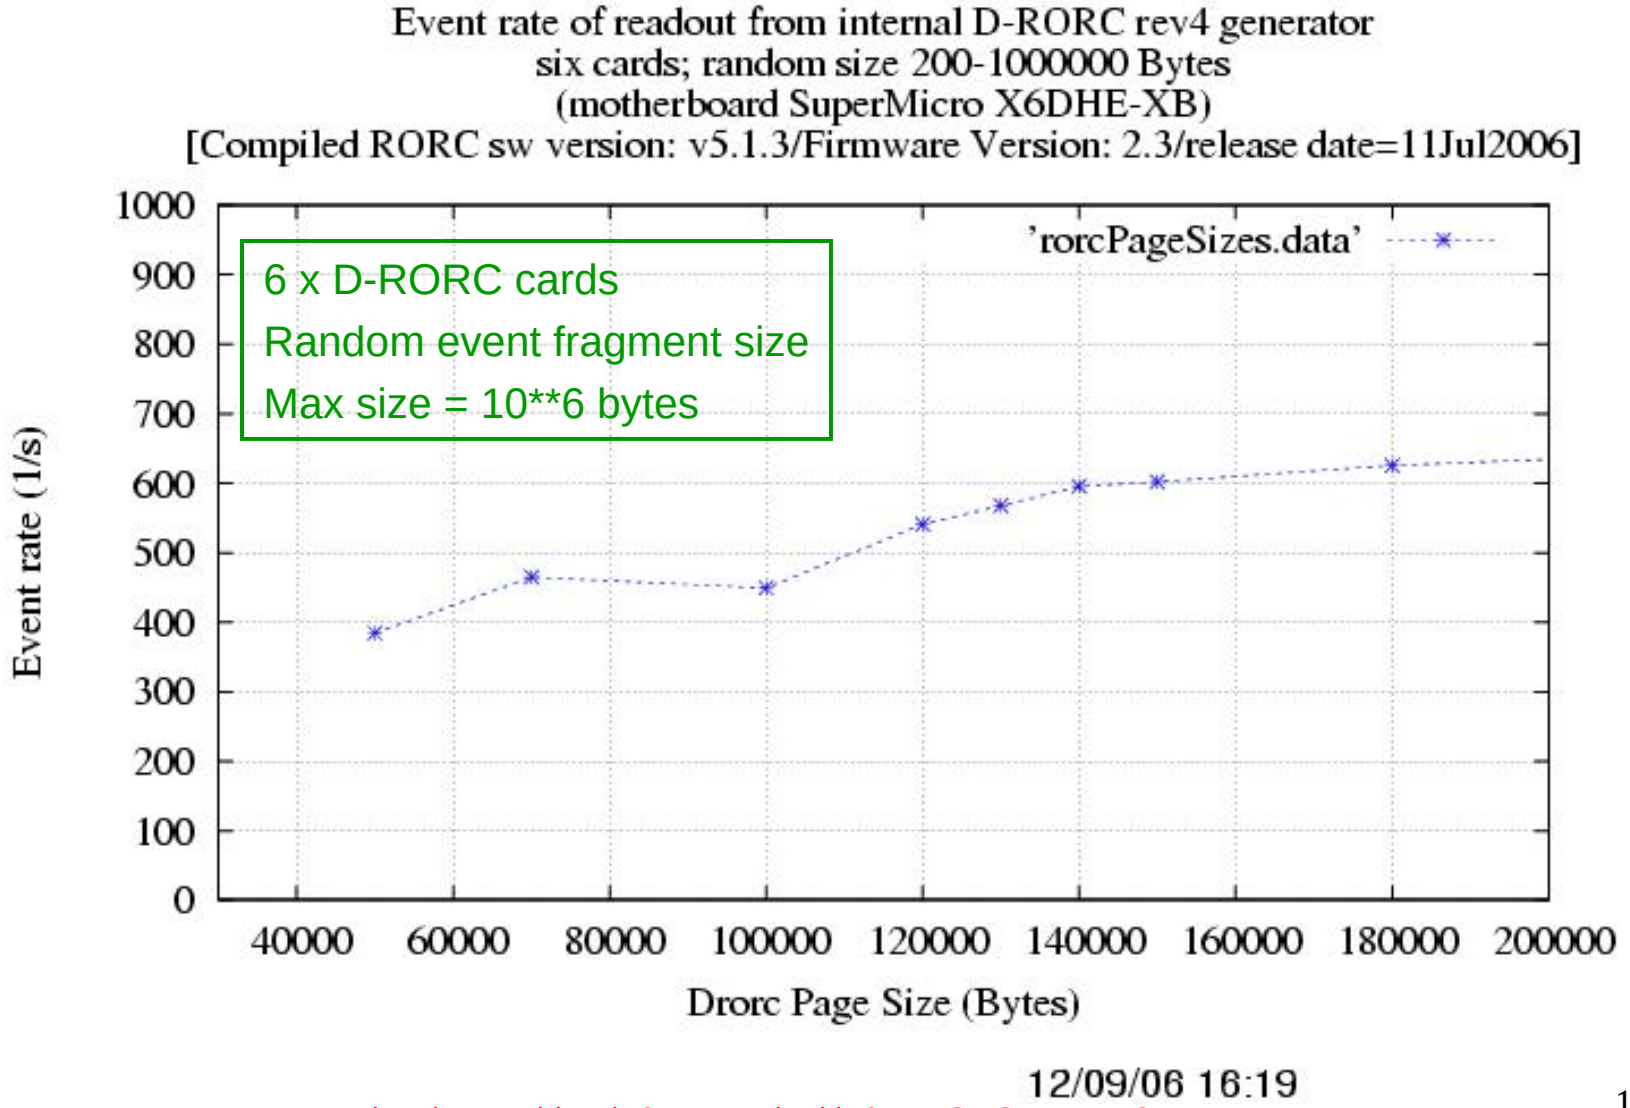

Testing the machine (X6DHE-XB) with 6 D-RORCs, X7DB8+

Bandwidth: Influence of D-RORC page size

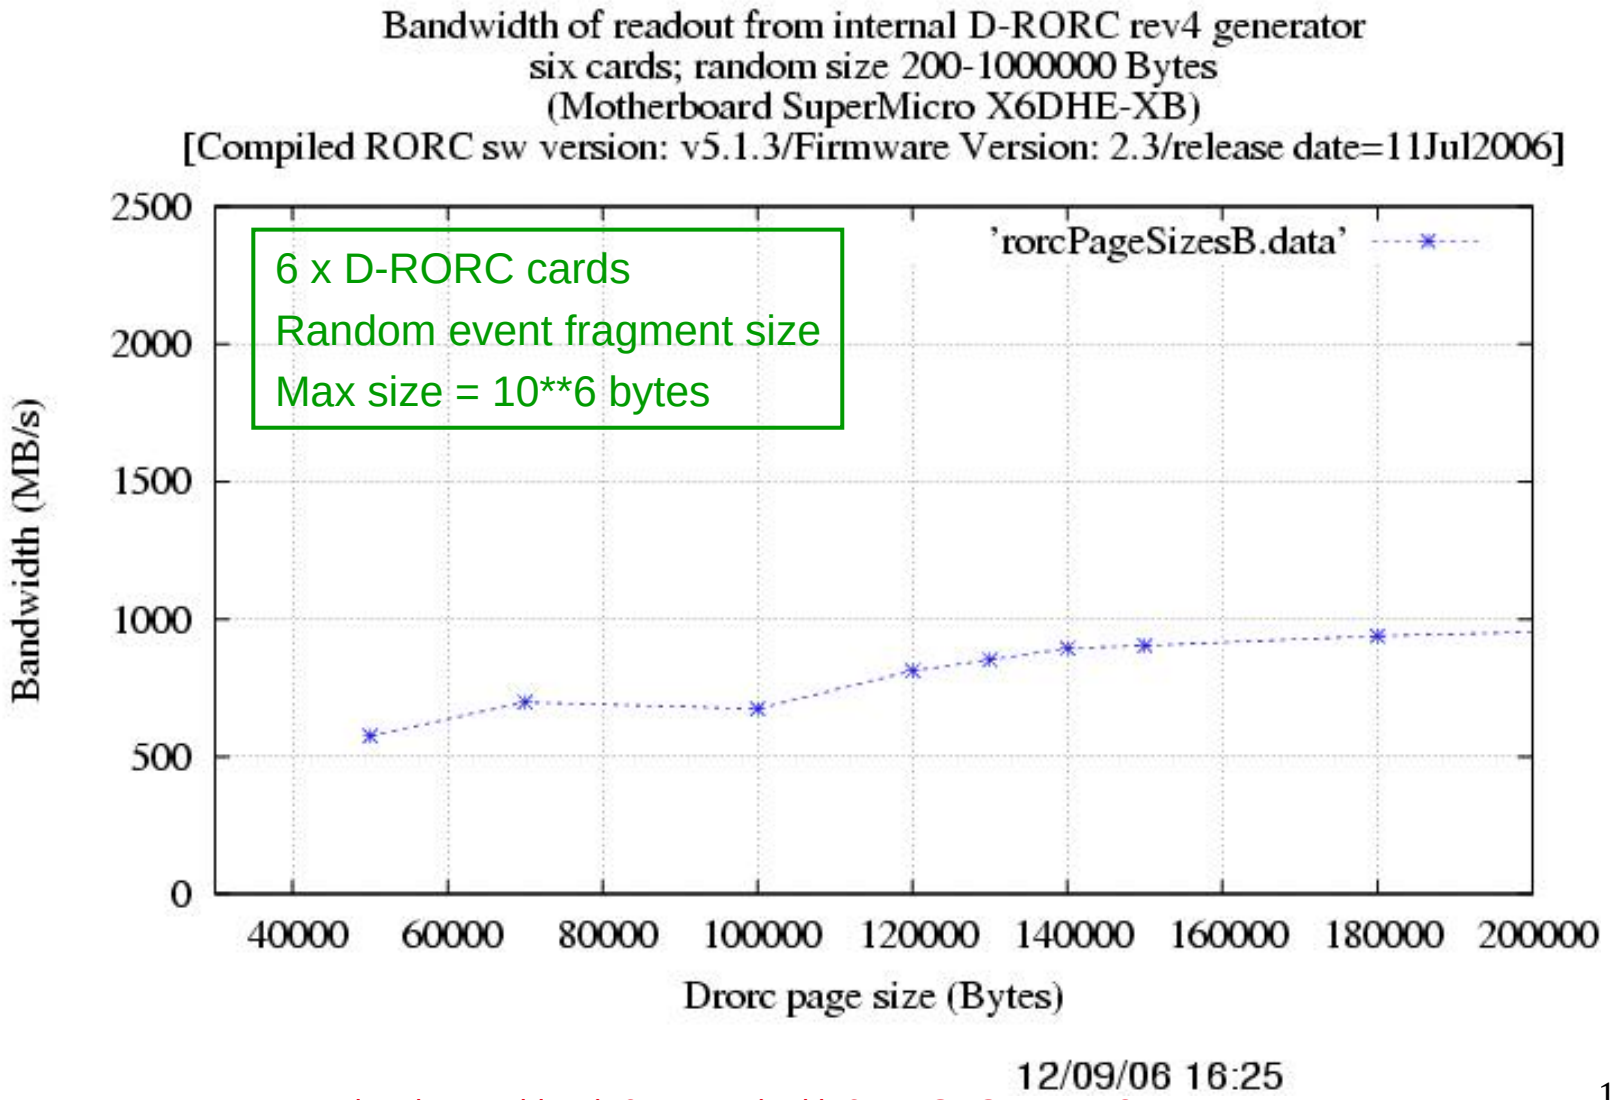

Testing the machine (X6DHE-XB) with 6 D-RORCs, X7DB8+

Test of Dual-Core Supermicro X7DB8+

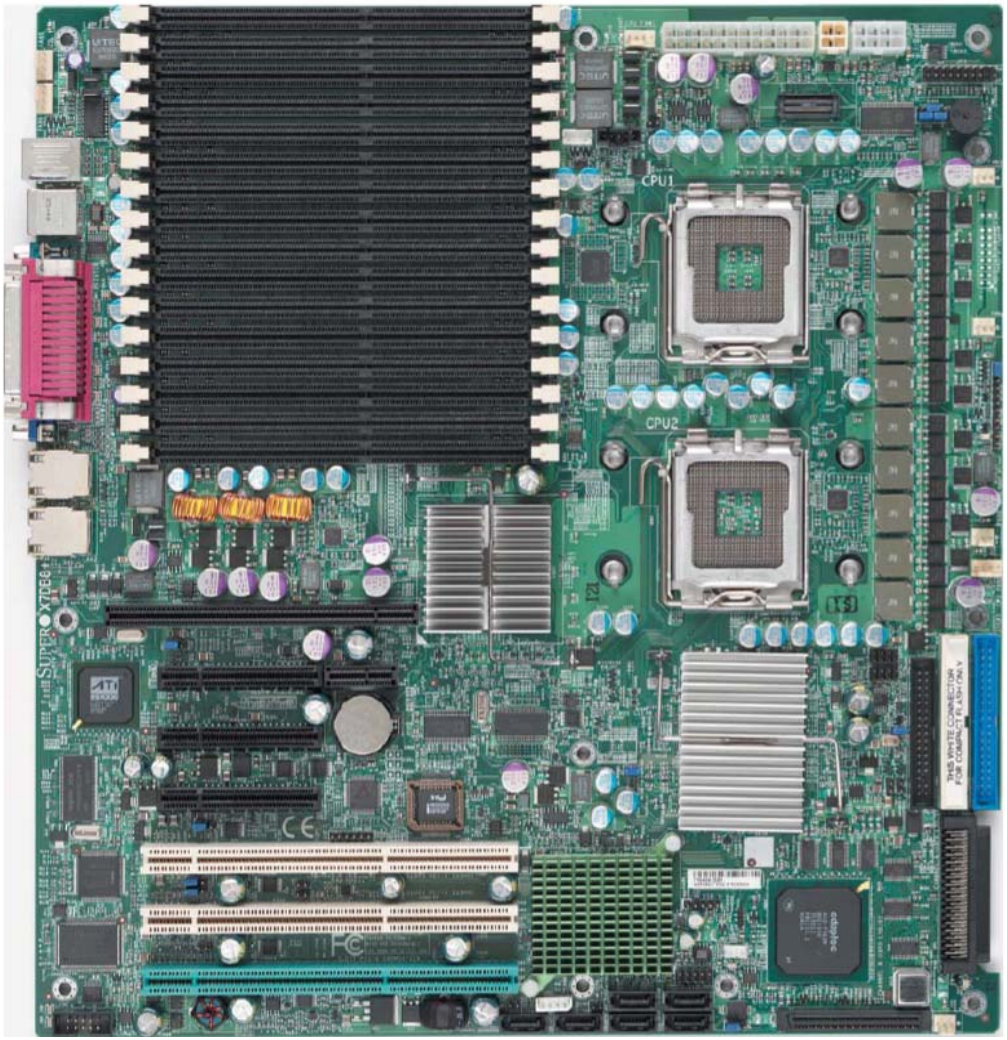

CPU	2 x Xeon LGA 771 pins (Dempsey) **dual core, cache - total 4 MB**
Chipset	Intel 5000P (BlackFord)
FSB	1066 MHz
Memory	4 GB (max 64GB), 667/533 MHz,DDR2
I/O	2x PCI-Express slots (2x X8;1x X4) 3x PCI-X (1x 64/100; 2x 64/133) 1x PCI-e (x4) SIMLP IPMI 2.0
Network	2x GbE, BCM5721
Disk	120 GB, SATA
Graphics	Rage XL 8 MB
Periphery	2x USB front, 2x USB back Serial, PS/2
Chassis	5U, 1350W PSU 3+ 1

X6DHE-XB vs X7DB8+

CPU	**2x Xeon** 2.8 GHz, 1 MB
Chipset	E7520
FSB	800 MHz
Memory	4 GB, DDR 333
I/O	**2x PCI-X (64/133)** **4x PCI-X (64/100)** 1x PCI-e (x4)
Network	2x GbE, BCM5721
Memory bandwidth (mem2mem)	1.5 GB/s

CPU	**2x Xeon** LGA 771 (Dempsey) **dual core**, 2.33 GHz, 4 MB
Chipset	Intel 5000P (BlackFord)
FSB	1066 MHz
Memory	4 GB, DDR2 533 MHz
I/O	2x PCI-Express slots (2x X8;1x X4) **3x PCI-X (1x 64/100; 2x 64/133)** 1x PCI-e (x4) SIMLP IPMI 2.0
Network	2x GbE, BCM5721
Memory bandwidth (mem2mem)	3.8 GB/s

Recommendation and conclusion

- **Recommendation: if you need to allocate D-RORC cards in the machine and like to have maximum bandwidth for X6DHE-XB:**
  - it is better to avoid to use PCI slot 1 & PCI slot 2 at the same time for high bandwidth information sources.
- **Conclusion**
  - In the test it was shown bandwidth (1.6 GB/sec) when all 6 D-RORC cards are in operation and event fragment size is more 10**5 bytes.
  - Randomly distributed event fragment size with maximum 10**6 bytes gives total bandwidth 814 MB/sec during two days run.
  - The machine **X6DHE-XB** is excellent to be used as LDC with 6 D-RORC cards. Selected a single-core LDC for the ALICE DAQ
  - The machine **X7DB8+** might also be used as LDC with max 3 D-RORC cards. First test of a dual-core LDC from Supermicro. Only 3 PCI-X slots. The 6 slots version will be tested when available.

Thank you!

Klaus Schossmaier, Sylvain Chapeland, Ulrich Fuchs and other crew of the PH/AID