Blog Export: c0t0d0s0.org, http://www.c0t0d0s0.org/

Tuesday, December 16. 2008

NFS at 2 GByte/s

Brendan Gregg published an follow-on article to his 250.000 IOPS article. In "Up to 2 Gbytes/sec NFS" he shows how far he can drive the Sun Storage 7410 in a single-node configuration. I like his approach to talk about an benchmark-special result at first, to show it´s pitfalls afterwards just to present a more realistic number. As in the last test, he factored out the harddisks by using a workingset within the size of the main memory (100 GB on a 128GB system)

He was able to get up to 2 GByte/s from a 7410 with two 10 GBit/s interfaces and 20 clients. (with some upside potential, as the CPUs weren´t loaded). In a more realistic test he was able to yield a little bit more than 1 GByte/s over a single 10 GBit/s interface.That's 1.07 Gbytes/sec outbound. This includes the network headers, so the NFS payload throughput will be a little less. As a sanity check, we can see from the first screenshot x-axis that the test ran from 03:47:40 to about 03:48:30. We know that 50 Gbytes of total payload was moved over NFS (the shares were mounted before the run, so no client caching), so if this took 50 seconds - our average payload throughput would be about 1 Gbyte/sec. This fits.

10 GbE should peak at about 1.164 Gbyte/sec (converting gigabits to gibibytes) per direction, so this test reaching 1.07 Gbytes/sec outbound is a 92% utilization for the 7410's 10 GbE interface. Each of the 10 client's 1 GbE interface would be equally busy. This is a great result for such a simple test - everything is doing what it is supposed to.An excellent result!

Posted by Joerg Moellenkamp in English, Solaris at 13:21

1 Gbit/s over a single 10 Gbit/s interface - how impressive

Mind the bytes vs. the bits.
Anonymous on Dec 16 2008, 15:16

Aber du weisst schon, das das GB fuer Gigabyte steht .... nicht fuer Gigabit? Ich habs fuer dich aber trotzdem mal ausgeschrieben ....

Anonymous on Dec 16 2008, 18:22

Ich denke mal es ging um den Satz "In a more realistic test he was able to yield a little bit more than 1 GBit/s over a single 10 GBit/s interface.".

Dort ist das 1 Gbit - hoffentlich - ein Tippfehler.
Anonymous on Dec 16 2008, 18:58

Ich habe extra aus Deinem Text zitiert. Ich weiss nicht, was daran eindrucksvoll sein soll, 1 GBit/s über ein 10 GBit/s Interface auszugeben.

Und GBit/s steht m. E. für GBit/s - nicht für GByte/s.
Anonymous on Dec 16 2008, 19:08

Klar ist das ein Tippfehler ... wie auch aus dem zitierten Text hervorgeht Habe ich korrigiert.
Anonymous on Dec 16 2008, 19:22

Wie auch schon im zitierten Text ... es handelt sich um Gigabytes ... nicht um Gigabits. Der Fehler ist korrigiert.
Anonymous on Dec 16 2008, 19:23