

Tuesday, September 22. 2009

Updated Sun Storage 7410

The Sun Storage 7410 got a little update. You can get it with up to four six-core CPUs, thus up to 24 cores in total.

Posted by Joerg Moellenkamp in English, Oracle, Solaris at 16:39

Hey, Joerg:

This is someone who worships Sun software engineering.

Sun really has to dig deep down why they expect someone to pay 74K for a Dual Socket Istanbul System with 64GB. Even with startup essentials and say 50% off list price, it is still way off there.

These days, I can get Dual Nehalem-EP L5520s with 72GB ram from Dell for \$5k now. Tag on a couple of X25-E 32GB as ZILs and a couple of X25-E 64GB as L2ARCs, and 22x Western Digital RE3s, the whole system is only \$10K or so. And yes, it is running OpenSolaris.

So, there is no wonder why Sun is losing 100 Million a month, since their MBAs working at Sun doesn't understand shit about Volume Economics.

Seriously, this is a friend trying to tell you that you are just pissing off potential customers. Drop the Stec Dipshits and UltraStars and start designing for volume economics, which is what ZFS is about right? Jeff Bonwick understands the power of volume economics, why can't Sun's marketing department?

Look at the following specs; this is what people want:

Dual Socket Nehalem-EP L5520 with 72GB ram
Intel X25-E ZILs, Intel X25-E L2ARCs
\$20K price point. (I am giving Sun a 100% premium over market price for this one)

And, lastly, stop screwing around with ridiculous list prices. Be transparent like Dell and HP, and you will sell a lot more units to amortize Jeff Bonwick's salary (worth every penny by the way)
Anonymous on Sep 23 2009, 01:40

You know about the existence of the 7110, 7210 and 7310 ???
Anonymous on Sep 23 2009, 07:00

Sure,

7110 is shit because of 8GB ram max.
7210 is a Thor rebadge, not bad in terms of number of drives, but the Dual AMD CPUs are certainly not Nehalem-EP grade. (Hint, it is time for Sun to get a Nehalem-EP Thor with IOP348s)
7310 is not bad, still way overpriced for a dual socket Opteron

What I want is a x4170 based Open Storage server with Intel SSDs as configurable ZILs and L2ARCs. Work with Intel on that. And hit the \$20K price target with ZIL and L2ARCs included.

Just a hint, look at last picture of this link:
<http://www.semiaccurate.com/2009/09/22/larrabee-breaks-cover-last/>

1 million IOPS 200GB+ SLC on a single PCI-e. That pretty much handles FusionIO and Sun F-20 hands down. (It will probably cost 2500 dollars considering how much the X25-E costs. It is 7x X25-Es raided over IOP348)
Anonymous on Sep 23 2009, 09:02

BTW, with Intel offering 1 Million IOPS on a single PCI-e card for \$3K, do you think the F5100 with 1 Million IOPS for 300K would have a market? I am asking this rhetorically.

Oracle can probably rip out the F-20 accelerators out of their Exadata v2, and replace them with the intel card, and get probably triple the performance...call it Exadata v3
Anonymous on Sep 23 2009, 09:14

Oh ... this card is nice ... but there are several points, that make this card problematic:

1. No super-cap or battery to protect the data in the caches.
2. I didn't found any data how their measured the 1 Mio IOPS. Sorry, but when they put the equivalent of 7 SSD on this board and we further assume 35k IOPS (that's the rate of the X25-E at 4k blocks then i'm at 245000 IOPS. Perhaps the 1 Mio is the speed of the DRAM cache or they used 1k blocks for benchmarking. Want more data.
3. The F5100 has 4 TB flash. How much has this baby?

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

You see, too much missing information just to draw a single conclusion and this baby looks like a prototype ... the F5100 is an existing product.

Anonymous on Sep 23 2009, 10:46

1. 7110 - do you really need more than 8 GB Cache in a 2/4 TB filer ?
7210 - why should there Intel IOP on the board ... the gear doesn't need RAID controllers. And realworld I/O-wise i wouldn't wipe Opteron from the desk.
7310 - have you compared this to a fullfeatured NetAPP?

2. You can build the el-cheapo filer on you own. Take OpenSolaris, the plain-vanilla server, the plain-vanilla storage and all your knowledge. This is the variant for people having time but no money. You have a S7000 minus the fishworks interface and all the knowledge outside solaris and the hardware.

Or you could buy the variant everything ready to go inclusive tuning and the appliancification of the device, QA, further development. That's the extra you pay for. That's the method for people with money, but no time. For the enterprise.

3. Of course you build an el-cheapo filer with other components: But the STECs for example have some advantages, like having 12000 sustained write IOPS instead of 3500 IOPS max. Of course you could use a RE3 but i don't think you get it cheaper than a UltraSTAR with all the contractual conditions between a Tier-1 vendor and a supplier.

4. I assume, that the SSD will cost vastly more than 2-3 thousand dollars, Lower volume, more components etc. And as i wrote at another occasion: To make this a product, you need much more components. You can't simply multiply the numbers ... that's armchair engineering, sorry ...

Anonymous on Sep 23 2009, 11:59

1. 8GB cache is a joke. No SSD ZILs or L2ARC for the 7110, so all the nice things that accelerate ZFS are not working in the 7110. So yes, if Sun thinks anything less than 64GB-72GB of ram would sell in their Storage 7000 series, I would be disappointed.

2. I have mentioned the components a couple of times. But Dell and HP gear isn't "el-cheapo". It is far cheaper than Sun, but it doesn't mean people who bought it have taken a step down in quality. In fact, behind the shiny new chassis, all of them are Supermicro gears underneath(I know what I am talking about) Sun still hasn't figured out that people who have the intelligence to understand ZFS's architectural beauty also have the intelligence not to pay one more zero on the bill for the hardware.

3. You obviously haven't taken the hint. I am talking about putting 4x of the Intel PCI-e in there as ZILs+L2ARCs in my next build. So much with Stec's 12K random writes when you have 80 times that on the Intels. As to the UltraStars, I have yet to see data showing higher reliability than RE3/RE4.

I know F5100 has 4TB, the Intel one has probably 200-400GB per card.(7x64GB or 7x32GB). You still don't get my point. You can put 9x Intel PCI-e in a HP DL370G6 and get the 3.6TB of flash and get nearly 9 million IOPS per 3.6TB of capacity. The intel card has higher IOPs per GB than the F5100. and lower \$ per GB or IO than the F5100.

I think I saw a line of SuperCaps on the bottom of the Intel PCI-e card right next to the IOProcessor. Right now, nothing is certain. We will just have to take a look at PCB when that beaut is released. I am sure if Intel promises 1 million IOPs on that card, you will get supercaps on it.

Anonymous on Sep 23 2009, 21:19

1. Of course more cache is nice, but you have to take into consideration, for which tasks an 7110 is used. And sorry to disappoint you ... AFAIK we sell the 7110 quite often ...

2. Oh ... i've read your suggestion of using the Intel SSD as using the "Intel X25E" not this prototype. I'm sure our engineering will have a look at that device when it's ready. But i'm not in fishworks engineering, thus i can't really comment on that. And if i knew more, i would write it here

3. HP cheaper: 1TB drive for DL320 G5p 609\$... just looked at HP store.Dell: The 1TB SATA Disk costs you 549\$ for the PowerEdge R710. Looked at online store a few minutes ago. Sun 1TB Hard Disk for X4275 \$659.00 at Sun Store ... sorry ... but that doesn't look much cheaper too me.

4. Sorry ... but the Sun gear isn't supermicro.

5. I didn't said that a RE3 or RE4 is less reliable than the UltraStar. I've just said i have my doubt that you can buy it cheaper from a Tier-1 vendor with all the contractual agreements between the vendor and the harddisk supplier.

5. I have my doubts, that those silver cans are power caps. They doesn't look like the super caps i know from WIMA, Maxwell (they are blue) or NESSCAP (they are black and grey)

6. I quite puzzled that you brag about a product that you just saw a picture and have no other data than the one number from Intel marketing. You extrapolate the costs just out of X25E multiplied by seven.

7. I did some research how they got to this numbers. The 1 Mio IOPS number isn't the number of one of this cards, you need 7 of them to get this number of IOPS. http://www.theregister.co.uk/2009/09/23/insane_ssd_performance/ . Furthermore the test was done on 4K blocks. 142957 IOPS per card ... roughly equal with the number you get from other storage controllers via PCIe. Finally this numbers make sense.

8. This is a card in a lab of an intel researcher don't expect it to be a product soon.

9. And don't make price estimation now. Just think about the list price of a Fusion io idrive and the larger ones use MLC not SLC

**** Anonymous on Sep 23 2009, 23:29

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

1. You guys sell the 7110 quite often, but not often enough to amortize your salary right? Not to be talking in a condescending way, but people who buy 7110 are idiots.
2. Look into it. I would be excited to see that Intel PCI-e in future Sun Storage systems.
3. That's why you guys are sick fockers, chargin 500-600 for a 1TB drive. Most people who use Dell and HP source their "additional" drives from newegg.
4. Pretty much all motherboards are sourced from Taiwan. They might be designed in the US, but all trace back to the same few manufacturers.
5. I have to apologize, I got the 1 million IOPS number from Charlie the moron. Shouldn't have trusted him. It does make more sense now that it is an aggregate from 7x PCI-e cards, so about 140K IOPs per card. Still doesn't change the economic dynamics vs the F5100 or F20. We will just have to wait and see how greedy Intel wants to be with this. All in all, I think you already know that I am bearish on Stec and Sun F20.
Anonymous on Sep 24 2009, 00:28

Regarding #3: because people source harddrives from newegg, dell decided to lockout these 3rd party harddrives in their latest PERC controller
Anonymous on Feb 24 2010, 12:30