

Wednesday, December 31, 2008

Happy new year!

As the first of my readers enter 2009 in a few minutes (e.g. the readers in New Zealand), i want send some greetings from my side:

Posted by Joerg Moellenkamp in General at 11:46

FUD from the Linux Foundation or: Mr Zemlin again ...

I wrote it before in my blog, but once (when i was young and wild) i threw a sales rep out of my office, who tried to get into the account (me) by FUDing their competitor. I was responsible for an project in the range of 10 Million Euro, thus i assume they got to much "Deutsche Mark" signs in their eyes. I had a basic simple rule for conversations with vendors: You have to shine for your self. If you need to point to non-product related weak points, there is a high probability, that the business relation will be an unpleasant ride.

I wouldn't dissect HP or IBM balance sheets at a customer meeting, albeit i would dissect performance claims and hint to quirks. I wont use the unclear future of FuSi directly at a customer. I fight on a technical level ... everything else is for sissys not knowing their shit ...

Using FUD is a good fear detector in my daily business. The amount of spreaded FUD is proportional to the amount of fear. Thus i have to assume that Mr. Zemlin of the Linux Foundation is really afraid of Solaris.

In the article "Linux in 2009: Recession vs. GNU" wants to make a point for Linux again. It's Zemlin FUD time again. I tend to award him the "Golden FUD catapult 2008" for this stuff and his infamous involvement in this story

At first: When all the commercial products are to expensive for a customer, he could opt in recession times for the free stack at Sun. For example running your systems with Opensolaris, Sun Java Application Server, Mysql or the Sun Web Stack for free without support. And when your company have weathered the recession and your budget isn't such a sad story, just call the Sales Rep. Or you could use the equivalent to a debian like structure. OpenSolaris 2008.11 with Glassfish V3 with Mysql Community Edition. Recession IT budgets doesn't equals Linux. But i wrote about that before. I wont write about it again here ...

I was upset by another comment. He states : Zemlin also sees FOSS as remaining strong on the server, with Linux continuing to be the major player. Most of Linux's growth in 2009, he said, will be "at the expense of Sun Microsystems, which is floundering in its business model right now. People look at Linux, and they say, HP, IBM, Dell, Intel and AMD -- these are collectively not going to go out of business any time soon. Then they look at Sun Microsystems, and they say, 'Whoa! This company has some serious financial difficulties, they have an uncertain future -- that's not a safe bet for me.' Nobody is really growing much, but where there is growth, it's going to in Linux." I want to dissect this statement.

At first. AMD and Intel are somewhat OS agnostic. They make their money with supporting Windows, not with this niche market called Unix x86 market. Sorry ... the Q1CY2008 x86 server market was 1.9 million systems large. The market for MacOS X based systems was 2.6 systems(perhaps this is the reason, why they got the "65nm chip in a 45nm casing" custom build from Intel for the MacBook Air). Server x86 is a an intersting business for both it it doesn't pay their bills. Or the other way round: If you don't like Sun x86, you just can user IBM, HP or Dell as well with Opensolaris x86. And furthermore: IBM, HP or Dell aren't married to Linux. When the three think that there is money to make with Solaris, they will sell it. And they already sell it. And they will drop Linux faster than you can write "penguin" when they come to the conclusion, that it's not cost effective to sell and support Linux.

BTW, Mr. Zemlin, Sun has no financial difficulties ... it has a problem with its stock price ... the financial reality and the stock price isn't really correlated (if it's correlated at all, i have my doubts about it)

Okay: Opensolaris ... the code is opensourced, the cat is out of the bag. Anybody can build it's own business on the basics. Just a thought game: Let's assume Sun would go out of business tomorrow. Without a warning. Jonathan says: "Dear shareholders, here is the money. Let's call it a day. And thank you for the fish". Now take into consideration, that

the installed base of Solaris as large as billions and billions. Business is like political power. It don't like a vacuum. Allowing everybody to participate in the development is just one side of opening the source. The other side is the fact that you enable everybody to support the code. The whole business model of Red Hat or Suse is largely based on this point. Thus in the theoretical case Sun support for Solaris would disappear, it would take a few days until another company fills this vacuum. There is money to earn. Much money. But this whole discussion is hypothetical: There is no vacuum to fill, as Sun won't disappear. And Sun wants to earn this money. Anyway: The complete discussion and the idea introduced by Mr. Zemlin is just utter nonsense, classic spreading fear, uncertainty and doubt. FUD 101.

I could make a similar prediction about Linux: In the course of the recession more users will opt for free distributions like Debian, CentOS and OpenSuSE as their budgets don't allow them to buy support for every system. The sales of subscriptions will decrease. This will lead to financial pressure to the major vendors of Linux distributions. Because of the single-trick ponyness of both vendors (RedHat and Novell) it put both companies on the verge to oblivion. They will cease to exist if purchased by IBM or HP. Furthermore this will bring problems for the core development of Linux (not the drivers, the real core) as the both companies employ many of the core developers.

Doesn't sound reasonable? Well ... then tell why it's more rational to assume a company with hundreds of billions installed base, billions of revenue a year, positive cash flow, 2 billions at the bank, almost no debts should go out of business or purchased. The last thing is extra ridiculous. It's credit crunch time at the moment. The hedge fund locusts don't get money to do their business at all even for smaller deals, everybody keeps it's money to have liquid money in case of a longer credit crunch phase and we didn't talk about the anti-trust regulations at all.

At the end Mr. Zemlin just spreads FUD. Large amounts of it. Dear Mr. Zemlin, we can talk about the importance of features like in your last FUD marketing attempt, we could talk about technical advantages of Linux and Solaris. But dear Mr. Zemlin, stop to discredit yourself as a FUD thrower ... otherwise you are burden for the community you represent ... again ...

Posted by Joerg Moellenkamp in Solaris at 00:28

Tuesday, December 30, 2008

Goodbye 2008, hello 2009!

2008 is over. Almost. I don't think, that the last few days will change it. At least i hope so. But ... i don't see any game-changing development for the last few hours of this years.

One of the big headlines of this year was "loss": The biggest loss at first. I've lost my grandma at the beginning of this year. My grandma was already in an age of 60 when i was born. That's an age where the death starts to take away our grand parents. So the additional 34 years were a gift. Each year.

In the middle of the year i've lost a friendship ... at least i thought she was a good friend. But errors are human and i'm just a human. I really thought that that this loss would hurt me. I don't make friendship easily. But at the end the exact circumstances of this situation created opposite feelings. At the end it's better that way: Friendship is something mutual ... and if one person think of a relation as a friendship and one other as a loose acquaintance it's better to keep it the second way. I still do not fully understand the dynamics of that situation but i stopped thinking about it a while ago. You can't look into the brain of other people.

At the end of the year I've lost another friendship to a large heap of misunderstandings. But without trust there is no friendship ... and at the end there was no trust. It's just a very loose acquaintance now from my point of view. It was my fault as well as the the fault of this friend. But well ... this happens. Shit happens.

The other big headline was "win": I've won new friends. Albeit i never met Marianne, we have developed a strange kind of friendship in the mean time. I've met other great persons this year. I would call them friends right at this moment, but they are already important to me.

Additionally i thought, i've lost the very special relation to another person, but the end of this year proved otherwise. There are very rare persons in your life who are so special, that even not seeing them for a while and huge hassles in both lifes doesn't change a thing.

Personal targets? The last two years weren't good ones in regard of my weight. Too much work, too much frustrations, less sports. This has to change, but i started to do more sports. So this will be just a matter of months after. Christmas is over, so no further problematic days in sight. Want back to my old weight.

The blog got more and more important to me this year. Personally and professionally. My objectives for this year were 5000 Visits per day , 750 subscribers und 600.000 Pageviews. The results of this year: Up to 10000 visits per day (with large variances the, at the weekend the visit count drops to 1000, monday and friday are the best days), 1175 subscribers (right before cristmas) and 750.000 page views and 15 million served requests by the webserver. Objectives for the next year? Well ... up to 20000 vists, 2000 subscribers and 1.5 million page views. I know ... very high targets but targets have to be hard to be challenging.

The "Less known Solaris features" series was a large success. 15.000 downloads of the pdf versions. The web log versions of the document's articles were loaded 105.859 times in the last year. My plans for the next year? I'm already working on a german version. I will publish it as a book as planed at first. I will publish it as a free .pdf-book again. In the next few weeks i will publish a revised version (with the help of Ceri Davis, who helped with the english grammar and typo checking) of the english version. I hope i can iron out all the tex warnings in the next few weeks. Such a tedious job...

Professionally i will plan my next steps in the next few weeks: I'm member of the '09 SEED mentoring programme as mentee and i was one of the first ones matched to a mentor. I think i had luck with my mentor.

I plan to talk more often on public conferences (hope my papers are accepted, one is a live LKSF session, another a talk about my view to virtualisation and some additional smaller sessions) and i'm planing for a new career step (I'm staying at Sun, at least when Sun doesn't think, i should be part of the 6000 people, so it's something internal to Sun) but it's much to early to talk about that.

I wish you all a good start in the next year. May all your hopes for 2009 come true. May all your fears in 2009 turn out as unnessessary over the course of the year.

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

Posted by Joerg Moellenkamp in [Braindump](#) at 09:41

Monday, December 29, 2008

Less known Solaris Features: About crashes and cores - Appendix B: Live crashdumps

With gcore you can create core dumps of processes without the need to stop them. You can do the same with the crash dumps as well. You can trigger a crash dump without rebooting the system. The command for this task is the savecore command. The normal task of this program is to take the content of the crash dump device and to create files in the regular filesystem to make them persistent for further studies after the reboot of a system.

How to trigger a live crash dump? A less known option of this command is the -L option. With this option you can command the system to create a live crash dump. The system isn't paused while creating the dump, so the live crash dump isn't self-consistent (the state of the system at the end of the creation may be different than the one at the beginning of this process. Nevertheless a live crash dump is still a useful tool. Triggering the crash dump is really easy as the following example shows.

```
# savecore -L
dumping to /dev/zvol/dsk/rpool/dump, offset 65536, content: kernel
100% done: 155962 pages dumped, compression ratio 3.32, dump succeeded
System dump time: Mon Dec 29 22:13:41 2008
Constructing namelist /var/crash/central/unix.0
Constructing corefile /var/crash/central/vmcore.0
100% done: 155962 of 155962 pages saved
After the creation of the crash dump you can analyse it with the standard
toolset like the mdb:
# mdb /var/crash/central/unix.0 /var/crash/central/vmcore.0
mdb: failed to read panicbuf and panic_reg -- current register set will be unavailable
Loading modules: [ unix genunix specfs dtrace cpu.generic uppc pcplusmp scsi_vhci zfs
ip hook neti sctp arp usba uhci fctl sd md lofs audiosup fcip fcp
cpc random crypto smbsrv nfs logindmux ptm ufs sPPP nsmb ]
> ::ps
S  PID  PPID  PGID  SID  UID  FLAGS      ADDR  NAME
R   0   0    0    0    0 0x00000001 ffffffbc29c30 sched
[...]
R   7   1    7    7    0 0x42000000 fffff00c5c42188 svc.startd
R  418   7  418  418    0 0x4a004000 fffff00ca5a0018 ttymon
R  391   7  391  391    0 0x4a014000 fffff00cd995e08 sac
R  395  391  391  391    0 0x4a014000 fffff00cd9938e8 ttymon
>This enables you to analyse the state of your system without having to aim for a moving target.
```

Posted by Joerg Moellenkamp in Solaris at 22:11

Bullshit in regard of Openoffice

Hmm ... long time ago i've subscribed a Google news alert about "Sun Microsystems". Sometimes it gathers really excellent articles at locations i wasn't aware of before. But most of the time it gathers pure bullshit. On of this "pure bullshit" articles is this one: "OpenOffice bound for orphanage?" written by Mr. Barlow.

At first: Openoffice wasn't developed by volunteers the last 20 years. The earliest versions were written at by Star Division in Hamburg. In 1999 Sun acquired Star Division. We opened up the code in October 2000 and since this time there is OpenOffice and StarOffice (an enriched version of OpenOffice). Many Staroffice developers paid and pay their bills by working for Star Division or Sun.

About the orphanage thing: I have a good insight into that topic ... i'm working in the Sun Office in Hamburg. Staroffice Development and the Sun Office are in the same building. At a part even in the same floor of a multi-floor building. And i don't see any sign of orphanage ...

I won't comment on the other bullshit in the text. So, dear Mr. Barlow, please just don't comment on things without doing any basic research (a short search on wikipedia would be sufficient). But i don't think that was his intend ... i smell a pro-Microsoft agenda in this text.

BTW: The article "Measuring the true success of OpenOffice.org" referenced by Mr. Barlow is written by a Novell guy

working on a code-fork of Openoffice and he is promoting it in his text. Hmm ... i'm not that deep into the openoffice development structures, so i will leave it to a fellow blogger from the Staroffice Development to comment that article. But some parts reads as a Novell vs. Sun thing at some parts. Obviously it's a pain in the a... for Novell that the sole usable Office suite for their flagship product SuSE Linux is a product of a development community lead by Sun. But that is just my impression.

Posted by Joerg Moellenkamp in Sun at 19:39

Andy at the 9th ICFA seminar

Andy held a presentation at the 9th seminar of the International Comitee of Future Accelerator: The Solid State Storage Revolution. The pages 25 and 26 are especially interesting.

PS: Other topics were Neutrino factory and muon collider R&D or X-band high-gradient and CLIC R&D for future colliders and other strange stuff like that ...

Posted by Joerg Moellenkamp in General at 18:26

Requalified for FTL status

Only 45 flights this year ... thus it was a little bit harder to requalify for FTL. At the end it was the last flight of the year in conjunction with a retroactively submitted boarding pass. Reached it with the margin of a whisker. Now i'm FTL until 2012 ...

But i have no idea how to get 93.261 miles in just two days to reach Star Alliance Gold status Well, first class to New York and back should do it ... but that's a little bit expensive for an airline status with dubious real benefit. But i assume, it's impolite to write that into the message beneath the actual value.

Posted by Joerg Moellenkamp in Business Travel at 15:32

Sunday, December 28, 2008

Logical Domains 1.1

The developers of LDOMs gave the users of the T1 and T2(+) based system a nice christmas present: The 1.1 version of the LDOMs. This update provides some interesting new features:

a new XML based interface to manage Logical Domains

"Guest Domain Migration provides the capability to migrate a guest logical domain from one server to another compatible server. If active, the domain on the source server is suspended, and its configuration and run-time state are transferred to another server, where the domain is recreated and resumed."

Virtual Input/Output (I/O) Dynamic Reconfiguration (DR) provides the ability to add and remove virtual I/O services and devices without rebooting

Enhanced Networking subsystem"Network Interface Unit (NIU) Hybrid I/O provides support for a virtualized I/O path, a hybrid I/O path, and better performance and scalability. "

"VLAN Support provides the capability to configure and use virtual local area networks (VLANs) in logical domains"

Enhanced Storage subsystemVirtual Disk Failover adds support for disk multipathing

Single-Slice Disk Enhancement allows installing the Solaris OS on a single-slice disk.New and Changed CLI Idm Subcommands

You can download the new version of LDOMs at at the Sun website.

Posted by Joerg Moellenkamp in Sun at 13:18

Thursday, December 25. 2008

Hohlräume

Interessanter Kommentar von Norbert Bluem zur aktuellen Wirtschaftskrise im Hamburger Abendblatt: Die gegenwärtige Wirtschaftskrise bietet manche Erfahrung, aus der man klüger werden kann: Keiner der großen ökonomischen Koryphäen des Neoliberalismus hat die Krise kommen sehen. Die etablierte Wirtschaftswissenschaft hat sich blamiert bis auf die Knochen. Der "Chefökonom" der "Bild", Professor Sinn, das Beste, was die Zunft offenbar bieten kann, im Nebenberuf Generallobbyist der Privatisierung des Sozialstaates, gab noch Mitte des Jahres Prognosen von sich, die so treffsicher waren wie eine Wettermeldung in New Orleans gewesen wäre, die beim Ansturm von Katrina eine Schönwetterperiode angekündigt hätte. Ökonomische Blasen sind virtuelle Hohlräume, die mit nichts anderem als mit Erwartungen gefüllt sind. Mit den Realitäten von Arbeit und Wertschöpfung haben sie so viel zu tun wie Wetten mit Leistung.

Posted by Joerg Moellenkamp in General at 19:53

Updated SDN article about pfexec

After some comments from readers Marina and I updated the article about pfexec at the Sun Developer Network. We integrated some clarifications and changed the wording in some parts.

It's difficult to write a text, when the jargon is a little bit ambiguous. After a writing "root capabilities" several times, you want to use another word ... as i think of having root rights as a privilege (you are a trusted and knowledgeable user), i like the phrase "root privileges". But as privileges are already a common phrase in Solaris 10, it's a little bit problematic as the capabilities of root are the result of a set of privileges, not a single privilege. So there is no "root privilege" (set this privilege and you are root) in the sense of the least privileges concept. At the end we opted for leaving the phrase "root privilege" in the text at some parts.

Posted by Joerg Moellenkamp in Solaris at 11:01

Tuesday, December 23, 2008

Solaris Zones from a security standpoint

Glenn Brunette and Jeff Victor published an interesting document about Solaris Zones with the Blueprint "Understanding the Security Capabilities of Solaris Zones Software": So why does the world need yet another article about Solaris Zones? Simple. Most publications and sites focus on the consolidation benefits of Solaris Zones. While server and service consolidation is a key use case for Solaris Zones, there is so much more to the technology. Other materials focus on system administration practices related to configuration, installation, management, and troubleshooting. This is incredibly useful information, but there is still an important gap. Namely, many people do not have a full appreciation of the security benefits enabled by Solaris Zones, and sparse root zone configurations more specifically. Definitely a must-read document!

Posted by Joerg Moellenkamp in Solaris at 17:25

Recommended listening: Collide - Two headed monster

This recommendation is from the darker side of my musical taste. It's Collide ... an electronic band from California. I've wrote about this band in this block back in 2005. How can I describe Collide at best. Hmm ... well ... think about the more melodic pieces of Nine Inch Nails (like Ruiner, Heresy and Perfect Drug) with female vocals.

The both musicians of Collide published their new album end of September. I really waited for their new album, as their remix album Vortex was somehow addictive and albeit their spin off project "The Secret Meeting" produced excellent results it made the waiting time for the new Collide album a little bit longer for me, as Collide is still a tiny little bit better. As usual with my taste of music: You like it or you hate it. I recommend Tongue tied & twisted, Spaces in Between and Shifting for a first taste.

Posted by Joerg Moellenkamp in Music at 16:30

Best wishes

I'm driving to my parents in a few hours from now. I'm already packing my bags for a few days at the moment. But i won't leave Hamburg without sending you my best wishes for Christmas! I hope you will have some quiet and stress-free days with your loved ones!

Posted by Joerg Moellenkamp in General at 09:08

Another comment about OpenSolaris 2008.11

A really positive article about Opensolaris 2008.11. Steven Lawson writes in OpenSolaris 2008.11: Its Time Is Coming: I was attracted to OpenSolaris 2008.11 in the first place by a couple of other internet articles, one of which posed a question, the other a bold statement.

The question was: "Will OpenSolaris 2008.11 attract Linux users?"

And the statement was: "OpenSolaris tackles Ubuntu dominance".

Let me deal with the question first by answering simply, 'Yes, if they have any sense'. And as for the statement, my reply is: 'No, it doesn't, not yet - but one day soon it should'. and OpenSolaris 2008.11 is well worth investing some time and disk space on. I think we will be hearing a lot more about this operating system in the future. Of course he sees room for improvement, but at the end the improvements are just "We work on that" issues: More packages (in this article multimedia applications) and more wireless drivers.

Posted by Joerg Moellenkamp in Solaris at 08:14

Monday, December 22, 2008

Road ahead ...

Posted by Joerg Moellenkamp in Photographie at 20:25

Recession friendly software

Julian Exenberger has an interesting point of view in regard of Sun software. He wrote about it in "It's recession time and Sun has the best recession kool aid.": Well I don't know if anyone has noticed but one of the coolest recession friendly enterprise software stacks around is being offered by that company we all love to hate: Sun Microsystems. That's exactly the stuff i'm telling to my customers in many presentations (of course without the hate stuff). An everoccurring answer is: "But my ultraimportant application needs BEA and Oracle. It runs that way for a while now and i don't want to change that". Okay ... i'm fine with that. Keep it this way. But: There are tons of small services in a corporation that doesn't need BEA or Oracle. If you take this kind of applications from BEA and Oracle to Glassfish and Mysql or Postgres, you get something really precious: Licenses for cores for your Oracle/BEA/WebSphere/younameit-dependant application because your licenses are perpetual and you can reuse them as you like to do (at least when you didn't made some strange agreements). Just look at the per-core price of an Oracle license to see the possible savings. So you don't have to spend for more licenses when you opt for a larger server, just use the spare ones. And as there is commercial-grade support available for Glassfish, Mysql and Postgres you don't base your operations. on newsgroup reading.

I don't want to be misunderstood: Glassfish V3 is at least (!) on par with any other commercial appserver. As Julian writes it: Sun has a really cool application server which is fully buzzword compliant in the form of Glassfish. Besides the fact that its pretty scalable with old Grizzly, it comes with a rock solid ESB and messaging system, A decent compatible web service stack, and on top of that it is dead simple to administer. And then of course lets not forget the awesomeness that is Glassfish 3.0. It's not about getting cheaper, but less capable software what i'm proposing here. It's just about breaking some habits. Now what is really cool about this whole stack is that everthing I've mentioned here is either open source or free, with support contracts available. If you wanted you could download all individual components and not cost to you, that's right no cost. How recession friendly is that.

Update: Okay, i won't write articles (the time in the article is the start time, not the submit time) that late again after being awake since half past four in the morning ... this article was full of typos and grammar errors (hopefully i have corrected most of them now). It's difficult enough for me to write a decent english when i'm fully awake ... it seems to be impossible when i'm already on my way to the land of dreams. Sorry for that!

Posted by Joerg Moellenkamp in Sun at 20:08

Road to Death Valley

Posted by Joerg Moellenkamp in General at 16:13

Crash twittering

As you may have read in the news, there was a plane crash in Denver. A Continental 737 left the runway in the wrong direction at take off. But this one is a little bit different: Mike Wilson - one of the passengers - twittered about his experience.

I don't think i would have the nerves to even remember that there is something called "Internet" after surviving plane crash. I'm glad that there wasn't any fatalities at this crash. As far as a read the incoming news they directly crashed in front of a fire station.

Posted by Joerg Moellenkamp in Aviation at 08:46

Sunday, December 21. 2008

Does happiness have a price tag?

Posted by Joerg Moellenkamp in Fundsache at 16:55

Gedanken zum Fest 2008

Ich habe die letzten Jahre eigentlich immer ein paar Zeilen zu Weihnachten geschrieben. Aber irgendwie habe ich dieses Jahr noch weniger Weihnachtslaune als sonst. Null ... nix ...

Ich mag dieses Fest irgendwie nicht. Obwohl: Weihnachten an sich ist ganz schoen. Die ganze Meute fällt bei meinen Eltern ein, frisst ihnen die Haare vom Kopf und lenzt die Getränkervorräte bis zur Bilge und trollt sich dann nach ein paar Tagen wieder in die Republik nachdem die Vorräte wieder von uns aufgefüllt wurden. Nebenbei wird die Internetverbindung so gluehen, das mein Vater schon überlegt, für Weihnachten stärkere Leitungen anzumieten. Grossfamilien sind doch was Schoenes ... manchmal

Nur manchmal? Ja, machmal. Wenn man klein ist denkt man ein wenig anders. Bei endlicher Investitionssumme des Weihnachtsmanns heisst Grossfamilie halt kleinere Geschenke. Liegt in der Natur der Sache. Und obwohl ich nie zu kurz gekommen bin, so nimmt die kindliche Logik manchmal doch eigenartige Wege. Ist man dann irgendwann endlich mal auf den Trichter gekommen, das Weihnachtsgeschenke eigentlich nur Weihnachten stressig macht, merkt man das der Punkt hinter Weihnachten eigentlich ein ganz anderer ist. Familienzusammenführung. Das fällt einem aber erst dann auf, wenn sich die Familie durchs Erwachsen werden in alle Windrichtungen zerstreut.

(BTW: Mama, Papa ... es ist kein Schockgeschenk wie im letzten Jahr geplant ... ihr koennt euch beruhigen Ja ... meine Eltern lesen hier auch ... sagen immer, man erführe schneller im Blog, was mit mir los ist, als das man darauf wartet, das ich anrufe)

Aber warum stellt sich bei mir keine Weihnachtsvorfreude ein. Weihnachten ist ja ein Fest, das sich mittlerweile von der generellen Erstverfügbarkeit von Weihnachtsfressalien bis zum Fest so mindestens zwei Monate lang hinzieht. Geschenke werden bereits einen Monat vorher daheim aufm Sofa bei Amabay und Quellotto gekauft. Wenn irgendwann der Einzelhandel uebrigens weniger Geschäft macht, dann könnte es damit zusammenhängen, das es eben nicht reicht, etwas Weihnachtsdeko aufzuhängen, die Servicekraefte mit Weihnachtsmützen rumlaufen zu lassen, die Preise am Hauptspielzeug etwas runterzudrehen und die Marge für Batterien ins Unsittliche zu erhoehen, um an Weihnachten zu verdienen.

Die Leute sind bequemer geworden und Geschenke werden nicht nach Preis bemessen sondern nach Menge und Groesse der Kartons ... zumindestens so in der kindlichen Denke. Einfach mal so ein paar Jahre zurueckdenken: Worueber hat man sich mehr gefreut ... ueber das einzelne teure Geschenk oder wenn man vor nem Haufen Kartons stand. Okay ... bei mir was das etwas anders. Mein erster Computer war sooo teuer, das es Weihnachtsgeschenk und Geburtstagsgeschenk in einem war, es war nur ein Karton, aber gefreut habe ich mich trotzdem bis Meppen. Aber ich war schon immer eigenartig. Wenn sich das weihnachtliche Investitionsvolumen durch Interneteinkauf in noch mehr Kartons umsetzen laesst, dann wird das getan. Das hat nix mit "Geiz ist Geil" zu tun, das ist einfach nur die Anpassung der Eltern an die kindliche Mathematik des Freuens. Ein Geschenk wird für ein Kind nicht schoener, nur weil man es beim lokalen Elektroverhoekerer für die UVP gekauft hat, anstatt irgendwo zum Angebotspreis. Interessiert nicht. Hat es mich zumindenstens nicht ...

Das vorweihnachtliche Freuen ... das stellt sich bei mir immer weniger ein. Vielleicht waere es besser, wenn Weihnachten etwas konzentrierter kommt. Zwei Wochen vorher die ersten Spekulatius, eine Woche vorher den ersten Glühwein. Vielleicht wuerde dann auch wieder Vorfreude entstehen. Weihnachten muesste irgendwie wieder intensiver werden, und nicht in diesen 2-3 monatigen low-intensity conflict ausarten.

Genug das Fest der Familie, der Geschenke, des Batterienwuchers, des ungehemmten Stromverbrauchs für Lichterketten, als würde des keine globale Erwärmung geben, schlecht geredet. Ich wünsche Euch eine schöne restliche Vorweihnachtszeit.

PS: Und immer dran denken ... durchhalten ... in spätestens einem Monat können wir die ersten Schokoosterhasen kaufen ...

Posted by Joerg Moellenkamp in Braindump at 06:00

Saturday, December 20. 2008

links for 2008-12-20

Bundestagsdrucksache 16/10220: Stand der Planungen für den weiteren Ausbau der A 20

(tags: A20)

Posted by del.icio.us in del.icio.us at 12:00

Take off from Las Vegas

Posted by Joerg Moellenkamp in Photographie at 08:22

Caffeine

That was an successful end of this year yesterday. Customer more than satisfied with performance. Additionally he learned a thing or two. And i learned a thing, too. I will never touch the coffee again (at least not three cups). I felt like Roger Rabbit on Speed err ... on caffeine yesterday evening ...

Posted by Joerg Moellenkamp at 07:32

Thursday, December 18. 2008

Neue Events

In der linken Seitenleiste stehen wieder eine Reihe von Terminen von Events, die entweder von Sun ausgerichtet werden oder bei denen Sun beteiligt ist. Ich bitte um freundliche Beachtung. Sollte bei einem interessanten Event kein Link zu Anmeldung sein, wendet euch bitte an den Sun oder Partner VB Eures Vertrauens.

Posted by Joerg Moellenkamp in Sun at 20:52

24 hours of air traffic

Posted by Joerg Moellenkamp in Aviation at 20:26

links for 2008-12-18

Using Solaris RBAC to only allow scp/sftp

(tags: Security solaris sysadmin ssh rbac)

Fernverkehrsseiten von Marcus Grahnert

(tags: travel businesstravel Bahn Wagenreihung)

Posted by del.icio.us in del.icio.us at 12:00

About networking performance in Solaris

I found two really interesting articles about the networking in Solaris: The first one - Examining Large Segment Offload (LSO) in the Solaris Networking Stack - examines the effect of large segment offloading to the cpu utilisation. The second one - CPU Utilization in networking - the TCP transmit path discusses how the cpu utilisation changes with different block sizes and number of connections.

Posted by Joerg Moellenkamp in Solaris at 11:30

Delamination dementi

There were some rumours last week about the delamination of the wings, albeit the Flightblogger was able to gather some other informations that this rumours was nonsense. But after almost 15 years in the computer business i've got a fine sensor for half-a... dementis. Ben Sandilands writes in Boeing avoids the 'D' word about a statement he was able to gather from Jon Ostrower (the flightblogger) which he got from Boeing: Our tests continue to confirm that the composite construction of the 787, including the wing, is sound. The 787, like every commercial jetliner, is going through a rigorous series of tests to verify the safety and certificability of its structure, systems and operations. The FAA will review the results of this compliance activity prior to certifying the airplane for commercial use. I don't think, that's a good dementi. I would like to hear "There is no delamination. Period." instead of "When we get this plane through certification, you can be sure that there is no problems with the wings"

Posted by Joerg Moellenkamp in Aviation at 09:38

New scrub

One of the annoying things at zfs was the restart of the scrub when you take an snapshot while the scrub runs. This issue was fixed some builds ago, but Matthew Ahrens published an interesting article about the fixed code yesterday:

New Scrub CodeWe have solved a nasty bug by rewriting the scrub code to do its work in syncing context, visiting datasets in order from oldest to newest. I had to introduce a few new on-disk data structures so we can quickly determine this order. And there are a few operations that the scrub code needs to know about and account for. This work lays a bunch of infrastructure that will be used by the upcoming device removal feature. Stay tuned!

Posted by Joerg Moellenkamp in General at 05:43

Wednesday, December 17, 2008

Linux, Solaris and the alleged conflict

Jim Grisanzio wrote an excellent piece about the alleged areas of conflict between Linux and Solaris in "It's still too late": So, is it too late to catch Linux? I'm not sure it matters much in that context. It's a big world, and there is room for all of us to fit. We simply have too much work to do learning from Linux in some areas where they are strong, focusing on some of our clear advantages in other areas, transforming the existing Solaris base into an open community, and reaching out to new users and developers who have never even heard of us. It's not too late. Not by a long shot. I just don't view projects from that perspective. I share this opinion. This is the very reason why I don't like the question "Sun is the bad guy because they have their own Unix and don't work at the Linux kernel". I'm a Solaris supporter ... but there is room for both operating environments.

Posted by Joerg Moellenkamp in Solaris at 18:43

Virtualbox 2.1 released

We have announced a new version of Virtualbox. 2.1 contains some new features and bugfixes. You will find the following add-ons in it: Support for hardware virtualization (VT-x and AMD-V) on Mac OS X hosts
Support for 64-bit guests on 32-bit host operating systems (experimental)
Added support for Intel Nehalem virtualization enhancements
Experimental 3D acceleration via OpenGL
Experimental LsiLogic and BusLogic SCSI controllers
Full VMDK/VHD support including snapshots
New NAT engine with significantly better performance, reliability and ICMP echo (ping) support
New Host Interface Networking implementations for Windows and Linux hosts with easier setup (replaces TUN/TAP on Linux and manual bridging on Windows) You will find the changelog at [here](#). The binaries are available at the usual location.

Posted by Joerg Moellenkamp in Sun at 16:55

UltraSPARC T2 Reference Design Kit

Sun announced an Reference Design Kit for CMT: It's an UltraSPARC T2 at 1.17 GHz on an ATX board. You will find more information at the [datasheet](#) and the [Sun website](#). As this board is intended for hardware developers planning their own UltraSPARC T2 system, the kit contains such things like board schematics, gerber files, board layout design and the bill of materials for the board.

I don't have an idea what this board will cost at the moment, but this looks as a way to custom-build a CMT workstation

Update: Too expensive for a workstation at home ...

Posted by Joerg Moellenkamp in Sun at 16:21

NAND-Flashes from Sun and Micron

News from the SLC flash front - Micron, Sun take NAND endurance to one million cycles: Micron Technology Inc. (Boise, Idaho) has said it has worked with Sun Microsystems Inc. (Mountain View, Calif.) to develop a single-level cell (SLC) NAND flash memory technology that extends the lifespan of flash-based storage for enterprise applications and reaches a million write cycles.

Posted by Joerg Moellenkamp in Sun at 15:39

links for 2008-12-17

no solaris ... no life ! - Life with Solaris

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

(tags: solaris opensource packages)

Posted by del.icio.us in del.icio.us at 12:00

Tuesday, December 16, 2008

NFS at 2 GByte/s

Brendan Gregg published a 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 a benchmark-special result at first, to show its 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 Solaris at 13:21

Cisco selling blades?

Cisco seems to look for new markets - Cisco planning significant data center assault: Perhaps the most important example of that will be a new Cisco blade server system expected next year. This will take the company into the data center compute space, right up against longtime stalwarts — and up to now, Cisco partners — IBM and HP. This matches to rumours and hints I've got in in the last few weeks. But I don't believe that Cisco will play a large role in this business. There are already several players in this game ... IBM, HP, Dell, FSC and Sun. And more important: People think of Cisco as the supplier of their networking equipment ... not of their servers ... and Cisco hasn't a long track record as a supplier of servers. Cisco needs a real differentiator ... I'm curious what this differentiator might be ... but at the moment I'm not convinced about the success.

And there is another problem: Network admin departments, storage admin departments and server admin departments have separated supplier spheres in many companies. And Cisco is a supplier for the network admin departments and they rarely connect to the server departments. And: Did you know Sun sold a hardware load balancer for a while? No? Cisco will have a similar problem ...

Posted by Joerg Moellenkamp in The IT Business at 10:46

Monday, December 15, 2008

Attracting Linux users?

Interesting conclusion of Charlie Schluting in his article "Will OpenSolaris 2008.11 Attract Linux Users?": If OpenSolaris spent the next 1-2 years focusing primarily on adopting standardized open source software rather than implementing shiny new features, the project would stand a much better chance of attracting long-time Linux users. Don't think so. The new features are the stuff that attracts Linux users ... you have to do both. Making it easy for Linux users to choose their favourite OS on technical merits than on usage habits. But there has to be a technological lead. Otherwise there wouldn't be any reason to think about changing. Often good-enough is good-enough. You need such features like dtrace or zfs or crossbow to create a demand. OpenSolaris doesn't need the exclusive community that Solaris enjoyed all these years, it needs the huge Linux community. I think it will happen; the question is, "how quickly?" ACK. I think Solaris will get more influence from the Linux side. We need the bigger linux community. The problem is: You can't get a linux-like standard on one side without violating some standards and alienating Solaris users. Solaris don't need the community at the price of its soul. We have to balance this.

PS: Anonymous is correct ... i've read this article in an incorrect way as i was distracted, so it led me to a wrong comment. I've reread the article again and edited the article.

Posted by Joerg Moellenkamp in Solaris at 14:49

Solaris at Microsoft?

There is an interesting comment at Slashdort. Sadly it's an anonymous comment ... BitTorrent is an excellent intranet content-distribution tool; we used it for years to push software and content releases to 600+ Solaris servers inside Microsoft (WebTV).

Posted by Joerg Moellenkamp in General at 13:05

links for 2008-12-15

YouTube - Landing San Francisco from Cockpit 747 - (part-1)

(tags: aviation)

YouTube - Landing San Francisco from Cockpit 747 - (part-2)

(tags: aviation)

Posted by del.icio.us in del.icio.us at 12:00

About Crossbow

After the integration of the first bits of Crossbow, Sunay found some time write about some of the advantages of this network stack: Crossbow - Network Virtualization Architecture Comes to Life.

Posted by Joerg Moellenkamp in Solaris at 11:55

Sunday, December 14, 2008

A380 vs. Boeing 787 delay

After the last delay in the 787 program many news outlets write, that the delays in the 787 program reached the delays in the A380 program. There is only one important fact, they forget to report. The A380 was late just three months at first flight. The problems leading to the two years delay hit them afterwards. The 787 has already accumulated 2 years of delay before this aircrafts first flight. To be exact: At the moment nobody knows, if this airframe will fly ... there are just computer models and experience suggesting that this heap of plastic will take of at the end of the runway. They made a break-test of the wing box ... not the complete wing. So ... it's 3 month delay in the A380 program compared to 2 years in the 787 program.

But Boeing has a good chance not to gather that much delay after first flight for the first delivery as the 787 is a much smaller plane. On the other side there are the first rumours and discussions, that Boeing isn't able to build enough planes fast enough for their ultra-aggressive schedule for flight-testing. But that's a different story.

Posted by Joerg Moellenkamp in Aviation at 11:05

Saturday, December 13. 2008

Yesterday

There are days that feel like used ones. Somebody already took the best parts out of it and left only the bad ones. Yesterday was such a day. To start my meetings at a normal time in Dresden (normal time for the customers, the fact that i'm coming from Hamburg isn't the problem of the customer) i have take the 06:30 am flight. Check-in is closing at 06:00, urban train to airport starts at 05:31, so i have to leave at 05:20. I need a little bit less than an hour from deep sleep mode to fully operational so my alarm clock wakes me at 04:30. This isn't much of a problem for me, as i'm waking up between 05:00 and 06:00 in any case ... so it's just one hour less sleep. Nevertheless i like this hour ... and i like to stay in bed drinking tea, reading news and often my first mails for the job are written in my bed.

So i was really upset to find my aircraft delayed almost two hours. Instead of 06:30 this aircraft had an estimated 08:30 departure. At the end it was 08:15 ... but didn't made it better. At the end this led to an 30 minutes delay at my customer, so it wasn't such a problem for him ... it was more a problem for me.

Without FTL status i would be more angry. A 10 Euro voucher for breakfast as an excuse isn't really much money. Every business traveller knows the ridiculous prices at airports. It buys you a latte macchiato and perhaps a second one ... with FTL status you can at least enter the lounge and have some breakfast (free of charge) there and it's much more comfortable than at the gate. Especially when it's gate A28 in Hamburg. A28 in Hamburg are just stairs to the apron in the middle of a large hallway. The most uncomfortable gate at HAM.

At 08:10 or so we boarded the apron bus ... and not much later we were above the clouds. I was still removing slides from my presentation, as due to the delay i had already postponed and shortened the meeting .

The plane on the picture wasn't my plane ... that's the Continental shopping shuttle directly from Hamburg to Newark.

Okay, the meeting was okay ... i won't talk about it here because this isn't the place to talk about customer meeting. That's independent from it's outcome. So i don't say it was a bad meeting, and i don't say it was a good meeting. It was okay . But i thought long enough about this meeting just to keep my mind busy until i got to sleep. I've learned a thing or two yesterday.

Back to airport shortly before 02:00 pm. My aircraft was scheduled for a departure at 04:25 pm, so boarding starts at 15:55. Okay ... just waiting 2 hours in the DRS FTL lounge with bad coffee, some fruits and cookies. But there was something missing in the lounge

They removed the curtains there, so it feels like in a aquarium ... there was just missing a sign in front of the lounge "frequentius travellerus vulgaris" and people in front of it staring at you while you eat, work and wait for the time to pass ... after that you know how animals feel in zoo. BTW: As the Star Alliance Gold lounge is a different room, you could not look it it ... so you didn't need a sign "frequentius travellerus senatoris".

I know i shouldn't talk bad about the airport in Dresden. But this airport is outright depressing. Not much traffic. You know traffic is sparse when the next bird on the departure table is your own in 2 hours. And the weather wasn't capable to make it a brigher day.

But the day wouldn't be that bad if there weren't other delays and other hassles. At 15:55 the airline told us that the aircraft got to late to Dresden. 30 minutes additional delay to prepare the aircraft.... back to the lounge. It was nearly half past four when boarding started. I was glad that i changed my seat in Hamburg at the Quick-Checkin-Terminal. It would be a comfortable journey home. At least i thought so. At the boarding i've got a new seat. Do understand my "anger" you have to know something about the Dornier 328 (the aircraft on this relation). It has just three seats per row. One at left side of the aisle and two on the right sight. I'm checked in on the left side at one of the last rows to get one of the single seats (originally they booked me at 5F, but you can override it at the checkin). I choosed 11A. But Lufthansa had

no mercy. How could i question the infinite wisdom of the german flag carrier. They took away my single seat and gave me the other seat. So i sat dense packed with another passage in my size. On 5F. FscK. And just a second after realizing that 11A was still empty another passager placed himself on this seat.

But, i'm to fast. We (the passenger) entered the bus on the apron after boarding. The Dornier 328 is a small plane, so it's apron parking in any case. It was cold, my mood just went downhill. And then the woman of Gate 7 came to the aircraft and told us: "Sorry, you have to wait additional 10 minutes in the bus. The aircraft wasn't unprepared ... it was broken, so they had the change the aircrafts. Dear Cirrus (they operate the flight for LH) ... the business class meals in economy are nice ... but i would prefer more reliable planes.

At the end all this delays summed up to an hour. I arrived in Hamburg after an hour long flight at a little bit after 6:00 pm.

The only comfortable part of this journey was the new urban train. I've stepped out the of the aircraft at 06:01pm and was at home at 06:27pm. That's good ... that's excellent and it's free of charge for me as i have ticket valid for year for the most of the lines in Hamburg. As i like to sit on the observer terrace of the airport just to look after the planes i assume that i will use it on a regular schedule.

With all that time waiting at various location i had some time to think. I've learned an important thing yesterday. I'm an avid observer of the world surrounding me. I'm observed the people around me in the airport, in the train. I observed the people in the new shopping mall in the airport.

I thought for some time to move to the "better" parts of Hamburg ... Eppendorf for example, i even looked for something affordable near the harbour of hamburg. I have postponed it because i was sure at some point of time that we would see credit rates really soon. Yesterday i had an important thought: It was the right decision out of the wrong reasons. A good friend of me was right ... i would fell well there ... i wouldn't like it. I don't like the people walking around in the inner city of hamburg (you know ... Neuer Wall or Jungfernstieg) ... especially those one with a large heap of arrogance in their eyes. I wouldn't like to live in the middle of them. Why should i live near of people who hide their empty eyes with expensive Gucci sunglasses. Hide themself in expensive clothing. I don't need that. I know ... only a few people are that way ... but i can smell arrogance like sharks smell blood over miles and miles ... and i don't like the smell of it. I hate it.

Pretending to be someone else is a sure way to get an Prozac junkie. I know in the following sentence is to much pathos ... but i found a simple truth while looking outside the aircraft. It's knowledge that drives me, it's the curiosity that drives me, it's the will to distribute my knowledge that drives me, it's at foremost the existence of my family and my friends that drives me. Everything else is just irrelevant, redundant, futile.

I don't need a flat in one of this hip parts of Hamburg, half as large and double as expensive as my current one just to be one of them. I don't need to spend 50.000 Euro for a car just to be one of them. The one for 18.000 is as fast as the expensive one for the trip to the airport. And what's the rationale for such a car just to place it at the airport parking lot. Whats the rationale for a 4000 Euro chronometer when even the 10 Euro quartz timer is more precise and the 100 Euro radio-controlled clock is a second late in one million years. Besides of having a sign of status? When a person measures me by such signs, the person is absolutly irrelevant for me. It's that easy. And by the way... my cheap Fossil looks better than a 4000 Euro Breitling Omegaglasshuette ... but that's my personal opinion.

I get more satisfaction from learning something new or seeing the "Wow, this guy is right ... i've learned something" in the eyes. That's cool ... that is satisfying ... just try it on your own. That's why i like to hold presentations. The people learn from me and in every presentation i've held so far (and i held hundreds of them) i've learned a thing or two. That's why i'm blogging. I've learned a vast amount of new things from it ... and i hope you've learned a thing or two, too.

But ... well ... i needed some years to learn this, too at 27 i found it cool to drive a 250 PS Volvo C70. You don't see the cool design of this car ... or its motor. You just see the numbers rotating at the gas station. At 30 i thought about buying one of this expensive clocks, but my tendency to loose things kept me away from this (wrong reason, correct decision). And until the middle of this year i thought about moving.

There is a deep truth in the way the buddism handles the run for success or wealth. Stopping to run after it doesn't make you more lucky ... it make you less unlucky. Of course it's great to have luxury ... but a few days later you are used to the luxury and you recognize that you have to work even harder to stay at your current level.

By the way: The new shopping plaza at HAM airport ... well ... FAIL! EPIC FAIL! At least for me. Just all this expensive

shops selling expensive perfume, expensive clocks and expensive clothing. The usual suspects. The complete plaza looks so pre-financial-crisis. But enough for philosophy, today

I will close this article now. To add insult to injury i found out, what i forgot the last few months ... i found an envelope with my business travel receipts ... for 5 months ... shit ... this will a hell of a walk to Canossa. But to much money to simply ignore it. Fsch.

PS: There is another thing i found out on this business trip. For all the people out there sitting to long in front of a computer ... at the end the internet is just an empty place.

Posted by Joerg Moellenkamp in General at 22:24

Recommended listening: Diary of Dreams - Nekrolog 43

After my recommendation for the Amy MacDonalds album i've got some question from friends what's had happen to my usual favourite music. Nothing, my dears, nothing. I just extended my taste in music with that scottish artist. So i wan't to recommend a new album. It's the perfect opposite to the last album. Dark, really dark perfect for my usual mood in the winter. When my soul catches the colour of the time when i wake up. And to the others, which are unaware of my old music taste. Well ... that's normal ... i'm not depressed or something like that. I just like such music. Other friends speculated because of this music taste when would start to transform the guest restroom to a terrarium for flying foxes or bats. But that is a different story.

Today i want recommend Nekrolog 43 from Diary of Dreams. Yes, this music again ... i've talked about an earlier album last year.

Very dark lyrics with lines like Hey, prison cell
you can not keep me here
Silent waters are deep you knowin The Plaque or My home is where my heart died
Don't listen to what they say
I may not be your best friend
and I know you feel the samein the track Remedy Child. Surely not music for everybody or depressive people. But i like it It's a really good album. At first you should try The Plague. If this music is your style you should try the titles already mentioned. When you don't like this track, you shouldn't go farther

Posted by Joerg Moellenkamp in Music at 17:49

Sun X4600M2 vs HP ProLiant DL785 G5 revisited

I further digged down into the topic of the seemingly vastly slower HP Proliant 785 G5 when compared with a Sun X4600M2. In some discussions colleagues told me: You have to compare two systems with the same os and same database. This discussion even found it's way to the german heise.de forum. Okay, that's correct for a final proof but i don't expect a HP running the SAP SD benchmark on Solaris and MaxDB, so we will never know it

So i've studied the SD 2-tier benchmark list out of curiosity to find some hints to explain the performance difference between the both benchmark results.

Hypothesis 1: The different database and operating system explains the benchmarking differenceThis would be the most obvious explanation. The Sun benchmark used Solaris 10 and MaxDB and the both benchmarks done by HP uses Linux/Oracle respectively MS Windows/MSSQL. But then i found the following benchmarks in the benchmarking list. There was a benchmarking run (SAP Certificate: 2006079) with an HP system on November 1, 2006 performed by HP: HP ProLiant DL380 G5, 2 processors / 8 cores / 8 threads, Quad-Core Intel Xeon Processor X5355 2.66 GHz, 64 KB L1 cache per core and 4 MB L2 cache per 2 cores, 32 GB main memoryThis system running with Windows Server 2003 Enterprise Edition and SQL Server 2005 delivering 8970 SAPS.

On February 12, 2007 there was a benchmark (SAP Certificate: 2007028) with a server configured as following: HP ProLiant DL380 G5, 2 processors / 8 cores / 8 threads, Quad-Core Intel Xeon Processor X5355 2.66 GHz, 64 KB L1 cache per core and 4 MB L2 cache per 2 cores, 32 GB main memory You see, it's the same system. This It was

operated with SuSE Linux Enterprise Server 10 and Oracle 10g. This system yielded 9000 SAPS.

The difference of two identical configured system with different OS and DB was just 0.34 percent. It's just an educated guess, but the choice of database or operating system just seems to have a minor impact to the SAPS number, at least when comparing Linux+Oracle and Windows+MSSQL.

Hypothesis 2: The MaxDB database from SAP delivers more performance than Oracle 10g or MS SQL. This would be the next obvious conclusion. Sun used a different database than HP. But there is some data in the benchmarking list that suggest something else. On April 19, 2008 HP performed a benchmark (SAP certificate: 2008025) on the following configuration: HP ProLiant DL380 G5, 2 processors / 8 cores / 8 threads, Quad-Core Intel Xeon Processor X5460, 3.16 GHz, 64 KB L1 cache per core and 6 MB L2 cache per 2 cores, 32 GB main memory. This benchmark was performed with MaxDB 7.7 and yielded 11400 SAPS.

On October 30, 2007 performed an benchmark (SAP certificate: 2007064) with the same configuration, but used Windows Server 2003 Enterprise Edition and SQL Server 2005. This yielded 12,180 SAPS.

It's yet another educated guess, but the usage of MaxDB seems to slowdown the system. There is no proof for the hypothesis that MaxDB is better than Oracle 10g or MS SQL.

Conclusion: There are two remaining hypotheses. The first one: "Solaris 10 explain the difference between the Sun SAPS number and the HP SAPS number and is vastly better than Windows and Linux in regard of SAP". I can't support or dismiss it with numbers, because there is no SAP benchmark in the list which only differ in the operating system but not in hardware and database. The second one: "The X4600 is vastly better than the HP ProLiant DL785 G5 in regard of performance". As there is no benchmark of both system with the same combination of OS and database i can't support or dismiss it as well.

Anyway ... the proof of any of the both both hyptheses would be very good for Sun. Either our OS is vastly better or our system is vastly better. But both would proof: Engineering isn't futile.

Posted by Joerg Moellenkamp in Sun at 15:50

links for 2008-12-13

5.99 clocks per byte in C++! | The Skein Hash Function Family

(tags: cryptography skein)

Posted by del.icio.us in del.icio.us at 12:00

Friday, December 12. 2008

My first Sony ...

Inspired by local, i searched for a really old device, and i found it ... Sony HitBit 75B. My first computer. It wasn't the first one i used. That was a Sinclair ZX81 (Just for the younger people: That was a Zilog Z80 based system with 1024 bytes of main memory. But the computer just consisted out of 4 chips). But the Sony was my first own computer.

I've talked enough about myself ... what was your first computer?

Posted by Joerg Moellenkamp in Braindump at 20:27

Extreme Overclocking of an AMD Phenom II

There is an interesting report at xtremesystems.org with some photos about overclocking an AMD Phenom II to almost 6.3 GHz. Okay ... you need liquid nitrogen to do so, so it isn't really relevant for most people, but 4,13 GHz with air cooling is interesting as well. An observer writes: - We were able to run Crysis @ 6,1 GHz (2 loops CPU-test)
- Max CPU-Z screenshot and validation we got with 4 cores was 6280 MHz (no dirty tricks)
- With air cooling (Cooler Master Z600) highest we got Crysis stable was @ 4,13 GHz (1,58 V) Looks like the AMD Phenom II has some potential in regard of it's frequency.

PS: By the way ... somebody out there with a spare UltraSPARC T2 system? I would purchase some liquid nitrogen

Posted by Joerg Moellenkamp in The IT Business at 19:38

Memory wall? Pin Wall!

There is some discussion about many-cores hitting a "memory wall". Robin Harris writes about it in his article Many-cores hit the memory wall. But i think most of this article think not for enough. Who says, that you just can have many cores. You can have many-memorycontrollers as well. The future solution for memory is perhaps not the fastest bus. It's the bus that transport as much data as possible with the fewest possible amount of socket pins. The real wall may be the pin budget of modern processors as the number of memory controllers is limited by the possible connections to the outside world.

I don't think there is a wall at 16 cores in processor design because of memory bandwidth mandated by some basic principles, there is just a economic limit of manufacturing costs that prevents you from simply increasing the pin count and thus the numbers memory channels.

And so the problem of such articles like Jon Stokes "Analysis: more than 16 cores may well be pointless" is just the point, that they think in x86 technology and x86 economics. And even there he might be only correct when you just multiply the cores. There is a reason for connecting three channels of memory to an Intel Nehalem.

At the end the memory-wall discussion is somehow a validation of a Sun technology. You have to do something different. The idea of putting 4 memory controller on a T1 for example. The threading concept in the UltraSPARC T1 and T2 was introduced for a reason, too. The basic idea of this concept was "Live with the limited memory bandwidth and with the latency of memory and make the best out of it". The T1/T2 just switches to another thread in case of a memory stalled thread.

Posted by Joerg Moellenkamp in The IT Business at 15:33

Opensolaris at Heise

Nice ... Opensolaris on the home page of Heise.de, one of the largest IT-news portals in germany:

Posted by Joerg Moellenkamp in General at 14:34

I´m a little bit upset ...

Okay ... the headline was my yearly "vast understatement". I´m really pissed off. I got out of bed at 04:30. Walked to my urban train station to airport at 05:15. There is a new train line to airport ... really nice. But: The quick checkin system told me: Your flight is delayed ... not for 10 minutes ... not for an hour. The delay is 2 hours! At least i don´t have to sit around in the plane as at my last flight. Fsck!

Posted by Joerg Moellenkamp in Business Travel at 06:25

Wednesday, December 10, 2008

Opensolaris 2008.11 launch event at 9am PST

At 9am PST on opensolaris.com.

Posted by Joerg Moellenkamp in Solaris at 18:05

Official announcement of OpenSolaris 2008.11

The .iso's for OpenSolaris 2008.11 are available for download now for a few days, but today we had the official announcement: Sun Microsystems Launches Latest Version of OpenSolaris; Unveils Time Slider Visualization Tool and New Partnership with Toshiba to Build Line of OpenSolaris Laptops.

Posted by Joerg Moellenkamp in Solaris at 17:56

PR-Probleme ...

Es kursieren ja viele Scherze ueber chinesische Schriftzeichen-Tattoos, die bei Chinesen nur zu lauten Gelächter führen, weil der des Chinesischen unkundige Träger nicht weiss, das sein Bein üble Beleidigungen über die Mutter und die weitere Sippschaft enthält (dieser Gag wurde mir mal von einer Sinologin erzaehlt) . Bei der Max-Planck-Gesellschaft hätte ich aber erwartet, das zumindest mal ein chinesischer Gastwissenschaftler sich das anguckt: Max-Planck-Gesellschaft wirbt für heiße Mädchen

Posted by Joerg Moellenkamp in Fundsache at 15:16

HHOSUG

Nachdem ich jetzt einige Antworten bekommen, habe ich mich entschlossen, das mit der Hamburger (Open)Solaris User Group durchzuziehen. Mehr in den nächsten Tagen und für jedewede Unterstützung waere ich sehr dankbar!

Posted by Joerg Moellenkamp in Solaris at 14:57

8 socket benchmarking revisited

The fight for the best 8 socket benchmark in SAP is in it's second round. In the left corner ... HP "the inventing garage" ProLiant DL785 G5 with an fighting weight of 72 kg or 160 pounds. The system presents:HP ProLiant DL785 G5, 8 processors / 32 cores / 32 threads, Quad-Core AMD Opteron Processor 8384, 2.7 GHz, 128 KB L1 cache and 512 KB L2 cache per core, 6 MB L3 cache per processor, 128 GB main memoryIt runs SuSE Linux Enterprise Server 10 and Oracle 10g.

In the right corner ... Sun "The revenge of Andy" Sun Fire X4600M2 with an fighting weight of 40kg or 88 pounds. The Sun system is equiped with:Sun Fire X4600M2, 8 processors / 32 cores / 32 threads, Quad-Core AMD Opteron Processor 8384, 2.7 GHz, 128 KB L1 cache and 512 KB L2 cache per core, 6 MB L3 cache per processor, 128 GB main memoryThe system is powered by Solaris 10 and MaxDB 7.6

After fighting the following results were gathered from the fighting judges: Sun X4600M2 with 39,270 SAPS and 7,825 SD users (Certificate: 2008070). The HP was able to gather 35,400 SAPS and 7,010 SD users. (Certificate: 2008064).

The winner by a clear margin: Sun "The revenge of Andy" X4600M2! It should be noted that the Sun system already used the extra soft boxing gloves called "unicode" , whereas HP used the old and outdated boxing gloves "non-unicode". There is already an discussion raging about the usefulness of results based on those old gloves. Albeit there deliver more punch you can't use it in more and more fights in the real world.

Okay, fun aside: This benchmark result is interesting, albeit you can't use this result to compare both machines. In the last benchmark the system HP system used Windows Server 2003 Enterprise Edition and SQL Server 2008. The HP system just had 88.24 percent the performance of the Sun System. In this benchmark the HP system had 90,15 percent

of the Sun performance. Two different operating systems, two different databases ... but a similar difference? It's just a speculation ... but maybe the HP implementation of 8 sockets is vastly less efficient. Why vastly? You have to remember, that unicode costs you 15% performance. It would be nice to see a result for the ProLiant DL785 G5 with 6.0 unicode, Solaris 10 and MaxDB to get a clearer view to the issue.

Update: I dugged down into this benchmark. Albeit i have still no hard proof for my thesis, the operating systems or the database seem to be just minor factors to the SAP benchmark. I summarized my findings in "Sun X4600M2 vs HP ProLiant DL785 G5 revisited".

Posted by Joerg Moellenkamp in General at 10:09

Tuesday, December 9, 2008

Airbus Auto Brake

Many people in flight communities say "Don't compare your MS Flight Simulator experiences with real flying" and they are correct about that. But sometimes i have the strange feeling that modern aviation gets nearer to the Flight Simulator (as long the computers are functional). Flight Global writes about an interesting feature ... the Brake-to-Vacate: The BTV system shows pilots two lines on the electronic airport diagram on the A380's navigation display. The first line shows the minimum stopping distance for dry pavement (assuming no reverse thrust). The second line shows the minimum stopping distance for a wet runway. Pilots can then select their desired turn-off point on the display using a cursor, and the system internally computes the required auto-braking deceleration profile. Contaminated runway scenarios will be addressed in later versions of the programme, says Tarnowski. The system even warns the pilot, if the remaining runway doesn't suffice to stop the aircraft. The pilot can order an go-around in that case. At the moment this feature is integrated into the A380, but the article states, that this feature will go into the A320 and A330/340 as well.

Posted by Joerg Moellenkamp in Aviation at 19:17

Sun servers shanghaied ...

Our servers are now available with the new AMD chips (codenamed Shanghai): Sun x64 Servers with New Quad-Core AMD Opteron Processors Boost Performance Up to 35 Percent for Enterprise, HPC and Web Applications

Posted by Joerg Moellenkamp in Sun at 18:11

Support for encrypted lofi devices

With the integration of the code developed for PSARC 2007/001 Opensolaris supports encrypted lofi devices: The lofi(7D) driver will gain the ability to encrypt/decrypt the raw blocks. This provides parity with Linux and MacOS X which both have encryption support in their equivalent to the lofi(7D) driver. You will find additional informations at the PSARCs onepager.

Posted by Joerg Moellenkamp in Solaris at 15:21

A more sane comment about Suns future

Interesting comment about Suns future at the Enterprise Storage Form - Sun May Surprise Its Critics: The truth is that Sun has much more time than the urgent tone of much of the coverage would have you believe, and has no need for a white knight or a federal bailout. The company's dramatic September-quarter loss was largely a non-cash write-down — the company's actual cash levels dropped \$242 million in the quarter. Not pretty, but with about \$2 billion more in cash and liquid investments than debt, the company could survive for a couple of years at its current cash burn rate. There are blue chip companies that would kill to be in that position right now.

Posted by Joerg Moellenkamp in Sun at 14:50

links for 2008-12-09

Oracle prices themselves out - Another guy named Mike

(tags: oracle)

Posted by del.icio.us in del.icio.us at 12:00

Monday, December 8. 2008

Hamburg (Open)solaris User Group

Ich fange ja so langsam an, meine Vorsätze für das nächste Jahr zusammenzustellen. Einfach mal die Dinge zu sammeln, die ich nächstes Jahr vor habe. Ein ganz wesentliches Ding rotiert da schon etwas länger in meinem Kopf. Ich überlege momentan, ein Projekt für den Hamburger Raum zu starten, aber bevor ich da weitere Gedanken reinstecke, würde ich gerne eure Meinung hören.

Ich möchte in 2009 eine Hamburger (Open)solaris User Group ins Leben rufen. Es geht hier hierbei um eine Möglichkeit die Solarisnutzer in der Metropolregion Hamburg untereinander zu vernetzen. Was haltet Ihr davon? Würde euch so eine Gruppe zum Erfahrungsaustausch interessieren? Wer nicht via Kommentar antworten möchte, kann seinen Kommentar via Mail schicken. Mailadresse steht rechts.

Posted by Joerg Moellenkamp in Sun at 18:42

Storage market in the third quarter of 2008

The IDC Worldwide Disk Storage Systems Quarterly Tracker report for the third quarter of 2008 showed a positive development for Sun. The Register reports: Although Sun is listed by IDC in the 'Others' category, it grew disk storage system sales 25 per cent in its second consecutive quarter of revenue growth, with 16.1 per cent growth in external disk systems. The company sold more disk storage for Unix servers than any other supplier has done, it claimed, for the past 20 consecutive quarters. If the new Open Storage products take off then it could start approaching Hitachi and NetApp. That would be a noteworthy feat.

Posted by Joerg Moellenkamp in Sun at 16:15

links for 2008-12-08

Wolfram Blog : The Incredible Convenience of Mathematica Image Processing

(tags: visualization visualisation tutorial tools trends)

Posted by del.icio.us at 12:00

Sunday, December 7. 2008

About Time Slider

This isn't something i expected to read in regard of Desktop Computing:Sun shows Microsoft how it's done -Time Slider is what Vista's Previous Versions should have been.

Posted by Joerg Moellenkamp in Solaris at 16:44

Home NAS mit OpenSolaris

In der gedruckten Ausgabe der aktuellen c't wird ja der Aufbau eines Filers auf Basis von PC-Technik dargestellt. Auch auf Solaris kommt die Sprache. So richtig lang war dieser Teil allerdings nicht: Ein wesentlichlich längerer Artikel ist im Onlineauftritt zu finden: OpenSolaris als Fileserver. Herr Dr. Diedrich schreibt dort unter anderem: Dass OpenSolaris für diese Aufgabe taugt, haben wir im Vergleich mit anderen Betriebssystemen gesehen: Mit einer Leserate von 111 MByte/s und einer Schreibrate von 86 MByte/s ließ das Sun-Unix alle anderen freien Systeme hinter sich, lag gleichauf mit dem Windows Server 2008 und erreichte Werte nahe dem in einem Gigabit-Netz maximal möglichen Datendurchsatz. Der Artikel enthaelt auch eine komplette Anleitung, wie man mit Solaris einen Filerserver aufbauen ann. Ein sehr wohlwollender und schoener Bericht zu diesem Thema.

Posted by Joerg Moellenkamp in Solaris at 13:53

Saturday, December 6, 2008

45.000 images

Wired reports about an interesting way to shoot a music video - just take thousands of photos: Eschewing a video camera, he took 45,000 photographs with a Nikon D200 DSLR (digital single-lens reflex) camera and stitched them together to create the illusion of video. With a DSLR and a computer the limit is just your imagination.

Posted by Joerg Moellenkamp in Photographie at 20:12

On heavy rotation at the moment

Strange mood at the moment ...

No fear, i will not do that stuff from the video ... i have music for that ...
(at iTunes)

Posted by Joerg Moellenkamp in Music at 19:35

in retrospect

May 2005 - i've lost my fear of flying

December 2008 ...

(inspired by a conversation with a good friend)

Posted by Joerg Moellenkamp in Aviation at 19:15

Friday, December 5, 2008

Ars Technica about Opensolaris 2008.11

A surprinsgly postive review about Opensolaris 2008.11 at Ars Technica - Hands-on: OpenSolaris 2008.11 a major step forward for Sun:OpenSolaris 2008.11 is a very big step forward for Sun, [...] Sun has some creative solutions to address that problem and clearly has a strong grasp of user expectations. [...] It offers some different technologies, a different community, and a different approach to development. All of those things bring value to the broader ecosystem of open source software platforms. I think that it is impossible to overstate the importance of competition as an instrument of progress. Ryan Paul just critize two points: Not enough packages and not enough drivers. Both points will vanish over the next time.

Posted by Joerg Moellenkamp in Solaris at 21:00

Flagday for Crossbow

One of the most interesting upcoming features of Solaris found it. The first parts of the results of PSARC/2006/357 Crossbow - Network Virtualization and Resource Management were integrated yesterday into Opensolaris Build 105. Crossbow provides the building blocks for network virtualization and resource control by creating virtual stacks around any service (HTTP, HTTPS, FTP, NFS, etc.), protocol (TCP, UDP, SCTP, etc.), or Virtual machine (Containers, Xen and LDom based hypervisor xVM virtual machines).As you can see from the year in the ARC case number, it's one of the large projects at Opensolaris, as it was under steady development since that time. You will find more information at the projects webpage on opensolaris.org.

Posted by Joerg Moellenkamp in General at 18:21

links for 2008-12-05

GlassFish Wiki: JRuby

(tags: jruby rubyonrails ruby rails glassfish sun server)

Ihc - Photocopier and Laser Printer Hazards

(tags: laserdrucker)

Particle Emission Characteristics of Office Printers - Environmental Science & Technology (ACS Publications)

(tags: laserdrucker)

Risiko im Büro - Gesundheitsrisiko Laserdrucker - Gesundheit - sueddeutsche.de

(tags: laserdrucker)

Posted by del.icio.us in del.icio.us at 12:00

Yet another delay in the 787 program

The Chicago Tribune reports:The Wall Street Journal reported Thursday that the first of the largely composite jets isn't likely to be delivered before mid-2010, more than two years behind schedule.

But a source close to the program said the jet's launch could occur even later, especially if flight-testing for the all-new aircraft takes longer than the aggressive schedule that Chicago-based Boeing has planned.When it rains, it pours.

Posted by Joerg Moellenkamp in Aviation at 09:03

Thursday, December 4, 2008

Flight to Munich

The flight to MUC yesterday wasn't really a nice experience. The meeting was good. But the travel to MUC yesterday was a test how long you are able to keep your countenance:

The pilot was correct with his words: "The airport of Hamburg was totally surprised by the sudden winter, as they had one last year already". We had to wait for one and a half hour in a fully seated plane.

There is a large insurance company in Germany. And there was a single person in the plane who was able to ensure singlehandedly that I will never opt for an insurance from this company. The first sin: arrogance. There are several reasons, why you are obliged to switch off your electronic devices while start. This man kept his laptop on while start and while landing. This idiot was well aware of this regulation, because he hid the active notebook from the flight attendants by almost closing the lid. Not because he had to finalize his mail. No ... he looked after landing into the screen for 2 seconds and closed it. What an idiot ...

The second sin: The same person wasn't aware of the fact, that reading an internal report isn't a good idea in a plane. I was able to read the report without problems. Interesting lecture ... but I was clearly not the intended audience.

I know ... I'm just a FTL and my ticket in the booking class L wasn't that expensive. But that isn't a to ram the trolley twice into my shoulder. That isn't a reason for the flight attendants to run into my shoulders. I didn't overlap my seat to the aisle because I'm that fat or because I love to be touched by the butts of flight attendants. When you place three people in the height range from 1,90 to 2,00 meter with wide shoulders this is simply an expression of laws of basic physics.

Dear Lufthansa, is it really necessary to add additional delay for passengers by parking at a napron position for a flight that is already one and a half hour late?

There is something called "situational awareness". Situational awareness is the awareness for the world around you. To know when a body moves in the direction x with the speed y that this body may be located at the place z and that's a dumb idea to move direct way of this body instead of waiting until the body has passed (basic collision avoidance). Is this situational awareness a rare gift or is it just outright disrespectful that leads such persons. Just my good manners prevent me to simply move that people out of my direction. I should ignore my fellow human beings as well. It's the default behaviour today. In airports, at train stations, in the urban train, in the supermarket. I really hate this behaviour ...

The online boarding pass is a nice idea. It is just problematic, when the email client of your mobile is so slow, that it's faster to reprint your boarding card at the Quick-Checkin terminal than to open the mail on your mobile.

Well ... enough ranting for today

Posted by Joerg Moellenkamp in Mindfuck at 21:57

Leaked Airbus "787 lessons learned" presentation

The flightblogger was able to publish an article in the "once-in-a-lifetime-scoop"-class. He was able to obtain an internal Airbus presentation, that leaked somehow. This presentation is named "Boeing 787 Lessons Learnt October 2008" and it's based on informations and rumours from suppliers and openly available information. Some of the slides are based on Boeing presentations. This is the reason, why some people cry "Foul!" at the moment ... but this kind of intelligence is quite normal in business (Sometimes all you need is to google for presentation to get very interesting informations about IBM newest FUD).

But the important point is something different: The 787 program seems to be in a horrible state according to this presentation. Just half of it would be really bad from my perspective. Let's wait and see, how Boeing reacts to this document.

Update: Two sources near to the A380 program told me, that this presentation is a fake, albeit a good fake.

Posted by Joerg Moellenkamp in Aviation at 21:29

A380 too quiet

Okay, we all laughed about the comment of a Boeing guy who stated that the A380 is too quiet. But there seems to be an "too quiet" problem in this aircraft. There is a crew restroom in the A380 for long distance flights. Normally it's located

in the front of the aircraft, but in Emirates machines have this room in the back. The problem: They can't sleep - the passengers are too loud and there are no engine noises to overlap the passengers. Flightglobal writes: The pilots say that the lack of engine noise in the A380's cabin compared with other long-haul airliners means they are constantly disturbed by sounds created by passengers, such as crying babies, flushing vacuum toilets and call bells. Passengers also mistake the rest area for a lavatory, and pull the door handle. (Thanks to Kris for the hint)

Posted by Joerg Moellenkamp in Aviation at 16:42

Overpopulated

Looks like i'm using too many small tools ...

from right to left: Spotlight, Clock, Language Setting, Volume, Airport, Bluetooth, Monitors, Spaces, Time Machine, PPPoE, Dial-In networking, Accessibility, iSync, Growl, Caffeine, Twitterrific, Nokia Uploader, Adium X, Jing, Camouflage, smcFancontrol)

Posted by Joerg Moellenkamp in Apple at 16:34

Deutschsprachiger Anwenderbericht zu Opensolaris 2008.11

Ein sehr wohlwollender, deutschsprachiger Bericht über Solaris 2008.11 ist bei computerseite-spezial zu finden: Mir als Heimanwender macht der Umgang mit dem Unix OpenSolaris viel Freude. [...] Die aktuellen Einschränkungen des Systems betrachte ich optimistisch als einen vorübergehenden Zustand. Für mich persönlich ist mit OpenSolaris tatsächlich eine ?neue Sonne? am Himmel der Betriebssysteme aufgegangen.

Posted by Joerg Moellenkamp in Sun at 14:51

Guided tour through new features in Opensolaris 2008.11

Roman Strobl created a screencast for the new Opensolaris 2008.11 release. Really worth a look.

Posted by Joerg Moellenkamp in General at 09:01

More than 250.000 NFS IOPS

Brendan Gregg did an interesting benchmark with our new storage systems to show what an 7410 single-head can deliver: A quarter million NFS IOPS. This number was generated with a benchmarking configuration - 1 byte blocksize. This isn't really realistic - it's just for benchmarking. But he delivers a more meaningful result in his text, too: I've reached over 145,000 NFSv3 read ops/sec - and this is not the maximum the 7410 can do (I'll need to use a second 10 GigE port to take this further). The latency does increase as more threads queue up, here it is plotted as a heat map with latency on the y-axis (the darker the pixel - the more I/Os were at that latency for that second.) At our peak (which has been selected by the vertical line), most of the I/Os were faster than 55 us (0.055 milliseconds) - which can be seen in the numbers in the list on the left.

Posted by Joerg Moellenkamp in General at 07:43

Wednesday, December 3. 2008

New features of Opensolaris 2008.11

There is a great list of all the new features of Opensolaris 2008.11 at the opensolaris.com website: [What's New in OpenSolaris 2008.11?](#)

Posted by Joerg Moellenkamp in Solaris at 20:11

Tuesday, December 2, 2008

Adam Leventhal about the storage business in the light of the HSP

Adam Leventhal wrote an really insightful piece about the potential effect of the Hybrid Storage Pool to the storage business in "Casting the shadow of the Hybrid Storage Pool":What we have today though integrates flash in a way that changes the landscape of storage economics and delivers cost efficiencies that haven't been seen before. If the drives manufacturers don't already, it can't be long until they hear the death knell for 15K RPM drives loud and clearWhen the performance of a storage architecture is ensured by SSD there isn't really a need for low capacity 15k discs ... 7200 rpm (or even lower) drives with high capacity are more than enough.

Posted by Joerg Moellenkamp in Sun at 20:52

IBM benchmarking - now with 100% more nonsense

There must be a standard course at the IBM marketing school. The tricks of the benchmarking are the same again and again. At least it doesn't get boring to read at the IBM website. Now i found another benchmarking stunt at their website. It's called "Move Up to IBM Power Systems".

It was really fun to read it. IBM Benchmarking - now with 100% more nonsense. One important disclaimer: The following article just contains the stuff that was obvious at my first read. I have a daytime job that isn't dissecting IBM marketing stunts

Standard Trick: Comparing different classes of systemsThe standard marketing tactic is: When you want to make a statement on price use the cheapest system, when you want to make a statement about performance compare with the most expensive high-end system.

This is a nice example from an actual IBM marketing campaign.

This table is from the IBM website. They want to suggest that the p570 and the T5440 are similar systems. But they compare the performance of several systems in this table. In one of the following sentence they write: Compare the 32-core 4.2GHz Power 570 to the 32-core T5440 and again the winner is clearly the 570 with 1.9 to 2.7 times the performance."They only forget to mention, that the 32 core 4.2 Ghz Power 570 is a 16 RU system, while a T5440 takes just 4 RU. A system comparable to the T5440 would be the p550. It's the in the four socket class, too. And what do you see. Half the SAP SD performance and half the SPECjbb performance.

On the top end they use the p570 to state a huge number for performance and to tell you some stuff about RAS. Of course the p570 is better here. It's a much larger system with different design goals. But they forget to state an additional fact in their marketing stunt: The price of the p570. They didn't state the price of the p570 throughout the document. Just to give you a hint: Look at the TPC-C benchmarks for p570 systems. The activation of memory or one or two additional procs may cost as much as a complete T5220 or an T5440 in a smaller configuration. BTW: When you buy memory (the hardware) for a p570 you cant use it. You have to buy a memory activation as well.

So make it short: Such a comparison is just outright nonsense.

A really strange pricing comparisonBy the way: The comparative pricing is really strange. The IBM states in the article, that an T5440 costs you\$235,705. But when you look at Sun.com you will find a different pricing: An T5440 costs you \$132,995.00 list price with 128 GB and 4 8-core procs at 1.4 GHz (that's almost full blown). The solution of this miracle can be found in the Footnote 7.

At first they state: Prices based on configurations used in the SPECint_rate2006 benchmark including processors, memory and disk.Almost at the end of the footnote, they state:Sun SPARC Enterprise used Sun StorEdge 6140 Array for disk.Okay, that's fine with me, albeit disk configuration is irrelevant for SPECint. I assume it was just ease of configuration for the benchmarking guys. But hey, it's an Oracle pricing comparison and a 6140 makes sense. So this pricing for the T5440 was increased by the price for the 6140. Additional they took the the 256 GB configuration.

But - and now it gets interesting - they state for their own system: All IBM systems used internal storage.The disk subsystem for pSeries 560: 4x146 GB SAS 15K RPM. Memory? 64 GB. Nice try ... really nice try.

The system isn't comparable at all. SPECint configurations aren't any hint for a Oracle configuration. Thus is pricing comparisons about the system configuration are totally nonsense at all.

I can promise you one thing: At any realistic database size the T5440 configured for this comparison system will outperform the p560 configured for this comparison by a vast margin. There is one sizing rule in database sizing different to SPECint. You can substitute memory just by more memory. And two boot disks aren't really comparable to a 12-disk shelf with redundant FC controllers and 2 GB of NVRAM cache. At any working set size beyond 64 GB the p560 would wait with all its gigahertz for the hard disk. And even with a workingset size smaller than 64 GB the 6140 would outperform the 4 disk configuration at write loads due to it's NVRAM cache

Software pricing - or: The Oracle ProblemIt's correct, that Oracle is one of the weak points of the CMT technology. Or to be exact - not the Oracle software is the problem - Oracles licensing model is in favour of a few big cores. IBM doesn't get bored to point on this issue. But they forget some important facts. At first: There is other software than just Oracle, and Oracle is one of the last strongholds of per-core licensing. You will see "per-seat", "per-employee" or "per-named user" everywhere. SAP? Per user! The Oracle "Application Specific Full Use" licenses for SAP? Per user! Lotus Notes or Lotus Domino? You can license it per user! Several other software products? Per User!

So it's not relevant how many cores you have, it's relevant how many people use a software at your company. And this enables you to use the most efficient technology over the long time.

And they forgot to mention another alternative at all. The T5440 is eligible for the Oracle Standard Edition. The advantage: A per-socket licensing. The p560 16-core is roughly comparable with the T5440 in SAP-SD but it's a 8 socket system, thus you have to use the Enterprise Edition.

You think this is a non-argument? Maybe ... but i saw an IBM presentation using the same argument to sell their QCM modules by using Standard Edition on their systems and Enterprise Edition on ours. We can play the same game.

You should think about an additional factor: Oracle licenses are perpetual. Let's assume you have an older system in your data system you want to migrate to a newer one (For example your older SAP system on a 6800). You have already Oracle licenses for some amount of cores in your house. You can transfer this old licenses to a new system.

Additionally we should consider here, that we compare list prices. Oracle sales reps are not necessarily friends of IBM because of the DB2 menace in their account. You could use that.

And at the end: Many applications can run on Mysql or Postgres as well. There are licensing models for Mysql allowing you to install as many systems with mysql as you want at the price of a single Oracle core license. When you really want to save money, you should think in this direction.

The 11% claimIBM likes to claim that they had an 11% increase of pSeries business. What they forget to write: They have merged iSeries an pSeries business to a new pSeries spanning both of the two old product lines. The iSeries lost 82%. Well ... i leave it to you where this increase of 11% may find its explanation.

IBMs AIX vs. Solaris claimsWell just a few notes: The point that Solaris implemented RBAC in 2000 isn't an excuse for implementing it in AIX just this year. The WPAR Live migration is just a hack, as i explained in this article and Volume Managemnt doesn't entered Solaris with ZFS, we had an product called Solstice Disk Suite dating back into the last century There are many "

ConclusionI will stop to dissect the IBM statements here. There is more utter bullshit in this website from IBM but read it for you self. You will find it at their website.

Posted by Joerg Moellenkamp in Sun at 19:12

Opensolaris 2008.11

The released version of Opensolaris 2008.11 is available at the usual locations for download. It contains several interesting new developments like the Time Slider, Fast Reboot and many many more new features ...

Posted by Joerg Moellenkamp in Solaris at 13:41

This made my day

From a discussion about NAT in IPv6: the only good NAT is a dead NAT

Posted by Joerg Moellenkamp in Sun at 12:56

INTJ

Nice tool to find out the Myers-Briggs personality type of your blog. Or better: The personality type you show in your blog as it "reads" your blog: c0t0d0s0.org is an INTJ. Okay, this matches my own experiences with such tests - they sort me consistently in the INTJ or ENTJ groups

Posted by Joerg Moellenkamp in Fundsache at 07:46

Monday, December 1, 2008

History repeating?

The TLB bug almost ruined the AMD Phenom franchise. But it looks like if Intel has a similar problem with the Nehalem. According to Fudzilla the Nehalem is haunted by a TLB bug too and refers to the errata document of this processor. It looks like history repeating.

At page 16 of the the Intel Core i7 Processor Extreme Edition Series and Intel® Core™ i7 Processor they write:AAJ1. MCI_Status Overflow Bit May Be Incorrectly Set on a Single Instance of a DTLB Error

Problem:A single Data Translation Look Aside Buffer (DTLB) error can incorrectly set the Overflow (bit [62]) in the MCI_Status register. A DTLB error is indicated by MCA error code (bits [15:0]) appearing as binary value, 000x 0000 0001 0100, in the MCI_Status register.

Implication: Due to this erratum, the Overflow bit in the MCI_Status register may not be an accurate indication of multiple occurrences of DTLB errors. There is no other impact to normal processor functionality. But in the clarification for errata AAJ1 on page 37 of the same document they state:In rare instances, improper TLB invalidation may result in unpredictable system behavior, such as system hangs or incorrect data. Developers of operating systems should take this documentation into account when designing TLB invalidation algorithms. For the processors affected, Intel has provided a recommended update to system and BIOS vendors to incorporate into their BIOS to resolve this issue. Doesn't sound like There is no other impact to normal processor functionality. I don't think, that this bug has the scale of the AMB TLB bug, but i think customers are a little bit sensitive today when they hear TLB and bug in the same sentence.

Update: AAJ42 on Page 27 is interesting, too:Under certain conditions when C6 and two logical processors on the same core are enabled on a processor, an instruction fetch occurring after a logical processor exits from C6 may incorrectly use the translation lookaside buffer (TLB) address mapping belonging to the other logical processor in the processor core.

Posted by Joerg Moellenkamp in The IT Business at 20:01

NAS at home

While preparing the next part for the configuration walkthrough of the Sun 7000 my mind was playing with an desire: I want this interface on an soho device. I played around with many NAS boxes but there wasn't a single interface in the class of the Fishworks interface. Intel Atoms are cheap enough to build real soho NAS devices from it ...

Posted by Joerg Moellenkamp in Solaris, Sun at 13:40

links for 2008-12-01

Code Instructions: Java Performance
A good answer to all this Java is slow trolls
(tags: performance java)

Posted by del.icio.us in del.icio.us at 12:00