

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

Monday, July 28. 2008

### **NYC**

This blog isn't dead. It's on vacation. I simply don't have an internet connection. So don't touch the unsubscribe button. Main blogging activities will resume on next sunday.

Posted by Joerg Moellenkamp at 14:30

Saturday, July 26. 2008

**Child of fortune ...**

I just came back from checking in my luggage. I've got an exit-row seat for the flight to JFK. Not business-class comfort, but much more space for my legs and that's all i want from my seat

Posted by Joerg Moellenkamp in General at 10:36

Friday, July 25. 2008

### **Depressing**

There was one fact about my trip to Andorra that was really depressing: The shoes were cheap, but no shoes in a size above 45. Today i walked to a shoe shop in Hamburg, saw a nice pair of Geox, walked to the shoe sales man, asked for a pair in my size. The answer: "Sorry, we don't have this shoe in sizes above 45" ... fsck ....

Posted by Joerg Moellenkamp in Mindfuck at 22:22

### **Randy Pausch died at 47**

Perhaps you remember the "Last Lecture" of Randy Pausch, a Carnegie Mellon University computer scientist. I wrote about it some time ago. Randy Pausch's battle came to it's inevitable end: He died at 47.

Here the video that made him famous to a large audience:

Posted by Joerg Moellenkamp in General at 19:59

### **The next president has spoken - hopefully**

I've viewed the Berlin speech of Barrak Obama yesterday evening on CNN (the german networks had terrible translators). I really think, i saw the next president of the United States. From my perspective it was a good speech. Refreshingly different than the other politicians i saw before (the US ones and the german ones as well). I really hope, your get it right this time, dear US voters ... i really had some doubts about the mental status of the US people at the last election.

BTW: I don't think that this age is a problem. Long and vast Experience is often not much more than an excuse for stubborn narrow-mindedness. You have advisors for this task, knowing much more about an issue than you will know ... you just take care, that your advisors haven't an own policy (but Bush (the younger one) got even relected with such advisors).

Posted by Joerg Moellenkamp in General at 07:56

### **9th Annual System Administrator Appreciation Day**

Today is the System Administrator Appreciation Day. Appreciate the sysadmins of the world. Without them, your display would be black, the banking systems with your money would broke down, you couldn't load porn through the internet, you couldn't buy shoes or gadgets via Internet and so many other everyday occurrences. So ... appreciate them.

Posted by Joerg Moellenkamp in The IT Business at 07:46

Thursday, July 24. 2008

## **Simpsons**

Es gibt mehr als 414 Folgen der Simpsons, doch Pro7 schafft es, immer wieder die selben Folgen zu senden ....

Posted by Joerg Moellenkamp in Movies at 18:55

## **links for 2008-07-24**

Universal Automatic Calculator - Wikipedia  
(tags: history)

Posted by del.icio.us in del.icio.us at 13:33

## **Linux, DTrace and the NIH syndrome**

I wrote about the the strong "Not invented here" syndrome in the Linux community a year ago in DTrace, systemtap and a brief history of "NIH". Nothing has changed so far ... systemtap is still not admin-usable, dtrace find it's way in more and more environments and the linux community is still discussing about their way. Paul Murphy reports about the reignited discussion in his article "DTrace and the Linux bunker mentality" (referring to Bryan Cantrills DTrace on Linux article)

There is so much talent in the linux community. I don't really understand why nobody is analysing DTrace. With analysing i don't think about "Oh, it measures things". I think about analysing all the concept, the ideas, and the implementation. And after this implementing something on par with DTrace. But that doesn't take place at the moment.

Posted by Joerg Moellenkamp in General at 11:49

## **Oracle and CMT**

Oracle and CMT are often a natural choice. Whenever you have many parallel requests and the latency isn't a key performance indicator, you should give it a try. But sometimes there are loads, that should scale well on CMT systems but they don't scale well. In most cases there are some quirks in the SQL statements that makes the code single- or few-threaded. Glenn Fawcett summarized some great tips for Oracle and CMT in a series of blog articles to overcome such problems.

(via: Stefan Hinker)

Posted by Joerg Moellenkamp in Sun at 11:15

## **Google searches**

Strangely i found more and more google searches with my name in the logfiles of my blog and in the referer list of Xing.com in the last few days. Just in case: I haven't done it and it's not my fault ...

Posted by Joerg Moellenkamp in Blogosphere at 11:08

## **Sun will support \*AMP**

In the press release issued yesterday was another intereting news: Sun released a new version of it's Coolstack. It consists out of: The primary components in the Web Stack include (partial list) the Apache HTTP Web server version 2.2.8, Apache Modules Memcached 1.2.5 (distributed memory object system), MySQL 5.1 Database, lighttpd Web

server v 1.4.18, Tomcat Servlet engine 6.0.16, PHP 5.2.5, Ruby 1.8.6, Rails 1.2.3, RubyGems 0.9.0, Mongrel 1.0.1, fcgi package, RedCloth (text parsing), Perl 5.8.8 and extensions, Squid proxy server 2.16.x. In the past it was somewhat problematic, the versions delivered with Solaris were supported, but they were selected for stability and were often bugfixed, but not updated. The versions in the Coolstack were really actual, but they were unsupported. But this changes now: Sun will announce an offer to support the complete webstack soon. At first (this quarter) for Solaris on x86 and SPARC, but later on for Linux and other operating systems.

Posted by Joerg Moellenkamp in Solaris at 07:10

Wednesday, July 23. 2008

## **Presto**

The supporting movie in front of Pixar movies is a good tradition. Presto is the one before Wall-E. Absolutly hilarious.

Posted by Joerg Moellenkamp in Movies at 21:28

## **Habitation**

What is habituation? When you borrow a portable navigation device from your brother, try to scroll through the menus and twitch through the maps like on your iPhone and think at first: "Fsck ... the display is broken" ... i want an iNav!

Posted by Joerg Moellenkamp in iPhone at 21:07

## **Thoughts about a customer visit ...**

Yesterday i visted a customer with severe space problems. You think: "Hey, my datacenter is full, too. Not a big deal". But those folks have really no space for euquipment, to be exact ... they have negative space for new equipment ... when they remove one or two systems the datacenter is absolutly full I thought at first, the people of the customer had a "paint it black" day when they talked about it at first ... but i have to admit, that they even vastly understated the problem. This challenge will be a hard nut to crack .... but is solvable .. i have already some ideas. Sometimes it's a good thing to work for a technology company, because some problem must be solved by technology. And can be only solved by technology.

Posted by Joerg Moellenkamp in Sun at 20:45

## **Sun Java System Web und Proxy Server opensourced**

Two further products gone the opensource path. Sun opensourced the Sun Java System Web und Proxy Server. Especially the webserver is interesting. It's not that known than Apache, but's nevertheless an excellent product.

Posted by Joerg Moellenkamp in General at 20:23

## **A great usecase for the DTrace IP provider**

Ben of cuddletech.com reported about the Dtrace IP provider as well. A commentator challenged him with "There is nothing unpleasent about the wonderfulness that is tcpdump! You'll need to put a lot of work in to match tcpdump's usefulness with Dtrace... ". And Ben answered to that ... with two simple examples of the usefulness of the IP provider. He shows in his answer how you can correlate a sent package to a application on your system or how to print a complete stack trace that led to the transmission of a package.

Posted by Joerg Moellenkamp in General at 13:11

## **Rambo**

Zumindestens Rambo - First Blood war aus meiner Sicht immer ein sehr unterschaezter Film. Ich glaube, das liegt auch daran, das man Sylvester Stallone erst nach Copland auch mal zugetraut hat in einem Film mitzuspielen, der nicht nur aus plumper Zurschaustellung von Gewalt bestand. Das man eben primär auf die physische Erscheinung von Herrn Stallone geguckt hat, als denn auf die Geschichte, die dort erzaehlt wird. Und ob man es glaubt oder nicht, zumindestens der erste Teil erzaehlt eine Geschichte.

Nichtidentisches hat nun mit dem aeusserst lesenswerten und gelungenen Text "Rambology - Mit John J. Rambo durch die Dialektik der Aufklärung" diese dumpfe Ahnung untermauert.

Posted by Joerg Moellenkamp in Movies at 12:58

## **About the effects of L2ARC**

Brendan Gregg wrote a good piece about the performance of L2ARC in ZFS L2ARC: The pool\_0 disks are still serving some requests (in this output 30 ops/sec) but the bulk of the reads are being serviced by the L2ARC cache devices - each providing around 2.6K ops/sec. The total delivered by this ZFS pool is 15.8K ops/sec (pool disks + L2ARC devices), about 8.4x faster than with disks alone. and Our average service time is between 0.4 and 0.6 ms (wsvt\_t + asvc\_t columns), which is about 20x faster than what the disks were delivering.

Posted by Joerg Moellenkamp in General at 12:36

### **DTrace IP Provider**

At the moment tcpdump, snoop and others are the predominant tool to analyse network traffic and network problems on a host. But with the IP provider Dtrace makes the first steps to offer a tool for this class of problems as well. It found it's way into Solaris with nv93. You should give it try and it's only one of the first steps ... more to come.

Posted by Joerg Moellenkamp in Solaris at 08:14

### **VMware announces free ESX - or: The gorilla effect?**

VMware announced a free version of their ESX server product. Well, this may be an effect of the 500 pound gorillas entering the business. Sun will announce xVM soon, Microsofts developed Hyper-V ... thus the market for it's hypervisor may get narrow. At the end, the auxillary tools earn VMwares money, thus giving away the hypervisor is a clever move.

Posted by Joerg Moellenkamp in The IT Business at 07:28

Tuesday, July 22. 2008

### **Project Mayhem**

Do you remember the end of Fight Club? The "Project Mayhem" culminates in the explosive demolition of the buildings of banking cooperations. Well, when you look about all this news about the financial sectors you may come to the idea, that you don't need a "Project Mayhem" anymore. You just have to wait for the implosion of the respective companies ... today Amex announced a decrease of 37 percent in their incomes. And Amex isn't really a mortgage company, it's more known for credit cards, and creditcards are somewhat related to the consume ... not a good sign ...

Posted by Joerg Moellenkamp in General at 21:56

### **Rumoren**

Eine gute Freundin meinte heute nachmittag: "Jörg, wann immer du mal ein paar Tage in deinem Blog nichts schreibst, rumort etwas in dir". Hmm, ja ... mag sein ... so fuehlt sich momentan auch mein Hirn an ... aber ich weiss noch nicht, was da am rumoren ist ... ich habe zwar ein paar Verdachtsmomente, aber nichts stichhaltiges ...

Posted by Joerg Moellenkamp in Braindump at 21:37

Sunday, July 20. 2008

### **Not really a weekend**

This weekend wasn't really a weekend. Worked on my part of the answer to a lengthy Request for Proposal ... hope that all this effort isn't futile.

Posted by Joerg Moellenkamp in Work at 22:27

### **The consequences of the Novell ./ SCO ruling ...**

There is some discussion in the media about the consequences of the ruling in the Novell vs. SCO lawsuit. Some media outlets like Inquirer think this may pose a problem. I think, this reports are the usual "only bad news are good news". IMHO nothing like that will happen. Opensolaris is an established member of the open source community. Novell has a business relying on open source. Trying to use this copyright to take out a competitor or trying to get a large heap of money would totally ruin a reputation in a core market that already got some scars by the Microsoft/Novell deal. This would be basically a second edition of the SCO/IBM lawsuit ... and you know what happened to the unix business of SCO. And besides of this ... normal companies don't sue each other, as this can lead to a mutually assured destruction thing too fast ...

BTW: In my personal opinion the industry should end this topic for one and all times and put the old unix source code into the public domain ...

Posted by Joerg Moellenkamp in Sun at 13:56

Saturday, July 19. 2008

### **Geekhumor - heute: Warum 70 werden ?**

Eben im Chat:

Ich:ach ... die midlife crisis habe ich schon

I.: jetzt schon?

Ich: ich werd doch nur siebzig ...

I.: Nur? Woher willst du das denn wissen?

Ich: ich will doch zumindestens sehen, was 19. Januar 2038 um 03 Uhr 14 Minuten und 8 Sekunden passiert dafür muss ich etwa 65 werden ... und ein paar Jahre extra wären ganz nett, um das restliche Chaos zu erleben ...

I.: LOLunglaublich ...Hat er doch tatsächlich Minuten und Sekunden dabei ... \*LOL\*

Ich: jo, 32 Bit sind halt endlich

I.: Jau ...

Update: Da ich eben schon per Mail gefragt worden bin, was es damit auf sich hat ... an jenem 19. Januar sind die 32 Bit des POSIX Zeitaehlers zu Ende. Zur Erklärung hier ein Ausschnitt der Wikipedia:POSIX zählt die seit dem 1. Januar 1970 abgelaufene Zeit in Sekunden. Am 19. Januar 2038 um 03:14:08 Uhr UTC wird die Anzahl der vergangenen Sekunden die Kapazität einer 31-Bit-Zahl überschreiten. Das 32. Bit wird laut Konvention dazu verwendet, positive und negative Zahlen zu unterscheiden (siehe Zweierkomplement), so dass die Zählung in den negativen Bereich springt und die Konvertierung zu Datum und Uhrzeit Freitag, den 13. Dezember 1901, 20:45:52 Uhr UTC ergibt.

Posted by Joerg Moellenkamp in General at 17:51

### **An established brand**

You know, that a brand is well established in the market, when the approx. 3-4 year old boy on the lap of his mother on the adjacent seat says: "The man has an iPhone". At least I think the boy mumbled something like that ...

Posted by Joerg Moellenkamp in iPhone at 16:06

### **The real reason for headscarves**

Some women with muslim faith wear a headscarf. As Germany is a very secular society(okay, besides some regions in Bavaria and seemingly near the town Vechta in Lower Saxony, but that's more a catholic issue ) it's relatively seldom. I thought in the past, that this is a public expression of a cultural and religious identity. I was totally wrong about that, as I saw the real reason for the headscarf today in a train: It's simply a hands-free mounting for the mobile. A woman in the row in front of me put her mobile between her face and the scarf. Travel really broadens the mind.

Posted by Joerg Moellenkamp in General at 13:31

### **links for 2008-07-19**

YouTube - George Lucas in Love  
(tags: starwars georgelucas)

Solaris iSCSI CHAP and RADIUS Configuration - FAQ - A Honeycomb Belongs in a Garden  
(tags: iscsi opensolaris solaris)

Posted by del.icio.us in del.icio.us at 13:31

Friday, July 18. 2008

## **8 days to NYC**

At this time in 8 days i'm in New York. I'm really looking forward to a week of vaction. This will be a real vaction. I won't take my notebook with me (at least i hope i will not take with me ) My business telephone will stay in Hamburg (that's easy). I'm eager at reading my travel guide at the moment (besides writing an architecture paper for a death march project). I'm sure this will be 7 cool days in NYC.

PS: Two non-touristic things are on my list at least ... going to the Apple Store and watching Wall-E ... it's just to long for me to wait until it's premiere in germany

Posted by Joerg Moellenkamp in Travel at 20:50

## **Recommended listening: Instrumental version of "Signal to noise"**

"Signal to noise" is one of the most intense songs of Peter Gabriel. Whoever attended a concert of Mr. Gabriel knows the goosebumps effect of this song. Normally it's a piece with vocals from Peter and Nusrat Fateh Ali Khan. When i heard it the first time while driving from Oldenburg to Hamburg after midnight i had to stop. Yesterday i found the instrumental version of this song by a fluke when i have looked for the Wall-E soundtrack. Even without vocals, this song works exceptionally well.

Posted by Joerg Moellenkamp in Music at 18:29

## **links for 2008-07-18**

Japan's AIST boasts of longer-life NAND flash memory - Engadget  
(tags: NAND flash)

Posted by del.icio.us in del.icio.us at 13:30

## **Copycat Container**

Over a year ago, we have announced the Sun Modular Datacenter 20. Some people asked about our mental status because of this announcement. But it doesn't seem so dumb at all, as HP has announced a datacenter container on it's own: HP Performance-Optimized Datacenter

Posted by Joerg Moellenkamp in General at 10:49

Thursday, July 17. 2008

### **Hiiiiilfe ... die Dieselwolke ist da ....**

Da stehe ich ganz unschuldig zum Teewasserautomaten degenerierten Kaffeeautomaten, und rieche Diesel .... Udine laesst gruessen.

Posted by Joerg Moellenkamp in General at 16:04

### **links for 2008-07-17**

Spreeblick - Warum man niemals, niemals, niemals seinem Online-Uebersetzungstool trauen sollte, Teil 372  
Never trust an automatic translation tool ...  
(tags: humor)

Add iSCSI iBFT Support (v0.4.9) [LWN.net]  
(tags: iBFT)

Samsung's low-power 128GB SSDs go mass production on the cheap - Engadget  
(tags: hardware memory ssd)

SanDisk introduces write-once WORM SD cards - Engadget  
(tags: flash storage)

Blocks and Files: Sun and Samsung's extended-life flash  
(tags: ssd sun samsung flash)

2008/427 - iSCSI Boot at OpenSolaris.org  
iBFT based iSCSI boot  
(tags: iSCSI boot Solaris)

Posted by del.icio.us in General at 13:30

### **A ray of light in an otherwise totally lunatic and ridiculous day**

Over 800 subscribers ... kewl ....

Posted by Joerg Moellenkamp in Braindump at 11:06

### **400000 terrorists**

I always thought, that terrorists are small groups of determined lunatics. But i wan't aware of the scale of this lunacy. The US administration thinks, there are 400000 terrorists in the world, as Bruce Schneier reports in a blog entry. I have the impression that the US admistrations sees terrorists everywhere.

Posted by Joerg Moellenkamp in General at 07:16

### **The upcoming Opensolaris 2008.11 for the storage admin**

Elektronkind wrote a nice summary about all the features interesting for a storage admin in OpenSolaris 2008.11 - A Preview For The Storage Admin. He highlights the most important features of the Solaris Operating Environment in regard of storage.

Posted by Joerg Moellenkamp in Solaris at 07:08

Wednesday, July 16. 2008

### **Solaris Crashdump Analysis Tool 5.0**

Dana announced the availability of the Solaris Crashdump Analysis Tool in it's 5.0 incarnation today. The CAT is a really cool tool dig down into the crash dump of a Solaris system. It was written by people who do dump analysis for their living. Really worth a try!

Posted by Joerg Moellenkamp in Solaris at 19:15

### **"The website is down!"**

Posted by Joerg Moellenkamp in Fundsache at 18:39

### **Moellenkampsches Gluecktheorem in Aktion...**

Is ja wie immer: Ich freue mich ja schon echt viereckig auf die Woche in New York ab dem 26. Juli . Aber muss die Lufthansa-Belegschaft gerade sich ueberlegen in einen Streik zu treten zu wollen? Geht das nicht auch ein oder zwei Wochen spaeter ?

Posted by Joerg Moellenkamp in Persoenliches Pech at 18:12

### **links for 2008-07-16**

Umweltforschung: 140 Liter Wasser fuer eine Tasse Kaffee - Nachrichten Wissenschaft - WELT ONLINE  
(tags: umweltschutz wasser kaffee)

Sun's ZFS/Flash initiative | Paul Murphy | ZDNet.com  
(tags: zfs flash ssd)

Posted by del.icio.us in del.icio.us at 13:39

### **More 787 delays imminent?**

The last three delays for the 787 were related to supply chain problems, now Boeing told the press that there are certification problems with the brake monitoring software as RedOrbit reports in: Boeing 787 Dreamliner Experiences Another Delay:Shanahan, who has led the 787 program since October last year, said the latest hitch has been delays in getting the software in the 787's brake control system verified to meet stringent certification requirements. [...] "It's not that the brakes don't work, it's the traceability of the software," Shanahan said[...] He said that Crane had to go back and rewrite certain parts of the brake control software to verify it for the certification process. My sources in the airline industry told me that the certification is the hardest part at building a plane and thus problem can haunt you for a while when you are not really cautious. In addition to that the damages part of the Dreamliner Four introduce further timing problems, as they have to rebuild that section completely.

Well, i hope that Boeing is able to get rid of the problem. Another delay doesn't really sound as beneficial for the company. They already sold near of a thousand airplanes and only a million dollars of additional contract fines per plane would cost them \$1bn now.

Posted by Joerg Moellenkamp in Aviation at 09:06

### **Preliminary results for Q4FY2008**

Looks like our quarter wasn't that bad as many analysts and investors expected. Reuters reports in Sun Micro shares up on preliminary earnings: Sun, the world's fourth largest business computer maker, said it expects fourth quarter net income of 5 cents to 15 cents per share, including a \$100 million restructuring charge.[...] Sun on Tuesday also said it expects to post fourth quarter revenues ranging from \$3.73 billion to \$3.8 billion, as compared with \$3.835 billion for the fourth quarter of fiscal 2007. [...] Excluding items, Sun expects to report earnings per diluted share ranging from 25 cents to 35 cents. In after-hour trading, the SUNW stock rose to \$9.88 (or 12.12%).

Posted by Joerg Moellenkamp in Sun at 07:21

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

Tuesday, July 15. 2008

### Good idea/Bad idea - today: IM and presentations

Good idea: Using instant messaging for communication at Sun.

Bad Idea: Not to quit the instant messaging client before you plug your laptop to the beamer.

Posted by Joerg Moellenkamp in Persoenliches Pech at 21:19

### Once-in-a-lifetime disaster

There are disasters events in your datacenter of the kind, where will tell stories about in 10 years at the coffee machine: Anfalas working at the datacenter of the University of Erlangen had such an diaster in his datacenter. The control rack of the USV burned down. Power outage. 800 servers down in the morning ... all systems online again at the end of the same working day. He reports about this event in his blog: USV Brand am RRZE (german). I'm impressed.

PS: Rolf, das waere doch etwas für ein "Systemhelden des Monats" auf systemhelden.de

Posted by Joerg Moellenkamp in General at 16:10

### links for 2008-07-15

How To Set Up Software RAID1 On A Running System (Incl. GRUB Configuration) (Debian Etch) | HowtoForge - Linux Howtos and Tutorials

(tags: admin boot debian documentation etch howto grub linux sysadmin raid)

Eee box gets price and package details - Engadget

(tags: eeebox eeepc)

MAKE: Blog: Help Mythbusters recreate Archimedes' death ray

Sh..... i need a reason to stay hat SMI Corp HQ in September

(tags: mythbusters)

Posted by del.icio.us in del.icio.us at 13:40

### The fourth anniversary of c0t0d0s0.org

Well, c0t0d0s0.org celebrates it's fourth birthday today. On the 15th of July 2004 this blog started. Four years is a long time. I didn't really thought that i would write here that long but it's still fun at the moment. I really think that this blog will at least celebrate it's fifth anniversary.

What has happened in the last 12 month? The "Less known Solaris Features" series was really successful. The pdf of this series was downloaded almost 6000 times. 40.000 pageviews on the various LKSF articles in the blog. The series is vastly more successful than anticipated by myself when i thought about writing tutorials.

The spin-off aviation.c0t0d0s0.org doesn't took off as expected, but was my fault as i didn't had the time to publish new articles. My reader seems to be interested about news in that topic, as the statistics indicate interest in the aviation related articles.

The November was a great month for me. I've won the Best Blog of CEC contest with my reports from the CEC 2007 in October. I didn't really thought that my blog was good enough to make the first prize but the jury decided different. A big

thank you to you, again.

BTW: I hope that i can attend the CEC this year again (Hope to sit on the right side of the plane at the approach this time) but i think i won't write that much to concentrate more on networking there. After 7 years at Sun i've recognized that i should focus more on that as the next possible steps in my career at Sun are dependent from that.

The big picture: It is somewhat strange to blog about big technological steps into the future and seeing our stock price declining. But as someone told me: The stock market isn't the right location to look up the value of a company as it mutated into a large playground for private and institutional gamblers. Well ... sometimes it's disappointing not to be an official leak for informations. Besides one or the other scoop this would simplify the task to explain some movements vastly. All i can tell to you is that Sun will announce and deliver some really astonishing products until end of the next year (and yeah, Rock is one of them and it isn't dead).

Some numbers for the last year: 760 subscribers, up to 6000-7000 vists on weekdays, 2000 on weekend days. 4527 articles in the blog consisting out of 3.634.327 characters. The blog starts with 3306 comments in it's fifth year. Thank you to all the reader that participated in my blog in the last year.

The plans for the next year: At first the numerical ones .. over thousand subscribers. 10.000 visits a weekday. 5500 articles. I know, high numbers, but i need high targets ...

There will be still no advertisements in my blog (besides the small one for the company of my brother) as parts of it are created at work. May be this is right time to thank my managers for tolerating and encouraging this diversion from the daily business. There will be no printed LKSF book in english, but chances are good that there will be one in german.

Furthermore there will be another tutorial series about features in the Sun world that doesn't fit into the LKSF moniker. For example an Grid Engine tutorial or an Sun Cluster tutorial ... but both are in the earliest planing stages.

I wasn't at CeBIT this year, so i hadn't the opportunity to meet some of my readers, who asked me if they could find me on the CeBIT booth

At the end, i want to tell you two stories about the importance of my blog for me.

Silke told me, that is was easy to find out, that i was relatively fine after the breakdown of the last year. You just have to look if i write in my blog. So people who care seem to look in my blog as it's my public life sign and thus it was one of the first habits i had resumed. Some people say, you have meet to befriend someone. But that isn't true. I've won a really good friend in the last year by writing this blog ... it's Local from localwurst.de (and i still don't like the name of her blog )

Posted by Joerg Moellenkamp in General at 00:00

Monday, July 14. 2008

## **Solaris Infiniband Update 2**

An important update for the Infiniband Stack in Solaris was released yesterday. Josh Simons reports in his blog: Yesterday, the Sun InfiniBand engineering team released Solaris 10 driver support for ConnectX (a.k.a. Hermon), the latest generation of InfiniBand silicon from Mellanox.[...]In addition to the driver, the update also includes a new flash updating tool for ConnectX, a uDAPL update, and several additional components, all of which is described in the documentation. You can get Solaris Infiniband Update 2 at the usual location ...

Posted by Joerg Moellenkamp in Solaris at 22:03

## **Nachteil eines etwas selteneren Nachnamens ....**

Ist ja echt schade, das mein Nachname etwas seltener ist ... sonst wuerde mir die Air Berlin vielleicht auch die Bordkarten für interessante Ziele anderer Leute anbieten .... Immer nur nach München mit AB zu fliegen ist auf die Dauer doch langweilig

Posted by Joerg Moellenkamp in Aviation at 21:34

## **Road photo**

Posted by Joerg Moellenkamp in Photographie at 21:23

## **Some benchmark numbers about the SPARC Enterprise M9000**

There are some interesting benchmark numbers in the press release from Sun about the new SPARC Enterprise M9000 with 64 SPARC VII ("Jupiter") processors:

The SPARC Enterprise M9000 server running Solaris 10 OS, the SAP ERP application Release 6.0 and Oracle Database 10g achieved a world record result on the SAP Sales and Distribution (SD) Standard Application Benchmark by achieving 39,100 SD Benchmark users. The SPARC Enterprise M9000 server result beats the IBM Power 595 POWER6 server with 35,400 SD Benchmark users by 10.5% and outperforms the HP Integrity Superdome SD64B with 30,000 SD Benchmark users by 30.3%.[...mabushi deleted...].[1]In addition to that, Sun announces some impressive numbers for the LINPACK benchmark. 2 Teraflops in a single operating system imageThe SPARC Enterprise M9000 server, based on SPARC64 VII processors, and using Sun Studio 12 software with Solaris 10 OS delivers a score of 2.023 TFLOPS on the Linpack's Highly Parallel Computing benchmark. The SPARC Enterprise M9000 outperforms the IBM Power 595 POWER6 by nearly 2X and beats the HP Integrity Superdome system by 2.7X.[2]This is interesting for computing jobs in need of a large single system image with short ways to to the memory. (No, the SGI Altix isn't a competition for that ... single image too, but completely different architecture). Not all jobs scales on a cluster machine

Posted by Joerg Moellenkamp in Sun at 16:09

## **1 TB tape drive**

Today we have announced the 1 TB tape drive: Sun Microsystems Announces World's First One Terabyte Tape Storage Drive. The the T10000B delivers 120 MB/sec throughput with 1 TB native. The next competitor (LTO4) has 800 GB native. The only disadvantage: This tape drives are enterprise hardware and build to last ... so the pricetag is \$37.000. So don't wait for it as you backup solution at home.

Posted by Joerg Moellenkamp in Sun at 15:59

## **First news about the SPARC64 VII announcement**

The first news about the upcoming SPARC64 VII appear at the usual suspects: e.g. cnet. There is another report at TheRegister but Ashlee Vance seems to be a little bit bitchy, as Sun canceled his call

Posted by Joerg Moellenkamp in Sun at 09:28

### **Sun Fire X4150 reviewed by Infoworld**

Infoworld did a nice review of the X4150 in Sun Server Delivers Heavyweight Performance: Even without its disks, the X4150 is a powerhouse of performance in a small profile.

Posted by Joerg Moellenkamp in Sun at 09:10

Sunday, July 13. 2008

### **Mailserver migrated**

Because of a major hw fsckup on the old mailserver and some preparations for another project (yes, Kristan, relief for your problem is near) , i've migrated my mail server to a new system. I hope my new mail anti spam configuration isn't too strict now ....

Posted by Joerg Moellenkamp in General at 23:18

Saturday, July 12. 2008

## **links for 2008-07-12**

Elliptic Curve Cryptography (ECC) performance - Lawrence Spracklen's Blog  
(tags: cryptography ECC CMT)

Posted by del.icio.us in del.icio.us at 13:36

## **The first photo with my new digicam**

As i don't walk through New York with all the bodies and lenses i've opted for a new digicam. It's a Fujifilm Finepix FD100fd and it has the reputation of being one of the best (or the best) compact camera in regard of the image quality. Okay, this is the first photo:

Posted by Joerg Moellenkamp in Photographie at 11:27

## **Software crypto on CMT**

Lawrence Spracklen wrote a really interesting article about cryptography done in software on CMT. This sounds counterintuitive at first, as cryptography is considered as a computational intensive task and thus considered as a tasks for fast superscalar cores. But according to the article from Lawrence this is a implementation issue. Take the strength of the CMT architecture, and the result is a little bit different: As a result, as the number of strands is increased, performance scales almost linearly. Indeed, for Niagara, per-core Kasumi performance is around 8 times the performance of a single strand, and per-chip Kasumi performance is close to 64X single-strand performance. Indeed, single-core Kasumi performance is around 1.3X the performance of a single-core of a 3GHz Xeon processor.

Posted by Joerg Moellenkamp in Sun at 10:01

Friday, July 11. 2008

## **FSC's future in question?**

(Disclaimer: This article is my personal opinion, i'm not aware of the official opinion of Sun Microsystems regarding this rumours)

Well, this is a really interesting development: According to the german news portal heise.de with their article "Bericht: Siemens kündigt Vertrag mit Fujitsu"(on german) and a report on Reuters Siemens wants to get out of the joint venture (JV) with Fujitsu as the Siemens CEO doesn't see FSC as a business unit that's profitable enough to stay in the company.

This is interesting for Sun out of a simple reason: FSC is not just that notebook company. FSC is a distributor of the SPARC Enterprise Systems (M-Class and T-Class) and those systems are identical to ours. In Germany FSC has quite a market share.

Okay, according to the reports there is an agreement in the JV contract: Whenever someone wants to get out of the JV agreement, the leaving partner has to offer it's stake to the other partner at first. But i don't think, that Fujitsu will buy the part ... especially not when there are signs of a adverse global economic downturn. Thus it looks reasonable that they will offer the part to someone else will buy it interested in the well known Fujitsu-Siemens brand for desktops and notebooks. It would be further reasonable to assume that a new owner would concentrate on the x86 business and thus won't take the SPARC business. And this would open an opportunity for Sun.

Well, i will go to bed now and think a little more about this development. But as usual: Interesting times ahead of us ...

PS: Maybe this was thee reason for announcing the Xeon based mainframe a few days ago ... this would give Siemens a opportunity to keep their mainframe business without the SPARC procs from Fujitsu.

Posted by Joerg Moellenkamp in The IT Business at 22:26

## **News from the tanker deal**

Well, Northop lost the deal after the GAO ruling, there will be an new Request for Proposals ... but i really think that Boeing won't get the deal in the next attempt either when there is no further political counterfire.

The problem is somewhat bigger: Let's assume Boeing has enough time to build a tanker on the foundation of the 777 airframe to get a bigger aircraft. They have another problem: They do not have the capacities to build it. The 777 is a successful plane and in addition to that they sell this planes as an interims solution for airlines with an urgent need for aircrafts since the 787 is vastly delayed at the moment. The USAF would be a customer like anyone else and do not believe Boeing can explain to their commercial customers that the next 150 airframes are reserved for the USAF ... which in turn would be nice for Airbus: More A330/A340 to sell. And they can't put the tanker 777 at the end of the line. At several locations in the net you can read that they need a replacement for their old tankers really soon.

The solution with the 767 wouldn't have this problem. It's an EOled planed for commercial customers and thus they have the capacities to produce it in the old assembly lines. So the point "the airforce wants a bigger tanker" is a real problem to Boeing. They are dependent to the old requirements in the request for an proposal. And in some news outlets in the US you can already see the lobbying against a modification of the requirements in favour of an bigger plane.

I think, the end of this year will really interesting for both companies.

Posted by Joerg Moellenkamp in Aviation at 18:55

## **72 TB X4540?**

No, this isn't an announcement about a new flavour of our storver, but i started to think of new capacity jumps at hard disk with the Thumper/Thor factor in mind: Segate announced an 1.5 TB harddisk ... such drives would give the X4540/J4500 a capacity of 72 TB. One and a half rack for one Petabyte ... and in the case you use the 32 MB Cache flavour of the drive you would one and a half gigabyte worth of cache on the disks alone. Wow ...

Posted by Joerg Moellenkamp in The IT Business at 11:38

### **The reports of my layoff are greatly exaggerated**

eWeek linked my blog in an article about the layoffs as being one of the employees laid off yesterday: And a series of Sun's more prolific bloggers weighed in on July 10 to say it was their last day at Sun. In particular, this one and this one are telling. As long as the eWeek don't have informations i'm not aware of, this isn't the truth. As i'm not important enough I really doubt that a news outlet in the US has any internal HR informations about a SE in Hamburg

So: I'm still working at Sun, i'm not aware of any actions in this direction. I only grieved about the farewell messages i've got in my inbox and the farewell blog entries i've got in my feedreader from people i've heard in telephone conferences or met at some occasions.

Posted by Joerg Moellenkamp in Sun at 09:43

Thursday, July 10. 2008

### **I hate days like this one ...**

It's not the first reduction-in-force at Sun since i started to work for Sun ... but i still hate those days with farewell mails and blog entries more and more. This is the day when an announcement has voices and faces and you are sorry for the people behind the mails leaving sun.

Update: When reached this article by the link on eWeek, please read this article as well - The reports of my layoff are greatly exaggerated

Update2: The author of the article corrected the error. Thank You!

Posted by Joerg Moellenkamp in Sun at 22:03

### **Bettina und die Haie ...**

Ich kenn jetzt Bettina schon eine kleine Ewigkeit. Deswegen musste ich auch erstmal überlegen. Hat sie wirklich nichts mit dem "Kleinen Hai" zu tun? Ist das nicht von ihrem Arbeitgeber eingefädelt . Kleiner Hai? Nun ... das ist wohl der legitime Nachfolger von Schnappi und ein neuer Höhepunkt des Fremdschämens.

Was hat nun Bettina mit dem "kleinen Hai" zu tun? Scheinbar nur wenig, Bettina arbeitet bei EKM-Consult, EKM-Consult hat ein Blog und weil es so gnadenlos bescheuert ist, wurde das dann auch dort verlinkt. Über einige Umwege und Verschwörungstheorien wurde dann vermutet, dass die Verlinkung nur die Spitze eines Eisbergs einer durch EKM geplanten Marketingaktion viraler Art war. Das ging jetzt soweit, dass man sich bei EKM genötigt sah, eine Pressemitteilung zur Gegendarstellung zu schreiben.

Nur: Zutrauen würde ich Bettina sowas ... ohne zu Zögern ... aber sie schwört Stein und Bein nichts damit zu tun zu haben ....

Wer war es nun wirklich? Wer steckt hinter diesem Anschlag auf die menschliche Intelligenz? Wann wird es die erste CD mit dem Hai geben? Tshirts? Klingeltoene? Wird sich Alemuel aka Alexandra Müller je wieder auf die Strasse trauen koennen?

PS: Ich habe ja eher die kleine Firma DITA-Consulting in Verdacht, aber das ist nur so eine Vermutung

Posted by Joerg Moellenkamp in General at 14:53

### **links for 2008-07-10**

TP: Wie man die Prostitution auch foerdern kann  
(tags: studium politics)

Posted by del.icio.us in del.icio.us at 13:40

Wednesday, July 9, 2008

### **Get rid of root - or: PSARC/2008/321**

An interesting PSARC case got it's approval: PSARC/2008/321 - No Root Login:  
It has been suggested that Solaris should permit "root" to be a no login account.

[...]

Add a "solaris.system.maintenance" authorization. Modify slogin(1M) toprompt for a username and password. If the username entered is authenticated by the password and has the "solaris.system.maintenance" authorization, enter system maintenance mode. If not, as before this project, deny access.

Besides of making root a role, this would be another way to get rid of a direct usable root account in Solaris. With this change you can use your own username and password to log into an "root shell" without knowing the root password ... there doesn't even have to be a root password. You only need the authorization solaris.system.maintenance in your user attributes.

(via ceri on twitter)

Posted by Joerg Moellenkamp in Solaris at 22:06

### **New storage from Sun: J4200/4400/4500 and X4540 Storage Server**

Sun announced some new storage components today: At first we announced the J4200 and J4400 disk arrays. They are JBOD's, just bunches of disks. The J4200 has up to 12 SAS or SATA disks in a 2 RU enclosure.

The new J4400 is a 4 RU enclosure with 24 SAS or SATA disks. It's a doubled up version of the J4200:

Well ... and then there is the J4500. As the name suggests, it's a JBOD that looks really similar to the X4500 - the Thumper - and it has similar characteristics: 48 harddisks on 4 rack units.

When you look at the rear side of the J4500 you will not see the Opteron based System controller. Instead of this you see an interface board with four SAS-connectors. Two of them are host connect ports. The other two are used to stack J4500 enclosures together.

And last but not least there is the X4540. It's our new storage server. It looks almost as the old one. The changes are internal ones. The system now uses quad-core chips instead of dual-cores, the system gives you 3 PCI Express slots instead of the PCI-X slots of the old system. The new X4540 has a slot for a compact-flash card to boot the operating system. Or to say it more shortly: We integrated the stuff our customers wanted to see in the next-gen thumper.

Posted by Joerg Moellenkamp in Sun at 17:48

### **links for 2008-07-09**

Intel's Larrabee GPU based on secret Pentagon tech, sorta  
(tags: Intel CPU GPU Larabee)

Posted by del.icio.us in del.icio.us at 13:36

### **Vertical Multiteleconferencing**

While writing the article about the UltraSPARC T1 on the train ride back to Hamburg yesterday afternoon, i had an idea for a solution to save large amount of time a day. I think, anybody knows the problem: You are in a conf call, anything was said, but not by everybody ... or someone uses his slot in the "Misc." session at the end in an abusive way to talk about the books he or she had read in the past or what people he or she had met. Or news from the last vacation. You know, that timewaisting stuff ...

But the solution is sooo simple: The teleconf selector and multiplexor. In every conference there is a person detecting a stalled teleconference. This state is reached, when someone talks but delivers no important information. Now you can simple switch to the unstalled telephone conferences and hear it. I think four or five telephone conferences in parallel is a good compromise ... there is no danger, that there will be good informations in two of four conf calls at the same time. And even when there is this situation put one of the good informations in the GIB (good information buffer) for a delayed replay when there is no other good information..

This would really save much time.

Posted by Joerg Moellenkamp in Braindump at 12:44

Tuesday, July 8, 2008

## **About some rumours surrounding the UltraSPARC T1/T2**

Every once in a while a blogger, sales rep of a competitor or a misinformed journalist write "The UltraSPARC T1/T2 cores are just UltraSPARC II cores running at 300 MHz". Well, this is a rumour consisting out of two misunderstood points of the architecture. Well informed people already know the stuff in this article but this time your are not the target group of readers. But i think it's time to write a little bit about all this rumours as i write or tell this answer again and again.

Okay ... the UltraSPARC T1/T2 cores are just UltraSPARC II cores ...": This part of the rumour is as old as the UltraSPARC T1 series. But it's really simple to see, that it isn't true. Okay, a SPARC CPU is a SPARC CPU is a SPARC CPU ... you will find similar functional units on all SPARC CPUs.

Well, to be honest ... there are some features of the UltraSPARC II that aren't available in in the UltraSPARC T1 processor: For example, the UltraSPARC II an up are superscalar designs, where as UltraSPARC T1/T2 are scalar designs. The UltraSPARC T1 has short 6 stage pipeline, the UltraSPARC T2 is a little bit longer with 8 stages , UltraSPARC II has 9 stages and the UltraSPARC III has a 14 stages deep pipline. The T1 is a single-issue CPU (it issues one command at a time at pipeline), where as the II/III are multiple-issues CPUs. The cores are pretty much different. So wherever you read this "The N1 core is just an UltraSPARC 2 core", this is simply incorrect.

Okay, you might ask, why didn't we simply used the UltraSPARC III cores and glued them together: UltraSPARC T1 and II/III/IV/IV+ were designed with different mindsets leading to different designs. The core in the UltraSPARC T1 CPU was developed from the ground up to reach two targets: To deliver a SPARCv9 compatible core in the least possible amount of space with the least possible power consumption (okay, in a given timeframe and budget). Processor design isn't a single-direction road ... it's more like a tradeoff game.

The first installment of the concept drove this concept even that far, that they outsourced the FPU commands to a single FPU. FPUs aren't simple, they take a lot of space and they use a lot of power. We want a fscking fast database and web server. We don't need fscking FPU commands. Tradeoff game starts. FPU ... there is the door ... please close it from the other side. We've learned that some application used FPU where nobody really expected it and now N2 has eight of them.

You want to reach a certain target, look at your given budgets (transistor budget, die size budget, time budget, people budget, money budget) and search for solution fitting in your requirements without loosing too much at another front. IBM wanted the 4,7 GHz CPU and sacrificed the Out-Of-Order-Execution on that way, Intel wanted the GHz crown as well and designed the ultra-long pipeline Pentium (anybody remembering Prescott?). VIA wanted a lowest-power x86 CPU and sacrificed performance for a low power consumption. Sun wanted a many core CPU and sacrificed the core complexity in the first steps. You can't have it all, especially with all this walls around us. The thermal wall, the GHz wall, the budget walls, the structure size wall.

BTW: At the the end the laws of nature will stop us all and i think we will see a development like with the race to reach the temperature of zero Kelvin. Every step to half the temperature needs the same energy (i hope the recollection of my school physics didn't left me there). Every step to get nearer to the walls imposed by the laws of nature for chip manufacturing will take the same amount of money. I think it's a save bet, that AMD or Intel won't build a electron accelerator in the size of the Large Hadron Collider just do produce a particle stream for the Higgs-Boson lithography to build Core10Quad. I think it's more probable that we will see a Core4Hekaton based on a stacked structure of multiple dies but with a structure size in the range of reasonable economic investments.

We getting to the point where nobody can afford to build a fab for smaller structure sizes only for one or two generations of processors. The timeline of Intel procs is really interesting at this part. Each process technology was in use for two generations. 90nm, 65nm and 45nm (at least when you look at the official roadmap). This is a ruinous game at the long run. And this is the reason why the entire industry put it's research dollars, euros or yen into manycores.

I strongly believe that all proc vendors will opt for manycore designs. The question is only how to handle the legacy single-thread code. You may end up with a few (4 or 8 ) heavyweight high-speed cores like Intel or AMD, 16 heavyweight high speed cores with scout cores like Suns Rock. So everybody has the same problem. I'm aware of the developments at Sun to solve this problem and they look really promising, i'm aware of other initiatives at other vendors to solve the same issue for their respective technologies.

Okay, i've lost the topic ... back to the deliberate tradeoffs:

For example: The UltraSPARC T1 has a 6 stage pipeline. For modern procs this is really small pipeline. But it was a deliberate design decision. We could easily use a longer pipeline to reach higher clocking, but here we get to the trade of game again: You may reach higher clock frequencies, but the longer the pipeline is, the higher the penalty for a thread switch gets. You have to empty the pipeline and reload the pipeline and very step in the pipeline takes a cycle.

Let's assume that you have a 15 stage pipeline. Let's assume that the current thread stalls in stage 10 as it waits for data in the memory. You can react on this problem in two ways: Wait for the thread to continue or switch to another thread waiting for execution. When you switch to another thread all commands in your pipeline are not longer valid. You have to unload it and start work through the the pipeline again. The penalty for the thread switch is 9 cycles, as you need nine cycles to bring the first command of the new thread to the pipeline stage before the stage, that had stalled before. Okay, and here begins the gamble: Do you expect that the thread will unstall within 9 cycles or do you think it will take longer. 9 cycles are in the pot. Hmmm, hard decision.

Now let's calculate this with a 6 stage pipeline. Let's assume, the thread stalls in stage 3. Your switch penalty is just 2 cycles. It's a safe bet, that you won't get your data within 2 cycles thus.

But there is a further difference between normal procs and the Niagara. It has an own register set for every thread. This has important reason two. When you want to switch from one thread to another thread, you have to save the register sets, load a new one and when you switch back you have reload the old one. This isn't efficient. Such a task can take easily several hundred clock cycles.

Given the fact that Niagara has four independent register windows to switch between threads at no costs, you can use a another trick to fill that bubble in. Try to issue command from every core in a round robin fashion. Even when the pipeline stage stalls, the next command in the thread is perfectly executable, as it's from another nonstalled thread. As long the thread is stalled no new commands will be issued to the decoder stage. This is meant by thread switching without any costs.

This design choice have an different, but really interesting implication as well. Let's assume we have a four core four proc x86 system and a 64 thread Niagara II system. Let's further assume you run an application with 64 threads. This leads to an interesting effect: Context switches are incredible expensive operations. A context switch takes place when you want to run a different process on your CPU. You remember the the cycle of storing thread A registers, loading thread B register sets, storing thread B registers, loading A register set. When you have 64 processes but just 16 cores and thus 16 register sets, you have to go through the cycle four times to execute all processes. When you have 64 processes and 64 threads and thus 64 register sets ... how many context switches do you have to do? Correct ... none. There are several benchmarks that suggests that a UltraSPARC T2 will keep the same performance over a larger thread count, whereas other architecture may start faster, but will break-in after the thread count is larger than the core count.

And this led to other design decision: Do you really need Out-of-order execution logic (complex, thus power hungry), when you can simply switch to another thread in case of a thread stall? The same for branch predictors. Is it necessary to have big caches, when you augment the thread switching with four memory controllers and a fast crossbar? You can't use your x86 knowledge and use it for SPARC. You can't use your SPARC knowledge for x86?

And there we get to the second part of the rumour. It's "... running at 300 MHz": I know several presentations of our beloved competitors stating that each hardware thread has only 300 MHz worth of computing power. I knew that this bullet would be fired by our competitors at the moment i saw the first presentation about this chip. They used the following logic: The proc (one version of it) runs at 1.2 GHz A core runs 4 threads in parallel. 1200 MHz divided by four is 300 MHz. I assume that is the source of the UltraSPARC II rumour. The USII was clocked in that range. Well ... this logic is correct ... to a very limited point. When all cores work at full speed and all data fits into the L1 caches this is correct. For example executing the incrementation of a register ad infinitum. The problem: This isn't a real workload. It isn't even near of a real work load.

Real workloads are: Loading data from RAM or disk, work with it, save it to RAM or disk. And now a normal CPU does something quite usual for modern CPU ... it waits. Memory is slower than the CPU. Period. And when you clock at 3 Ghz even the first level cache miss and second level cache hit is a high latency event. This is even the case at some extend for clocks at 1 GHz.

As i told you before the Tx series has some logic to skip stalled threads and it issues commands to the decoder round robin. The next cycle is one with effective workand not spend with doing nothing. It's all about handling the one truth: You have to wait for memory.

So: The calculation clock frequency divided by four is only theoretical because the real world teaches us that cores doing real work will wait most of their time. You can measure almost every application and will find large amounts of cache misses. Did i wrote "You have to wait for memory" before? Well ... nevertheless i write it again: You have to wait for memory most of the time. And it's an absolute amount of time you have to wait. Given the same kind of memory, it's not relevant if you core clocks with 1,4 GHz or 3,6 Ghz. The time until the data from reaches you, will stay the same.

You can handle this problem on multiple ways: You can do thread multiplexing like Sun did it. You can but giant caches on your die like Intel or IBM (An Power6 CPU has more cache on die than my first 486 PC had as main memory) You can do out-of-order execution, branch prediction, clever preloading of the caches, scouting. At the end you do all this to reduce the latencies introduced by the speed differences of memory and CPU. And this latencies costs you the most of your performance.

My point is: When you divide the clock frequency of the Niagara by four because of the thread multiplexing, you should reduce the clock frequency of x86 or other high clocked processors by all the ticks wasted to latency. And how do you factor in the cryptographic units of the T1/T2? They work parallel to the cores. The N2 cryptographic units writes back it's result directly to the cache and don't use the pipeline ... interesting question and many of the "300 Mhz" people don't even think about it.

I have to admit: Niagara isn't a system you can't use without thinking. At the end: It's a 1.4 GHz proc with a simplified SPARC core and the biggest advantage of the processor is the achilles heel of the system as well. When you measure single thread performance you certainly measure the wrong thing. When you only load it with a single thread, it's the wrong processor. When the execution time of a single thread on a system loaded with a single thread is important to you it's the wrong processor. I have a simple rule, when i support a sales rep: When i don't know the application, i do not recommend the UltraSPARC Tx series. Period. The problem: The most admins don't really now their application that good ... but the Coolthread Selection Toolkit solved this problem really nicely. Big kudos to the development team of this tool.

I assume that the people with a negative position to the UltraSPARC T1/T2 chips are exactly the people who tried to use an application on it with the wrong characteristics and were disappointed.

But: When you have several threads in parallel, when it's important for you that the performance stays at the same level even when 256 threads of Apache do heavylifting data, then UltraSPARC T's are the correct choice for you. When you want good performance at a high thread count, it's a good choice. No, it's the best choice.

You can compare it with a E10k. A UltraSPARC T2 is a E10k on a chip. I wrote in a wiki article some years ago: The T-class systems are huge SMP/near-SMP designs, and they want to be used as such. Don't let them confuse you by their size. Batoka will give you 256 threads on four rack units. Not long ago you could easily fill a mid sized datacenter with the machines needed for the same amount of cores.

There we get to another problem of modern IT: The availability of notebooks with processors with multitudes of GHz on the developers desktop led to a large heap of bad software. It's like with the F-4 Phantom. The F-4 phantom was the proof for the concept, that you can fly with a brick, when you put enough thrust at it. And the General Electric J79 for the modern software developer is the Core2Duo at excess of 3 Ghz. The software developer doesn't get punished for making his/her life easier by using the design pattern of the singleton instead of thinking about proper code scaling on multiple processors. On their notebook with a test load this brick will fly. I spoke to admins full of bitter words about their colleagues after they installed an app server per core to get some performance out of their machines. The "good" thing ... all vendors go the way of many cores and the need for good and scaling code isn't a problem of Sun ... it's an industry wide problem .. universal punishment for bad code is near. But: In my opinion the problem of bad code will haunt the industry for quite a while. All of us. It will cost all of many headaches and many R&D dollars to find workarounds for this problem.

Okay, the workload for the T class servers are specific ones: High thread count and many small requests but with large datasets, and the T class will run as hell ... This is the reason why we see such good results at customers for OLTP, for webserving, for mail, for enterprise backup any many other loads. This isn't a niche and it's the reason why we sell that many of the boxes. Sometimes this poses even a problem for us. Where we sold a big box in the past, we have sell just an T2 system. But better steal your own lunch than giving it undeliberately to someone else ...

I hope i gave you some insight into the differences of USII and UST1/T2 and what led to this design decision. As i'm not a chip guy, i hope i've got everything right but please correct me, if i'm wrong. I hope that you got at least an understanding, that the world in chip design isn't that easy to make simple and thus false comparisons.

Posted by Joerg Moellenkamp in Sun at 21:02

## **No Atomstrom for my Wohnhome**

Momentan versucht ja die CDU uns wieder die Atomkraft schmackhaft zu machen, obwohl es eigentlich seit Tschernobyl gesellschaftlicher Konsens ist, das die die Auswirkungen einer Katastrophe so massiv sind, das das Risiko einfach zu gross ist. Und über das Problem Endlagerung wurde schon vor meiner Geburt diskutiert und auch 35 Jahre spaeter gibt es da keine Lösung für.

Die Zeit hat hierzu uebrigens einen sehr guten Artikel geschrieben:Die »Sorge um das Sein«, die durch globale Risiken weltweit geweckt wird, hat längst zu einem Vabanquespiel um das Sein, zu einem Welthässlichkeits-Verdrängungswettbewerb der Großrisiken geführt. Die unkalkulierbaren Gefahren, die vom Klimawandel ausgehen, sollen mit den unkalkulierbaren Gefahren, die mit neuen Kernkraftwerken verbunden sind, »bekämpft« werden.Ich glaube ja ohnehin nicht, das es hier im Ökologie geht. Ich halte das ganze für Lobbyarbeit entsprechender Kreise, um die kurz vor der Abschaltung befindlichen AKW nach ihrer Abschreibung noch zwei bis drei Jahrzehnte äussert profitabel weiterbetreiben zu koennen. Es geht wie immer nur ums Geld.

Vielleicht sollte man sowieso in die entsprechenden Gesetze und Verordnungen schieben, das im Falle eines GAU das betreibende Unternehmen zur Deckung der Kosten sofort enteignet wird. Die wirtschaftlichen Risikobetrachtungen der Energiefirmen sähe sofort anders aus,da die Kosten der Katastrophe nicht mehr sozialisiert werden und wir muessten nicht ueber diesen Unsinn diskutieren.

Übrigens: In der Süddeutschen wird über die möglichen Endkundenersparnisse durch eine längere AKW-Laufzeit geschrieben: Es sind 50 Cent pro Monat ... der Austausch einer Glühlampe durch eine Energiesparlampe bringt mehr ...

Posted by Joerg Moellenkamp in policy of ... at 06:44

Monday, July 7. 2008

### **Der Fehler sitzt immer vor dem Monitor**

Da habe ich mich schon seit Monaten darueber gewundert, das mich einige Mails nicht erreichen ... hatte meinen Spamfilter in Verdacht ... den Viruscheck. Naja ... eben habe ich sie gefunden ... ich habe sie bekommen. Aber es ist schon ziemlich bloed, wenn saemtliche Mails als Absender eine Adresse tragen, deren Postfach man nicht abfragt .... ich geh mal meinen Postfix debuggen ...

Update: Klassischer Admintoobloed-Error .... kommt davon, wenn man glaubt, das man in einem WUI das gleiche machen kann wie mit einer direkt editierten virtual ...

Update 2: Ich habe da jetzt etwa 50 Mails ... mal sehen, welche ich davon noch beantworte ...

Posted by Joerg Moellenkamp in Persoenliches Pech at 23:00

### **747 crashes at start**

On May 25th a 747 of Kalitta Air broke in parts after a rejected take off after shooting over the end of the runway. It looks like Kalitta Air has a strain of bad luck at the moment: CNN reported that another 747 crashed into a house near Bogota: 747 freighter crashes in Colombia. The second crash of a Kalitta aircraft in 6 weeks. According to other news services it was the N714CK, a converted 747-200. This aircraft was delivered in 1981, thus it was 29 years old.

Posted by Joerg Moellenkamp in Aviation at 22:09

### **Hard Landing**

Mit besten Gruessen an meine Schwester Taz und meinen Bruder Florian, mit denen ich im August nach London fliege. Fuer beide ist es der erste Flug...\*evilgrin\*:

Posted by Joerg Moellenkamp in Travel at 21:55

Sunday, July 6. 2008

## **Nessun Dorma**

Ich finds ja auch gut, das die Werbung der Telekom Paul Potts entdeckt hat .... das ist jener Typ der bei Britain got talent eine gute Version von Nessun Dorma aus Puccinis Turandot (die Hamburger Inszenierung kann ich uebrigens sehr empfehlen) gesungen hat (hier als Video). Aber ich find das sch... das das jetzt alle paar minuten im Fernsehen kommt, da ich bei diesem Stück (am besten in der von Pavarotti, Domingo und Carreras gesungenen Version ... ist natuerlich eine ganz andere Liga, das ist nicht mal der selbe Sport) immer - aber auch wirklich immer - ein Gänsehaut bekomme ... und das nervt mit der Zeit

PS: Und erwiesenermassen ist es die beste Musik, um Menschen umzubringen, wie man seit der Endsequenz der Verfilmung von The Sum of all Fears bemerkt haben duerfte ....

Posted by Joerg Moellenkamp in Music at 21:14

## **History repeating**

This is the english translation of an article i wrote to answer an strange blog article in another german blog. While answering it, i found a case of "History Repeating".

In the actual installment of the Prozessorgeflüster (a reoccurring CPU technology article in the german computer magazine c't) Andreas Stiller discusses the the first facts about the new manycore processor "Larrabee" from Intel. Stiller states in this article, that this processor was discussed to release with with 16 to 24 cores, but will release with 32 cores when it appears in 2009. To the surprise of many experts, the cores are well known. The cores are nothing more than Pentium P54C cores. Well, the P54C was announced 1994. This would be similar to a Niagara T1 on the basis of the SuperSPARC II, which was announced 1994 as well.

Annother interesting fact reported in the article: The Larabee xPU will use roundabout 300 watts. That's much much more than a UltraSPARC T1 CPU. I know, Larabee is a GPU at start, but do you really believe that a manycore GPU with x86 compatible commands will stay on the graphics card for long time?

There is a small irony at this story. Sun learned many things about designing multicores when Sun developed the MAJC-5200 CPU and used it on the XVR1000 and XVR4000. The XVR4000 was the graphic card, that didn't used something earthly like a PCI-bus, you plugged it directly onto the Fireplane Interconnect of the V880 instead of the CPU. Sun learned so much about multicore that some of the findings will reappear only in future incarnations like the Rock CPU (speculative multithreading e.g.)

The funny (and "history repeating" part): Intel start s to use it's first manycore design on a graphics card as well. Like we did ... in 2002.

Posted by Joerg Moellenkamp in The IT Business at 17:11

## **CMT? Klar. Nur nicht auf dem Niveau von 1994!**

Der werte Herr Siering hat ja in seinem Blog angemerkt, das ich wieder mal Pseudoargumente fuer CMT finde. Ich habe mich dort in einem kurzen Kommentar ausgelassen, nichts ahnend welche Perle als Gegenargument mir Herr Stiller vom Prozessorgefluester einige Tage spaeter zuspielden sollte. Herr Siering ist ja der Meinung "CMT? Klar. Nur nicht auf dem Niveeau von 2001!" was in seiner Argumentation schon daher nicht treffend ist,, das das Gerücht die Cores der T1/T2 CPU seien nur UltraSPARC II Cores eben nur ein solches ist, ein Gerücht. Was zugegebenermassen von unseren verehrten Marktbegleitern oefters gestreut wird und in wenig Sun-nahen Kreisen leider auch zu oft geglaubt wird ...

Im aktuellen Prozessorgefluester werden die ersten Daten des ersten Manycoreprozessors von Intel mit dem Namen Larrabee diskutiert:

[..]ist Larrabee gedacht, ein Prozessor, der bisher mit 16 bis 24 Kernen gehandelt wurde, wahrscheinlich aber gleich mit 32 Kernen im nächsten Jahr debütieren wird – und zwar wie inzwischen durchdrang zur allgemeinen Überraschung

wohl mit genau den gut bekannten Pentium-Kernen: dem Pentium P54C. Hmmm ... P54C ... hmmm .... vorgestellt wurde dieser 1994. Das wäre ungefähr vergleichbar, wenn wir Niagara T2 auf Basis von SuperSPARC II (Super, nicht Ultra) gebaut hätten. Der Core wurde nämlich im selben Jahre vorgestellt

Vermutlich bezieht sich der Artikel im Intel@Research blog, den ich vor einigen Tagen verlinkt habe auf diese xPU, die Eigenschaften eines solchen Prozessors würden zu jenem dort geforderten Herangehen an Softwareentwicklung passen.

Interessant: Insgesamt soll diese xPU eine Gesamtstromaufnahme von 300 Watt haben (mithin das 3 bis 2,5 fache einer T2-CPU und fast das 4 fache einer T1-CPU). Ich weiss Larabee ist erstmal eine GPU ... aber glaubt jemand ernsthaft, dass eine GPU mit x86 kompatiblen Befehlssatz lange auf der Graphikkarte bleibt?

Es gibt hier übrigens noch eine kleine Ironie der Geschichte. Sun hat damaligerzeit viel über Multicores gelernt, als wir die MAJC-5200 CPU gebaut haben und dann auf den XVR-1000 und XVR4000 eingesetzt haben. Die XVR4000 war jene Graphikkarte, die sich nicht mit sowas läppischen wie einem PCI-Bus abgegeben hat, sondern direkt anstatt eines Prozessorsboards auf den Fireplane-Interconnect der V880 gesteckt wurde. Sun lernte damals soviel, dass manche Dinge erst mit Rock wieder in CPU fließen werden (Speculative Multithreading beispielsweise). Interessant ist nun, dass Intel ein solches Design jetzt auch zuerst auf einer Graphikkarte einsetzt. Nur: MAJC und die XVR4000 wurden September 2002 announced

Somit verbleibe ich mit einem Gruss an Herrn Siering und möchte mit einer genauso reisserischen Schlussformel enden: "CMT? Klar. Nur nicht auf dem Niveau von 1994!"

Posted by Joerg Moellenkamp in General at 13:46

Saturday, July 5. 2008

### **Good hotel in New York?**

I travel to New York end of this month for one week. Does anybody knows a good and affordable hotel in New York? I know, both for itself is not a problem, but i want both in one hotel ... liknow the Waldorf Astoria is excellent but this busts the budget of a poor SE

Posted by Joerg Moellenkamp in Travel at 22:07

### **links for 2008-07-05**

Screenshoot blurry in OS X? | creativebits

(tags: apple osx mac screenshots png screenshot terminal tutorial jpg)

Posted by del.icio.us at 13:35

Friday, July 4, 2008

## **Less Known Solaris features: Jumpstart Enterprise Toolkit - Part 15: Using flash archives for recovery**

You can't use flash without any further thinking for system recovery. As I wrote before, the new system is cloned and has a new personality, it's not a copy. But there is a trick. The basic trick My colleague Mario Beck showed it to me several years ago. The problem is the sysunconfig as it deletes the configuration information. The basic trick is simple. Before you create the flash archive, just do a backup of those files in one of the directories covered by the archive creation. Later on you can use this data to recover the configuration by copying those files in their original location.

Using an augmented Flash archive The flash archive is just a cpio copy of the complete system minus explicitly excluded parts. Everything on the system is included in the flash archive, not just the OS\footnote{otherwise the flash archives wouldn't make sense}. So we can easily rescue the personality of the system into the flash archive.

To simplify this task, I use the following script:#!/bin/sh

```
mkdir -p /var/opt/recovery
mkdir -p /var/opt/recovery/etc
cp -p /etc/hosts /var/opt/recovery/etc
cp -p /etc/shadow /var/opt/recovery/etc
cp -p /etc/passwd /var/opt/recovery/etc
cp -p /etc/vfstab /var/opt/recovery/etc
cp -p /etc/nodename /var/opt/recovery/etc
cp -p /etc/hostname. /var/opt/recovery/etc
cp -p /etc/dhcp. /var/opt/recovery/etc
cp -p /etc/defaultdomain /var/opt/recovery/etc
cp -p /etc/TIMEZONE /var/opt/recovery/etc
mkdir -p /var/opt/recovery/etc/inet
cp -p /etc/inet/netmasks /var/opt/recovery/etc/inet/netmasks
cp -p /etc/defaultrouter /var/opt/recovery/etc/defaultrouter
mkdir -p /var/opt/recovery/var/ldap
cp -p /etc/.rootkey /var/opt/recovery/etc
cp -p /etc/resolv.conf /var/opt/recovery/etc
cp -p /etc/sysidcfg /var/opt/recovery/etc
cp -p /var/ldap/ldap_client_cache /var/opt/recovery/var/ldap/ldap_client_cache
cp -p /var/ldap/ldap_client_file /var/opt/recovery/var/ldap/ldap_client_file
cp -p /var/ldap/ldap_client_cred /var/opt/recovery/var/ldap/ldap_client_cred
cp -p /var/ldap/cachemgr.log /var/opt/recovery/var/ldap/cachemgr.log
mkdir -p /var/opt/recovery/var/nis
cp -p /var/nis/NIS_COLD_START /var/opt/recovery/var/nis
mkdir -p /var/opt/recovery/var/yp
cp -R -p /var/yp/ /var/opt/recovery/var/yp
```

When you create a flash archive after running this script, it will include this directory in the archive, thus a newly installed machine with this flash archive will have this directory as well. So you can use it to recover the old status of the system. The process is simple. Just do a `cp -R /var/opt/recovery/* /`. The process of jumpstarting the server is identical to doing a normal flash install.

Posted by Joerg Moellenkamp in Solaris at 21:49

## **Less Known Solaris features: Jumpstart Enterprise Toolkit - Part 14: Using flash archives**

As I stated before, the installation mechanism of Jumpstart Flash is quite different from a normal installation. So I will start with a new template on a new server to demonstrate Jumpstart Flash. Creating a flash archive The first step to do a flash based install is the creation of a flash archive. Obviously you need an already installed system to create such an archive. To do so, I create a directory for the flash image at first:

```
bash-3.00# mkdir /flardir
bash-3.00# flarcreate -S -n "togusa" -x /flardir -R /flardir/togusa.flar
```

Full Flash  
Checking integrity...  
Integrity OK.  
Running precreation scripts...  
Precreation scripts done.  
Creating the archive...  
5144848 blocks  
Archive creation complete.  
Running postcreation scripts...  
Postcreation scripts done.

Running pre-exit scripts...  
Pre-exit scripts done.

bash-3.00#The second command generates the flar image itself. With this command, i generate the flash archive togusa.flar in the directory /flar. The -x option excludes the directory /flardir from the flasharchive. The \verb-R= specifies that the flash archive should contain all filesystems descending to /. The -S flag omits the size checks.# scp togusa.flar jmoekamp@192.168.10.1:/export/flar/togusa/togusa.flarAfter the creation of the flash archive, you have to transmit it to a server. It doesn't have to be the jumpstart server. it doesn't even have to be an NFS server. It just have to be reachable with HTTP or NFS by the server you plan to install. In this example we will use the Jumpstart Server for this task, thus we will use a share on this system.

Don't forget to share this directory via NFS:# echo "share -F nfs -o anon=0,sec=sys,ro -d \"Installation Images\" /export/install" >> /etc/dfs/dfstab

# shareall

Preparing the templateOkay, we want to install our node ishikawa with the flash image of togusa. At first like with all jet installation we add a hostname to our /etc/hosts

# echo "192.168.10.11 ishikawa" >> /etc/hostsNow we generate the template. We need only a small template for the basic installation.# make\_template ishikawa flash

Adding product configuration information for

+ base\_config

+ flash

Updating base\_config template specifics

Client template created in /opt/SUNWjet/TemplatesOkay, we have to fill the templates with standard settings.base\_config\_ClientArch=i86pc

base\_config\_ClientEther=08:00:27:DE:91:AB

base\_config\_ClientOS=sol10u5The use operating system here isn't important. You will not install this operating system. You will install the operating system contained in the flash archive.

Fill out the sysidsection of the system like with a normal system, you just need to fill the data to prevent the system from going interactive.

Now we get to the configuration of the flash install. You just define one or more locations of flash archives. When the installiion in your flash archive contains all recommended patches you can save some time at the installation and skip the installation by usig yes for the flash\_skip\_recommended\_patches parameter.

flash\_archive\_locations="nfs://192.168.10.1/export/flar/togusa/togusa.flar"

flash\_skip\_recommended\_patches="yes"

Now we have to create the Jumpstart environment for this client.# /opt/SUNWjet/bin/make\_client -f ishikawa

Gathering network information..

Client: 192.168.10.11 (192.168.10.0/255.255.255.0)

Server: 192.168.10.1 (192.168.10.0/255.255.255.0, SunOS)

Solaris: client\_prevalidate

Creating Client directory for ishikawa

Solaris: client\_build

Creating sysidcfg

Creating profile

Adding base\_config specifics to client configuration

Adding flash specifics to client configuration

FLASH: Modifying client profile for flash install

FLASH: Removing package/cluster/usedisk entries from profile

Solaris: Configuring JumpStart boot for ishikawa

```
Starting SMF services for JumpStart
Solaris: Configure PXE/grub build
Adding install client
Doing a TEXT based install
Leaving the graphical device as the primary console
Configuring ishikawa macro
Using local dhcp server
PXE/grub configuration complete
Running '/opt/SUNWjet/bin/check_client ishikawa'
Client: 192.168.10.11 (192.168.10.0/255.255.255.0)
Server: 192.168.10.1 (192.168.10.0/255.255.255.0, SunOS)
Checking product base_config/solaris
Checking product flash
FLASH: Checking nfs://192.168.10.1/export/flar/togusa/togusa.flar
-----
```

Check of client ishikawa

-> Passed....Please note the single line with FLASH: at the beginning. The JET framework checks for the availability of the flash archive. This prevents one of the most occurring problems with flash... a unaccessible flash archive at the installation.

When we look into the profile for ishikawa you will recognize, that all statements regarding cluster to install or similar stuff is removed. But there is a new statement. archive\_location specifies the location of the flar image and install\_type tells the system to do a flash install.# cat profile

```
#
# This is an automatically generated profile. Please modify the template.
#
# Created:    Sat May 24 12:50:20 CEST 2008
#
install_type flash_install
archive_location nfs://192.168.10.1/export/flar/togusa/togusa.flar
partitioning explicit
#
# Disk layouts
#
fileys      rootdisk.s0  free  /
fileys      rootdisk.s1  256  swap
# pwd
/opt/SUNWjet/Clients/ishikawa
#
```

While JumpstartingWhen you look into the logfiles of your system, you will notice a big change. You won't see a large amount of messages from the package installation. Instead of this you will just see the progress notifications of the extraction of the flar archive.

You will find this information in the middle of: /var/sadm/system/logs/install\_log[...]  
Beginning Flash archive processing

```
Extracting archive: togusa
  Extracted  0.00 MB ( 0% of 2512.13 MB archive)
[...]
  Extracted 2512.13 MB (100% of 2512.13 MB archive)
  Extraction complete
```

[...]The JET framework augments this with some housekeeping tasks like deleting the Solaris Volume Manager configuration. As usual you can look into the /var/opt/sun/jet/finish.log logfile to find the related messages:FLASH: Installing flash....

```
FLASH: Disabling /a/etc/lvm/mddb.cf -> /a/etc/lvm/mddb.cf.disabled
```

```
FLASH: Purging entries from /a/etc/lvm/md.tab
```

```
FLASH: Disabling mirrored_root_flag in /etc/system
```

```
FLASH: Cleanout crash dump area
```

```
FLASH: Clear out devfsadm filesAfter this you have a clone of your old system, but without the instance-specific configuration like ip addresses, hostnames and the volume manager configuration.
```

Posted by Joerg Moellenkamp in Solaris at 21:32

### Less Known Solaris features: Jumpstart Enterprise Toolkit - Part 13: Deep dive to post installation

As I told you before much of the configuration takes place after the installation executed by the original Jumpstart mechanism. We used several modules of the JET toolkit so far, thus this is a good moment to do a deep dive into the process that takes place after the normal Jumpstart installation.

The concept of this further installation steps relies pretty much completely on the concept of a so called finish script. Do you remember the rules.ok file? There was a finish script specified in that file for all installations: `bash-3.2$ cat rules.ok`

```
any any      Utils/begin  =      Utils/finish
# version=2 checksum=3114
```

The installation of the Solaris Operating Environment is equal to the normal Jumpstart process, as it relies on the same functions. But then JET comes into the game. After the installation has completed, the script `Utils/finish` is executed. But where is this file. It's relative to a directory we've specified before. Or to be exact, JET did that for us.

This is a snippet from the `menu.lst` for our system: `title Solaris_11 Jumpstart`

```
kernel /I86PC.Solaris_11-1/platform/i86pc/kernel/unix - install nowin -
B install_config=192.168.10.1:/opt/SUNWjet,sysid_config=192.168.10.1:/opt/SUNWje
t/Clients/togusa,install_media=192.168.10.1:/export/install/media/solaris/x86/nv
87,install_boot=192.168.10.1:/export/install/media/solaris/x86/nv87/boot
```

`module /I86PC.Solaris_11-1/x86.miniroot` The `Utils/finish` is relative to `install_config`, thus the executed finish script `192.168.10.1:/opt/SUNWjet/Utils/finish`. The NFS mount specified in `install_config` is one of the first mounts done on the new system and we can use the content of this directory throughout the installation process. By the way: This is the reason, why the `rules.ok` is located at this strange position.

We can study the further process in logfile of the installation. The complete log is located at `/var/opt/sun/jet/` in the file `jumpstart_install.log`. Let's start. At first the finish script starts to take copy some components from the Jumpstart server to the local disk. `Creating directory: /a/var/opt/sun/jet/post_install`

```
Creating directory: /a/var/opt/sun/jet/Utils
```

```
Creating directory: /a/var/opt/sun/jet/config
```

```
Creating directory: /a/var/opt/sun/jet/js_media/patch
```

```
Creating directory: /a/var/opt/sun/jet/js_media/pkg
```

```
Copying file Clients/togusa/sysidcfg to /a/var/opt/sun/jet/config/sysidcfg
```

```
Copying file Clients/togusa/profile to /a/var/opt/sun/jet/config/profile
```

```
Copying file Clients/togusa/host.config to /a/var/opt/sun/jet/config/host.config
```

```
Copying file Utils/solaris/releaseinfo to /a/var/opt/sun/jet/config/releaseinfo
```

```
Copying functions to /a/var/opt/sun/jet/Utils/lib
```

```
Copying file Clients/togusa/module_hints to /a/var/opt/sun/jet/config/module_hints
```

As you see, the JET copies over part of the toolkit as well as configuration files to the new position. But why are all those directories relative to `/a`. Well this is easy. In the netbooted mini root, the local disks are mounted relative to `/a`. The reasoning behind this copy is relatively simple. In the next boots the contents of `/opt/SUNWjet/` aren't available any longer, as the system doesn't mount it in the next steps. The post installation scripts rely on some helper function. The simplest way to ensure their availability under all circumstances (even when your installation disables the network) is a simple copy.

The next step is the mounting of the directories with patches and product. `Mounting`

```
nfs://192.168.10.1/export/install/patches on /var/opt/sun/jet/js_media/patch
```

```
Mounting nfs://192.168.10.1/export/install/pkgs on /var/opt/sun/jet/js_media/pkg
```

Now it gets a little bit complicated, but I will simplify it as far as I can. Depending on the complexity of the setup your configuration may use one or more so called products. A product in JET-speak is a JET module for the installation and configuration of a certain area of the operating system. In any case you will use the product `base_config` but there may be several ones. Our example uses the products `base_config`, `custom`, `sds` and `jass`. The JET framework gathers this information from the configuration template. It's stored in this line:

```
base_config_products=" custom sbd sds jass"
```

The framework takes this information and to execute the install script in any of this directories. For example it starts at first the install script in `/opt/SUNWjet/Products/base_config/solaris` as this is default for every installation, after this it will step forward by executing the install script in any of the product directories. The install script has two important roles. At first it installs packages, patches and files according to the configuration in the templates. At second it registers so called `post_install` scripts.

`Post installation scripts` `post_install` scripts are executed at the next boots. You can order this scripts by specifying a certain boot level. After the execution of all scripts in a boot level, the system reboots. For example all scripts with boot level 1 are executed after the first reboot, all `post_install` scripts with the boot level 2 are executed after the second reboot and so on. These scripts are installed in `/var/opt/sun/jet/post_install`

But how get this scripts to execution at the following reboots. JET copies a init.d script to the new system. On Solaris 10 it creates a matching SMF service. The function of this script is quite simple: Gather the actual boot level by reading the file `/var/opt/sun/jet/post_install/Platform`, execute all scripts in the boot level, increment the bootlevel by one and reboot the system.

An example for boot levels and postinstall scriptsWe've done the first boot. It booted from the network. The installation of the Solaris Operating Environment succeeded, thus the script from the finish column is executed by the netbooted installation system. After some configuration tasks the system starts to register postinstall scripts.SDS: Installing sds....

```
SDS: Copying file sds_functions to /a/var/opt/sun/jet/Utils/sds_functions
```

```
SDS: Creating directory: /a/var/opt/sun/jet/post_install/Platform/1
```

```
SDS: Register postinstall script 'create_fmthard' for boot 1
```

```
SDS: Register postinstall script 'set_boot_device' for boot 1
```

```
SDS: Register postinstall script 'create_metadb' for boot 1
```

```
SDS: Register postinstall script 'create_root_mirror' for boot 1
```

```
SDS: Register postinstall script 'attach_mirrors' for boot z
```

```
SDS: Register postinstall script 'create_user_devices' for boot 1
```

```
SDS: Register postinstall script 'attach_user_mirrors' for boot z
```

```
SDS: WARNING: unable to locate specified md.tab for SDS.
```

```
SDS: Enable md:mirrored_root_flag in /etc/system
```

You see, that the SDS product registered some scripts for boot level 1 and some for boot level z. Let's look further into the installation log. This happens after the first reboot:SDS: Running 001.sds.001.create\_fmthard

```
fmthard: New volume table of contents now in place.
```

```
SDS: Running 001.sds.001.set_boot_device
```

```
SDS: Setting OBP nvramrc rootdisk path
```

```
[...]
```

```
SDS: Create 3 copies on c1d1s7
```

```
metadb: waiting on /etc/lvm/lock
```

```
SDS: Running 001.sds.003.create_root_mirror
```

```
SDS: Setting OBP nvramrc rootmirror path
```

```
[...]
```

```
SDS: Installing GRUB on Alternate Boot Disk.
```

```
SDS: Running 001.sds.007.create_user_devices
```

Later on you you will recognize the scripts for boot level z

```
Running additional install files for reboot z
```

```
SDS: Running 003.sds.001.attach_mirrors
```

```
SDS: Attach d12 to d10
```

```
SDS: Attach d22 to d20
```

```
SDS: Attach d12 to d10
```

```
d10: submirror d12 is attached
```

```
SDS: Attach d22 to d20
```

```
d20: submirror d22 is attached
```

```
SDS: Running 003.sds.002.attach_user_mirrors
```

With this mechanism, you can implement installation processes with package or programm installations that need several boots to fulfill.

The end of the post installationAt the very end the init.d script is deleted together with the matching SMF service. The logfiles and the post intallation scripts stay on the local disk.

Posted by Joerg Moellenkamp in Solaris at 21:02

## **Less Known Solaris features: Jumpstart Enterprise Toolkit - Part 12: Automatic hardening with SST**

It's a best practice to harden a system before you place it into your production network to reduce possible attack vectors. Sun developed the Solaris Security Toolkit for this task to collect all the knowledge about hardening Solaris in a tool thats simple to use. I've wrote already about the usage of the toolkit in another installment of the LessKnownSolarisFeatures series.

It would be really neat to have an automatized hardening of new systems. The Jumpstart Enterprise Toolkit can do exactly this with the help of JASS module.

Preparing the Jumpstart for installation  
At first you uncompress and untar the JASS distribution.  
# copy\_product\_media jass 4.2.0 /export/home/jmoekamp i386  
Transferring package instance

Packages copied.

Okay, but we have to do another step. There is a patch for the version 4.2.0 of the Solaris Security Toolkit: 122608-xx. At first we have to tell JET that there is a patch for this product and version. We have to modify the file patch.matrix in /opt/SUNWjet/Products/jass:#

```
# Patch matrix for Solaris Security Toolkit (JASS)
```

```
#
```

```
# :::
```

```
#
```

```
10:i386:4.2.0:122608
```

Now it's easy to integrate the patch. I've unpacked the patch in the directory \verb=/export/home/jmoekamp/patch\_jass= before:

```
# copy_product_patches jass 4.2.0 /export/home/jmoekamp/patch_jass i386
```

Patches copied. Configuring the template  
Okay, you have to configure only a few basic variables to trigger the automatic hardening of your new installation.

```
base_config_products=" custom sbd sds jass"
```

```
jass_product_version="4.2.0"
```

```
jass_execute="secure.driver"
```

Please refer to the SST tutorial about the inner mechanisms and the concept of drivers in the Solaris Security Toolkit.

After Jumpstarting  
Okay, it's time to reboot the machine we want to install again. At first, all is like at the runs before. But then we see some further lines in the logfile.

```
JASS: Installing Solaris Security Toolkit (JASS) 4.2.0...
```

```
JASS: Installing SUNWjass from: /a/var/opt/sun/jet/js_media/pkg/jass/4.2.0/i386
```

Installation of was successful.

```
JASS: SUNWjass installation complete
```

```
JASS: Register postinstall script 'postinstall' for boot zlt's important to know, that the above configuration installed the SUNWjass package on the system, patches it there and then run runs the toolkit installed on the system.
```

The hardening of the system is started in the background. After a while you will recognize the work of the script. The backup files of the Solaris Security Toolkit are dispersed all over the directories.

```
bash-3.00$ ls -l /etc/*.JASS*  
-rw-r--r-- 1 root other 372 May 23 19:48 /etc/coreadm.conf.JASS.20080523195314
```

```
[...]
```

```
-rw-r--r-- 1 root sys 362 May 23 19:43 /etc/vfstab.JASS.20080523195420
```

```
bash-3.00$  
After the completion of the background JASS run, you have a automatically installed, patched, customized, mirrored and hardened system.
```

Posted by Joerg Moellenkamp in Solaris at 20:30

## **Less Known Solaris features: Jumpstart Enterprise Toolkit - Part 11: Mirrored boot disks**

Okay, in enterprise computing you wouldn't use a system without redundant boot disks (at least, when you haven't an application that can survive a loss of computing nodes without a problem). So it would be nice to automatically configure a mirror of the boot disks with the Solaris Volume Manager. I assume in the following part, that you have a working knowledge with the SVM. When this not the case, it isn't really a big problem, as this part it somewhat selfexplaining when you are aware of the concept of RAID.

```
# make_template -f -M -T togusa togusa sds  
Adding product configuration information for  
+ sds
```

Configuration in the template  
As in the last example i will just include the interesting parts of the configuration. You have to modify some lines:

```
base_config_products=" custom sbd sds"  
sds_product_version="default"
```

```
sds_root_mirror="c1d1"
```

```
sds_use_fmthard="yes"
```

```
sds_database_locations="rootdisk.s7:3"
```

```
sds_database_partition="s7:32"
```

```
sds_metadb_size=""
```

```
sds_root_alias="rootdisk"
sds_root_mirror_devalias_name="rootmirror"
sds_mirrored_root_flag="yes"
```

At first you have to include the sds module to the base\_config\_product line. Then choose the disk you want to use for mirroring, in my case it's /dev/dsk/c1d1. Line 4 orders the sds module to copy the vtoc of the first mirror to the second one. When there are only two disks in your system you have to specify the sds\_mirror\_root\_flag with yes in line 10. Solaris Volume Manager want to see the metadb replica copies on at least three disks. When you want to find out the correct state of an situation you need at least three copies to be sure. With only two copies, one or the other version may be correct, with three you have two correct copies, so there is a good chance that the two copies represent the correct state. Solaris Volume Manager likes to distribute those over at least three disk to ensure that the failure of a disk won't take out exactly the half of the copies. You have to tell the system that it shouldn't obey this rule. sds\_mirror\_root\_flag specifies that.

Effects of the configuration Okay, after the installation we can log into the system. You may notice after the installation a vast amount of accesses to the harddisk. The devices for / and the swap are not longer referenced by an device name. Instead of this you will find the names of the logical devices of the Solaris Volume Manager.

```
-bash-3.2$ cat vfstab
#device      device      mount      FS   fsck  mount  mount
#to mount    to fsck     point      type pass  at boot options
#
fd - /dev/fd fd - no -
/proc - /proc proc - no -
/dev/md/dsk/d20 - - swap - no -
/dev/md/dsk/d10 /dev/md/rdisk/d10 / ufs 1 no logging
/devices - /devices devfs - no -
ctfs - /system/contract ctfs - no -
objfs - /system/object objfs - no -
swap - /tmp tmpfs - yes -
-Let's look after the exact configuration of the volume manager.
```

```
-bash-3.2$ metastat
```

```
d20: Mirror
```

```
  Submirror 0: d21
```

```
    State: Okay
```

```
  Submirror 1: d22
```

```
    State: Okay
```

```
  Pass: 1
```

```
  Read option: roundrobin (default)
```

```
  Write option: parallel (default)
```

```
  Size: 530145 blocks (258 MB)
```

```
d21: Submirror of d20
```

```
  State: Okay
```

```
  Size: 530145 blocks (258 MB)
```

```
  Stripe 0:
```

Device	Start Block	Dbase	State	Reloc	Hot Spare
c0d0s1	0	No	Okay	Yes	

```
d22: Submirror of d20
```

```
  State: Okay
```

```
  Size: 530145 blocks (258 MB)
```

```
  Stripe 0:
```

Device	Start Block	Dbase	State	Reloc	Hot Spare
c1d1s1	0	No	Okay	Yes	

```
d10: Mirror
```

```
  Submirror 0: d11
```

```
    State: Okay
```

```
  Submirror 1: d12
```

```
    State: Okay
```

```
  Pass: 1
```

Read option: roundrobin (default)  
Write option: parallel (default)  
Size: 32836860 blocks (15 GB)

d11: Submirror of d10

State: Okay  
Size: 32836860 blocks (15 GB)

Stripe 0:

Device	Start Block	Dbase	State	Reloc	Hot Spare
c0d0s0	0	No	Okay	Yes	

d12: Submirror of d10

State: Okay  
Size: 32836860 blocks (15 GB)

Stripe 0:

Device	Start Block	Dbase	State	Reloc	Hot Spare
c1d1s0	0	No	Okay	Yes	

Device Relocation Information:

Device Reloc Device ID

c1d1 Yes id1,cmdk@AVBOX\_HARDDISK=VB37711a7b-00576cc9

c0d0 Yes id1,cmdk@AVBOX\_HARDDISK=VB97a90791-9d191449The default numbering scheme for the solaris volume manager is quite simple: The mirror is designated with the first number in a decade (e.g. 10,20,30), the parts of a mirror are numbered with the next free numbers in the decade. For example: The first mirror half of the first mirror get the number 11, the second number gets the number 12.

It takes a while until the mirrors are in sync, but after this you have a automatically installed, patched, customized and mirrored system.

Posted by Joerg Moellenkamp in Solaris at 20:14

## Less Known Solaris features: Jumpstart Enterprise Toolkit - Part 10: A polished basic installation

Okay, now we've done a basic installation. But mostly we do a standard set of customizations on every system we touch, like installing a few essential tools or integrating the actual recommended patch cluster. So we want to polish the standard installation a little bit. We will extend a basic Solaris 10 Update 5 installation with the following items:

- installing the SMCjoe package
- installing the recommended Patch Cluster
- configuring the Secure-by-default mechanism to limited

Adding the recommended patch clusterOkay, at first we have to copy the patches into the the jumpstart. JET provides the copy\_solaris\_patches command to do so.

```
# copy_solaris_patches 10_x86 /export/home/jmoekamp/10_x86_Recommended
```

Copied....The second parameter specifies where your unpacked your patch cluster. You don't have to configure anything in the templates. Every new Solaris 10 installation from now on with the x86 architecture will be installed with the matching recommended patch cluster.

Adding custom packagesOkay, almost everybody installs some custome packages on his/her system. For example, one of the first things i'm installing on new systems is \verb=joe= to have an WordStar compatible editor:# pkgtrans

```
joe-3.5-sol10-x86-local /tmp all
```

Transferring package instance

```
# copy_custom_packages /tmp i386 SMCjoe
```

Transferring package instance

Packages copiedjoe depends on the ncurses library. So we copy this package as well to our JET server.

```
# pkgtrans ncurses-5.6-sol10-x86-local /tmp all
```

Transferring package instance

```
# copy_custom_packages /tmp i386 SMCncurs
```

Transferring package instance

Packages copied

Extending the template Okay, we need more modules to fulfill this task. You don't need to delete the old one and retype all the data in the template. The `make_template` script can use an old template to create a new one while retaining all the old values. Basically you use the same name for the old and new template.

```
# ./make_template -f -T togusa togusa base_config custom sbd
```

Adding product configuration information for

```
+ custom
+ sbd
```

Updating `base_config` template specifics

Client template created in `/opt/SUNWjet/Templates` When you look into `/opt/SUNWjet/Templates/togusa` you will recognize your old configuration, with a large amount of new lines. But we have to change only a few ones: At first we change the operation system. We've used OpenSolaris in the last example, but there are no patches for this release. But we've copied a Solaris 10 media with the name `sol10u5` earlier: `base_config_ClientOS=sol10u5` Okay, now we want to install the additional packages. You have to add the names of the packages in the line `\verb=custom_packages=.custom_packages="SMCncurs SMCjoe"` You don't have to configure the Secure by default module, as this module configures the `\verb=limited=` service set when it's used in the template. Patching of the Solaris OE doesn't need configuration as well. So we have to change only these two lines.

The installation Okay, you can start the installation by starting a network boot. This time the installation takes a little bit longer. First the system starts to install the recommended patch cluster. As I've used a quite actual update, most patches are already installed. but a few ones will find their way into the installation. `BASE_CONFIG: Installing base_config...`

```
BASE_CONFIG: Product base_config started
```

```
BASE_CONFIG: Installing additional patches from : 10_x86_Recommended
```

```
BASE_CONFIG: Using patch_order file for patch installation sequence
```

```
BASE_CONFIG:
```

```
BASE_CONFIG: ----- Installing patches for product: 10_x86_Recommended -----
```

```
BASE_CONFIG:
```

```
BASE_CONFIG: Patch 120720-02 is already installed.
```

```
BASE_CONFIG: Patch 124458-01 is already installed.
```

```
[...]
```

```
BASE_CONFIG: >
```

```
[...]
```

```
BASE_CONFIG: Patch 122175-03 is already installed.
```

```
BASE_CONFIG: Patch installation completed.
```

```
BASE_CONFIG: No HW specific packages for platform i86pc
```

```
BASE_CONFIG: No HW specific patches for platform i86pc Okay, now the system starts to install the additional packages.
```

```
CUSTOM: Installing custom....
```

```
CUSTOM: Installing SMCncurs from: /a/var/opt/sun/jet/js_media/pkg/custom/i386
```

```
Installation of was successful.
```

```
CUSTOM: SMCncurs installation complete
```

```
CUSTOM: Installing SMCjoe from: /a/var/opt/sun/jet/js_media/pkg/custom/i386
```

```
Installation of was successful.
```

```
CUSTOM: SMCjoe installation complete
```

```
The following module doesn't do really much, as the configuration of the service profile is activated by the \verb=sysidcfg= file.
```

```
SBD: Installing sbd....
```

```
SBD: configured
```

Effects of the new modules Let's check the results of the installation. At first, we look for the custom packages:

```
# pkginfo | grep "SMC"
```

```
application SMCjoe          joe
```

```
application SMCncurs       ncurses
```

When we look for one of the installed patches, we will see its successful installation to the system:

```
# showrev -p | grep "120273-20"
```

```
Patch: 120273-20 Obsoletes: Requires: 119043-09, 121902-01, 122532-04 Incompatibles: Packages: SUNWbzip, SUNWsmagt, SUNWsmcmd, SUNWsmmgr
```

Posted by Joerg Moellenkamp in Solaris at 19:25

## **Less Known Solaris features: Jumpstart Enterprise Toolkit - Part 9: Basic installation of a new system**

At first we will do a really basic install. No tricks, just the pure operating system. Nevertheless this part will be a little bit longer as i will do a technical deep-dive into the process of installation in this example to show you the inner workings of JET with this installation as an example.

The template for the first installationAt first we have to create a template for the system. This is really easy:#

```
make_template togusa
```

Adding product configuration information for

```
+ base_config
+ custom
+ sds
+ vts
+ explo
+ flash
+ san
+ jass
+ zones
+ sbd
```

Updating base\_config template specifics

Client template created in /opt/SUNWjet/TemplatesOkay, this is too much ... at start we don't want all this modules right now. We can add them later, without loosing the configuration. Let's just use the module for the basic configuration:

```
# make_template -f togusa base_config
```

Adding product configuration information for

```
+ base_config
```

Updating base\_config template specifics

Client template created in /opt/SUNWjet/Templates

Even the basic template is quite long. I've reduced it for this tutorial by deleting all comments, all empty lines and all variables without a value.

```
base_config_ClientArch=i86pc
base_config_ClientEther=08:00:27:97:29:1E
base_config_ClientOS=nv87
base_config_client_allocation="grub"
base_config_sysidcfg_nameservice=NONE
base_config_sysidcfg_network_interface=PRIMARY
base_config_sysidcfg_ip_address=192.168.10.10
base_config_sysidcfg_netmask=255.255.255.0
base_config_sysidcfg_root_password="boajrOmU7GFmY"
base_config_sysidcfg_system_locale="C"
base_config_sysidcfg_timeserver=localhost
base_config_sysidcfg_timezone="Europe/Berlin"
base_config_sysidcfg_terminal=vt100
base_config_sysidcfg_security_policy=NONE
base_config_sysidcfg_protocol_ipv6=no
base_config_sysidcfg_default_route=192.168.10.1
base_config_x86_nowin="yes"
base_config_label_disks="all"
base_config_profile_cluster=SUNWCuser
base_config_profile_usedisk=rootdisk.
base_config_profile_root=free
base_config_profile_swap=256
base_config_ufs_logging_filesys="all"
base_config_profile_del_clusters="SUNWCpms SUNWCpmsx SUNWCdial SUNWCdialx"
```

```
base_config_dns_disableforbuild="yes"
base_config_update_terminal="yes"
base_config_enable_savecore="yes"
base_config_dumpadm_minfree="20000k"
base_config_noautoshtutdown="pm_disabled"
```

Let's dissect this template:

Line 1-3: These lines are the most basic ones. The first line defines the architecture of the system. The next line is the Ethernet-Address of the new system. The third one specifies the new operating system.

Line 4: This line specifies, how the new system gathers the most basic informations like it's own IP.

Line 5-16: Do you remember the part about `\verb=sysidcfg=`. The values for this files are defined in this part of the emplate. The `sysidcfg` will be filled with the values of this variables

Line 17: This line tells the system to suppress the start of the windowing system.

Line 18: Solaris needs a disk label on the disks for the system. This directive tells the system to write this label to all disks.

Line 19-22: Another known phrase ... profile. Here you specify the partitioning for the system and what packages will be installed on it.

Line 23-end: There are several further statements. Please the original file for an explanation.

Okay after this step, we have to generate the configuration for the Jumpstart mechanism. This is really easy: #

```
make_client -f togusa
```

Gathering network information..

```
Client: 192.168.10.10 (192.168.10.0/255.255.255.0)
```

```
Server: 192.168.10.1 (192.168.10.0/255.255.255.0, SunOS)
```

Solaris: client\_prevalidate

Solaris: client\_build

Creating sysidcfg

Creating profile

Adding base\_config specifics to client configuration

Solaris: Configuring JumpStart boot for togusa

Starting SMF services for JumpStart

Solaris: Configure PXE/grub build

Adding install client

Doing a TEXT based install

Leaving the graphical device as the primary console

Configuring togusa macro

Using local dhcp server

PXE/grub configuration complete

Running `'/opt/SUNWjet/bin/check_client togusa'`

```
Client: 192.168.10.10 (192.168.10.0/255.255.255.0)
```

```
Server: 192.168.10.1 (192.168.10.0/255.255.255.0, SunOS)
```

Checking product `base_config/solaris`

-----

Check of client togusa

-> Passed....

The nice thing about the `make_client` command: It doesn't just generate the Jumpstart configuration. It checks for the most dumb errors like forgetting to share the directory of your Solaris media with NFS. So you can detect many problems at an early stage. You don't have to wait until the jumpstart client comes up just to detect that there is no NFS or no DHCP config.

The generated Jumpstart configuration files Okay, let's look into the `\verb=/tftpboot=` directory at first. As the system uses `pxegrub` we need a `menu.lst`

```
-bash-3.2$ cat /tftpboot/menu.lst.0108002797291E
```

```
default=0
```

```
timeout=2
```

```
title Solaris_11 Jumpstart
```

```
kernel /I86PC.Solaris_11-1/platform/i86pc/kernel/unix - install nowin -B
```

```
install_config=192.168.10.1:/opt/SUNWjet,sysid_config=192.168.10.1:/opt/SUNWjet/Clients/togusa,install_media=192.1
```

```
68.10.1:/export/install/media/solaris/x86/nv87,install_boot=192.168.10.1:/export/install/media/solaris/x86/nv87/boot
```

```
module /I86PC.Solaris_11-1/x86.miniroot
```

```
-bash-3.2$
```

In the GRUB configuration we not only load the Kernel, we additionally name the location of the Jumpstartserver, the exact location and name of the `sysidconfig` file, the position of our installation media and at last the location of the

miniroot. In our example all locations are NFS locations.

Okay, the install\_config directory is the first important location. We find the rules.ok file there.

```
-bash-3.2$ cat /opt/SUNWjet/rules.ok
any any          Utils/begin  =    Utils/finish
# version=2 checksum=3114Okay, now let's have a look in the specified profile file:
-bash-3.2$ cat /opt/SUNWjet/Clients/togusa/profile
#
# This is an automatically generated profile. Please modify the template.
#
# Created:    Mon May 19 21:47:50 CEST 2008
#
install_type  initial_install
system_type   server
cluster       SUNWCuser
partitioning  explicit
#
# Disk layouts
#
filesystems   rootdisk.s0  free  /
filesystems   rootdisk.s1  256   swap
cluster SUNWCpm delete
cluster SUNWCpmx delete
cluster SUNWCdial delete
cluster SUNWCdialx delete
```

As I wrote before, we have to give the system an identity. The `\verb=sysidcfg` is responsible for this task, thus we find such a file in our directory. Our new system will use it when the installation has completed.

```
-bash-3.2$ cat /opt/SUNWjet/Clients/togusa/sysidcfg
name_service=NONE
root_password=boajrOmU7GFmY
system_locale=C
timeserver=localhost
timezone=Europe/Berlin
terminal=vt100
security_policy=NONE
nfs4_domain=dynamic
network_interface=PRIMARY {hostname=togusa ip_address=192.168.10.10 netmask=255.255.255.0 protocol_ipv6=no
default_route=192.168.10.1}
```

The installation boot

This leaves you to do one thing. Configure the system to start with PXE in the BIOS of your system. And the system will boot via network and starts to install a system. After a while the installation will be complete. The system will boot up. One tip when you use Virtualbox for the installation. When the system tries to reboot the first time halt the system and switch off the network boot again. Otherwise the system will install again and again. BTW: The default rootpassword is newroot.

You can look for the logfile of the installation at `\verb=/var/sadm/system/logs/`:

Configuring disk (c0d0)

- Creating Fdisk partition table

Fdisk partition table for disk c0d0 (input file for fdisk(1M))

```
type: 130 active: 128 offset: 16065 size: 33527655
```

```
type: 100 active: 0 offset: 0 size: 0
```

```
type: 100 active: 0 offset: 0 size: 0
```

```
type: 100 active: 0 offset: 0 size: 0
```

- Creating Solaris disk label (VTOC)

- Processing the alternate sector slice

#### Creating and checking UFS file systems

- Creating / (c0d0s0)

Warning: 1608 sector(s) in last cylinder unallocated

/dev/rdisk/c0d0s0: 31744440 sectors in 5167 cylinders of 48 tracks, 128 sectors

15500.2MB in 323 cyl groups (16 c/g, 48.00MB/g, 5824 i/g)

super-block backups (for fsck -F ufs -o b=#) at:

32, 98464, 196896, 295328, 393760, 492192, 590624, 689056, 787488, 885920,

Initializing cylinder groups:

.....

super-block backups for last 10 cylinder groups at:

30776480, 30874912, 30973344, 31071776, 31170208, 31268640, 31367072,

31457312, 31555744, 31654176

#### Beginning Solaris software installation

Installation of was successful.

[...]

Installation of was successful.

#### Solaris 11 software installation succeeded

Solaris 11 packages fully installed

SUNWkvm

[...]

SUNWsolnm

#### Customizing system files

- Mount points table (/etc/vfstab)

fd - /dev/fd fd - no -

/proc - /proc proc - no -

/dev/dsk/c0d0s1 - - swap - no -

/dev/dsk/c0d0s0 /dev/rdisk/c0d0s0 / ufs 1 no -

/devices - /devices devfs - no -

sharefs - /etc/dfs/sharetab sharefs - no -

ctfs - /system/contract ctfs - no -

objfs - /system/object objfs - no -

swap - /tmp tmpfs - yes -

- Network host addresses (/etc/hosts)

- Environment variables (/etc/default/init)

#### Cleaning devices

#### Customizing system devices

- Physical devices (/devices)

- Logical devices (/dev)

#### Installing boot information

- Updating boot environment configuration file

- Installing boot blocks (c0d0)

- Installing boot blocks (/dev/rdisk/c0d0s0)

Creating boot\_archive for /a

updating /a/platform/i86pc/boot\_archive

updating /a/platform/i86pc/amd64/boot\_archive  
You see that the system has gone through a perfect automated installation.

Posted by Joerg Moellenkamp in Solaris at 18:22

### **Less Known Solaris features: Jumpstart Enterprise Toolkit - Part 8: Preparing the first installation**

Now you have to copy your Solaris install media. Obviously you need the Solaris files before you can install them to a

system via network. You can do this on two ways.

From a mounted DVD media When you already have burned media, you can use it for the copy process. Just put it into the drive, the volume management mounts it and you can copy it from there.

```
# /opt/SUNWjet/bin/copy_solaris_media -d /export/install/media/solaris/x86/nv87 -n nv87 /cdrom/sol_11_x86
Copying Solaris image....
Verifying target directory...
Calculating the required disk space for the Solaris_11 product
Calculating space required for the installation boot image
Copying the CD image to disk...
Copying Install Boot Image hierarchy...
Copying /boot netboot hierarchy...
Install Server setup complete
```

Added Solaris image nv87 at the following location:  
Media: /export/install/media/solaris/x86/nv87

```
removing directory /export/install/media/solaris/x86/911
#
```

Let's dissect the command: -d specifies the target, where you copy the operating system. -n specifies a name for this media. From now on you refer this solaris media as nv87 in the templates for JET. At the end you specify the location, where the media is located at the moment.

From a .iso file Okay, you've downloaded your Solaris media. You don't have to burn it, you can use the \verb=.iso= files directly:

```
# ./copy_solaris_media -d /export/install/media/solaris/x86/sol10u5 -n sol10u5 -i /export/home/jmoekamp
sol-10-u5-ga-x86-dvd.iso
Created loopback device /dev/lofi/1 for /export/home/jmoekamp/sol-10-u5-ga-x86-dvd.iso
mounted /export/home/jmoekamp/sol-10-u5-ga-x86-dvd.iso at /export/install/media/solaris/x86/790/slices/s0 (of type
hsfs)
Copying Solaris image....
Verifying target directory...
Calculating the required disk space for the Solaris_10 product
Calculating space required for the installation boot image
Copying the CD image to disk...
Copying Install Boot Image hierarchy...
Copying /boot x86 netboot hierarchy...
Install Server setup complete
```

Added Solaris image sol10u5 at the following location:  
Media: /export/install/media/solaris/x86/sol10u5

```
Unmounting /export/install/media/solaris/x86/790/slices/s0
removing device /dev/lofi/1
removing directory /export/install/media/solaris/x86/790
```

The command line for using a .iso file is quite similar. You just specify with the -i that an .iso file has to be used and in which directory it should search for it. The last parameter is the name of the .iso file itself. The system mounts the dvd image by using the loopback facility of Solaris and copies the media to its target location afterwards.

Looking up the existing Solaris versions JET provides script to lookup the versions of Solaris you've copied to your Jumpstart Server. With the list\_solaris\_locations script you look up the version and the location of your solaris medias.

```
# ./list_solaris_locations
Version      Location
-----
nv87         /export/install/media/solaris/x86/nv87
sol10u5      /export/install/media/solaris/x86/sol10u5
```

Posted by Joerg Moellenkamp in Solaris at 17:35

## Less Known Solaris features: Jumpstart Enterprise Toolkit - Part 7: Installing the JET Server

After gathering all parts for the installation, we can install the first part of our new JET server.

Preparation of the systemA jumpstart server doesn't have to be a big server. But it's good practice to take the following into consideration:

You have to copy your installation media onto the server. A Solaris version needs up to 4 Gigabyte on disk. Additionally you need space for further products (like other applications) and patches. The actual recommended patch cluster for Solaris 10 x86 is over 300 MB large for example. Depending on the amount of Solaris versions you want to provide, it's a good idea not to spend the rest of an existing partition to your Jumpstart server.

It's a wise choice to have a fast network. As the system isn't hardened and unconfigured at the first install, the paranoid prefer to use a free Gigabit-Ethernet port at my jumpstart server and use a separated VLAN for initial configuration. Or you use your admin LAN for this purpose. Either ways, for an fast installation gigabit ethernet is nice. When you have found a good place for your JET server, we can start with the installation.

The installationAfter copying the JET package somewhere on our new Jumpstart server, we have to install the package:#  
pkgadd -d jet.pkg

The following packages are available:

- 1 JetEXPLO jet explo product  
(sparc) 3.1.11
- 2 JetFLASH JET flash product  
(sparc) 3.1.8
- 3 JetJASS JASS product  
(sparc) 3.0.14
- 4 JetRBAC JET RBAC product  
(sparc) 1.1.5
- 5 JetSAN JET san product  
(sparc) 3.1.7
- 6 JetSBD Secure By Default product  
(sparc) 1.0.2
- 7 JetSDS JET sds product  
(sparc,i386) 3.4.4
- 8 JetVTS JET VTS product  
(sparc) 3.0.11
- 9 JetWanBoot JET WanBoot support  
(sparc) 1.1.1
- 10 JetZONES JET Zones module  
(sparc) 1.1.12

... 2 more menu choices to follow;  
for more choices, to stop display:

- 11 SUNWjet Sun JumpStart Enterprise Toolkit  
(sparc,i386) 4.4
- 12 SUNWjetd JET Documentation  
(sparc) 3.3.1

Select package(s) you wish to process (or 'all' to process all packages). (default: all) [?,??,q]: all

All this modules implements in the JET framework:

JetEXPLO: Installs and configures the Sun Explorer

JetFLASH: Module to control Jumpstart Flash installes

JetJASS: executes the Solaris Security Toolkit on the new host

JetRBAC: configures the Role Based Access Control on the JET server to enable a role for all users allowed to trigger a Jumpstart Installation

JetSAN: configures the SAN framework of Solaris

JetSBD: configures the Secure-by-default setting

JetSDS: configures the Solaris Volume Management

eVTS: installs the Sun Validation Test Suite. It tests and validates Sun hardware by verifying the connectivity and

functionality of hardware devices, controllers and peripherals.

JetWANboot: configures the Jumpstart facilities for installation over the WAN

JetZONES: configures Solaris Zones on the newly installed zones.

As the package is quite small with half a megabyte, i always install all packages on a jumpstartserver.

Processing package instance from

Sun JumpStart Enterprise Toolkit(sparc,i386) 4.4

Copyright 2007 Sun Microsystems, Inc. All rights reserved.

Use is subject to license terms.

The selected base directory must exist before installation is attempted.

Do you want this directory created now [y,n,?,q] y

Using as the package base directory.

[...]

Processing package instance from

Secure By Default product(sparc) 1.0.2

#

This is all we have to do for the installation. Easy, isn't it ?

Posted by Joerg Moellenkamp in Solaris at 17:33

## **Less Known Solaris features: Jumpstart Enterprise Toolkit - Part 6: Prerequisites**

Okay, before working with the Jumpstart Enterprise Toolkit we have to prepare some stuff. We need systems, a network (as fast as possible to reduce installation time, but even 10 MBit/s would suffice), operating system ISOs etc. Okay ... We need several systems for this test. I've used Virtualbox in the preparation of this tutorial. You can use other virtualisation software, but one thing is important: It must support PXE boot of the guests.

I've populated my /etc/hosts with the following hostnames. In the case you wonder about the hostnames, these names are characters from Ghost in a Shell. And just in case you search for kusanagi ... this is the name for the system hosting the virtual machines:192.168.10.1 aramaki

192.168.10.10 togusa

192.168.10.11 ishiwakaIn my tutorial aramaki will serve as the Jumpstart server, togusa and ishiwaka are the installation targets.

Packages and ISOsFor this tutorial i've used the following packages and ISOs:Solaris 10 Update 5: This is the operating environment i will use to demonstrate automated patchingOpensolaris Community Edition Build 87: aramaki runs with this operating environment and it's no problem to jumpstart Opensolaris CE oder DE with JET or to use Opensolaris for the Jumpstart ServerRecommended Patch ClusterTo demonstrate automated patching, i've used the recommended patch cluster. You can gather it at sunsolve.sun.com.SUNWjet:The Jumpstart Enterprise Toolkit. You can gather it at hereSUNWjass The Solaris Security Toolkit. You can get it at here

Posted by Joerg Moellenkamp in Solaris at 17:22

## **Less Known Solaris features: Jumpstart Enterprise Toolkit - Part 5: The toolkit**

Now we get to the tool that will be the centerpoint of the whole tutorial: The Jumpstart Enterprise Toolkit. This tool was started as an internal tool by some project engineers to make their job easier. And with the time a really mighty tool was born.

The basic idea behind JETThe Jumpstart Enterprise Toolkit (JET) was designed as an add-on to the normal Jumpstart process. It solves two major challenges:

Jumpstart works exceptionally well to install the operating system and to do configurations in conjunction to the

operating system, but misses some easy to configure features to execute further configurations of higher-level configuration besides of having hooks for finish scripts.

Some people find it a little bit complex to configure. While i strongly believe, this isn't the case, i have to admit, thats hard to remeber the commands when this is not your everyday task

With both issues in mind, some Sun Engineers start to develop the Jumpstart Enterprise Toolkit. Basically it's a bunch of clever shell scripts.

The Jumpstart Enterprise Toolkit configures the Jumpstart Server accordingly to the configuration in a template.

Furthermore it provides a framework for application specific modules to install and configure them.

When using JET you don't have to hassle around with all this files explained before, making the installation easier.

Additional features of JETJET isn't just a config generator for Jumpstart. You can configure JET in a way to fully customize your system. Much of this capabilities comes from a link called S99jumpstart or the respective SMF service. This script has an important role: It executes further actions on the system on following boots. For example, the mirroring of harddisks isn't done by integrated functionalities of Solaris Jumpstart. Those would make this action unavailable on Solaris 8 for example.. It's done by a script located on the server and made available by NFS, but executed on the client by using the S99jumpstart script. This concept makes the JET a very flexible tool.

Posted by Joerg Moellenkamp in Solaris at 17:05

## **Less Known Solaris features: Jumpstart Enterprise Toolkit - Part 4: Jumpstart FLASH**

Sometimes you don't to do a new install of a system. You just want to clone a system. For example think about a webserver farm. Let's assume you have thirty of them. You've configured one and now you want to distribute this config to all of your system. You've tuned the system extensivly, you changed configurations throughout all components. And you don't want to do this 29 times again.

Full Flash ArchivesSolaris Jumpstart knows a special mode of operation for this task: It's called Jumpstart FLASH. The trick of Jumpstart flash is quite easy. At first a normal Jumpstart install and the FLASH install are identical. But when it comes to the installation of the Don't install the packages one by one. Instead jumpstart flash unpacks a archive of a running system on a new system. This archive is called FLASH archive. Technically speaking it's not much more than cpio archive of a runing system.

Differential Flash ArchivesThere is an interesting mode of operation for flash archives. You can create differential flash archives. Let's assume you created a basic flash archive and installed all your systems with it: your webserver, your mailserver, your database server. Most parts of the system are identical. Just a few additional binaries and configuration files differentiate your server from each other.

Let's assume you want to create flash archives from all systems. Of course you could create a full flash archive for each system, but this would be waste of disk space. The differential flash archive creation works relatively simple. It compares the content of a flash archive with the actual state of an installed system and just archives the changed parts. The next time you want to install the system, you use both archives. At first the full archive will be installed on the system, after this you use one or more differential flash archives to complete your installation.

old  
new  
Action

exists not  
exists  
File is included in archive

exists  
exists but different  
The file from the new state is included in archive

exists  
exists not

File will be deleted, when the diff archive is used on a server

flar creation is just a big wrapper around cpio, thus it's possible to some nifty tricks with it. The current states of the system doesn't have be the active one, and the old states doesn't have to be flar archives.

It's possible to compare an old boot environment and the actual boot environment from Liveupgrade to generate a differential flash archive. This differential can be used to update other servers. You could even compare a remote system via NFS, when don't squash root. I know this has some security implications, but hey ... you should limit the access for such stunts to your admin networks and you can deactivate it afterwards.

Challenges of Jumpstart Flash for System RecoveryFlash was designed with the task of system cloning in mind. So it removes the identity of the system after the installation by using the sysidunconfig command. The need for such a stepat system cloning is obvious: One part of the systems identity is the networking configuration. You can't clone the network configuration as TCP/IP hate duplicate adresses

sysunconfig deletes the entire configuration, that makes the installation an unique instance of Solaris:

- saves a copy of /etc/hosts and substitute it with a default one.
- removes any NFS mount from /etc/vfstab
- deletes NIS, NIS+,LDAP and DNS name service configuration
- removes the interface configuration of alle configured interfaces.
- removes the root password
- removes /etc/sysidcfg
- removes /etc/defaultrouter
- removes /etc/inet/netmasks
- regenerates the ssh-keys
- sets the timezones in /etc/timezone to PST8PDT

Albeit it's not designed for system recovery, there is a trick you can use to recover the removed information. The knowledge about the removed part is important for the trick, thus i've included a list of them in this tutorial. You will find a script at the end of this tutorial.

Posted by Joerg Moellenkamp in Solaris at 16:51

### **Less Known Solaris features: Jumpstart Enterprise Toolkit - Part 3: Files controlling the installation**

I will not describe the exact process of native Jumpstart in it's multitude of configuration files you have to modify, as the Jumpstart Enterprise Toolkit will do this job for us, but it's important to know some of the important internals of Jumpstart. I will describe the most important files. There are varius others like /etc/ethers or the dhcp server configuration, but as you don't touch them manually even with the native configuration of Jumpstart i won't describe them here.

The rules fileThe first important file for the automatic installation is the rules file. This files associates system with a installation profile. This enables you to define rules that for a certain class of machines a certain installation profile should be used, e.g. a profile for small, medium and large machines.

```
# rule keywords and rule values      begin script      profile      finish script
# -----
```

```
hostname aramaki setup webserver completion
```

```
any - - genericprofile
```

-The first rule can be divided in it's part like this: When the hostname of the new server is aramaki, start the script begin on the client before starting the installation. For the installation use the profile file webserver. After the installation execute the script completion.

The second line is a catch-all condition. The file is used top down and the process of matching a system to a profile stops at the first match. Thus an installation for aramaki would reach the second line. This line can be traslated like this. For any other host, use the profile genericprofile. There is no begin or finish script.

You can't use the the rules file directly. The Jumpstart server provides a script to do an syntax check on te \verb=rules=. When the file is correct, the script adds it get's renamed to rules.ok. You don't have to check this file manually, JET will to this job for you.

The profile fileThe profile file controls what we install on the system and how we partition the disks for the installation.

```
# profile keywords      profile values
# -----
install_type          initial_install
system_type           standalone
partitioning          default
filesystem             any 512 swap
cluster               SUNWCprog
package               SUNWman delete
cluster               SUNWCacc
```

You can have a multitude of profiles in your system. A profile for system with large disks, a profile for system with small disks, a profile with a selection of packages customized for a webserver, a profile customized for a developer workstation. The jumpstart framework will choose the correct one on the basis of the rules.ok

The profile is capable to control almost any important parameter for the installation on the disk. You can define packages for installation or the configuration of your partitions.

The sysidcfg fileInstalling the packages on the system isn't enough to get the operating system up an running. You have to give the system a unique identity. Parameters for this identity are:

- configuration of network interfaces
- locales
- initial root password
- time zone
- etc.

You can type in such information manually, but that wouldn't be a hands-off installation. The installation process of solaris has a solution for this problem. You can set all the parameters in a file called \verb=sysidcfg=. The solaris installer will use this file to configure the system accordingly. keyboard=US-English

```
timezone=US/Central
timeserver=timehost1
terminal=ibm-pc
service_profile=limited_net
```

```
name_service=NIS {domain_name=marquee.central.example.com
                  name_server=nmsvr2(172.25.112.3)}
```

```
nfs4_domain=example.com
```

```
root_password=URFUni9But: Whenever some essential information is missing, the installer will go interactive and ask for the missing information. This obviously is against our objective of an automated installation.
```

Posted by Joerg Moellenkamp in Solaris at 16:32

## **Less Known Solaris features: Jumpstart Enterprise Toolkit - Part 2: Basic Jumpstart**

One basic problem of automated installation is the provisioning of the operating system on the bare metal. Without OS any further configuration is obviously without any relevance. Since several releases Solaris provides a feature called Jumpstart. The basic idea of Jumpstart is the installation of a system by using the network.

The Jumpstart mechanism for PXE based x86The process of jumpstarting is relatively simple.

- The system comes up. It tries to gather informations about the network
  - item It's own IP, netmask and network address
  - item the boot server

It connects to the boot server via tftp, loads the network boot programm (in this case pxegruband the menu.lst.)The pxegrub boots the Solaris environment on the client. The menu.lst contains the locations of important sources:config server

- location of the sysid\_config file
- location of the installation media

location of the boot environmentThe mini root starts. From now on,you have a Solaris Kernel running on your systems.It gathers the rules.ok from the installation server.It tries to find a matching profile based on the rules in the rules.ok file.

Based on the jumstart profile it formats the disks, it mounts all filesystems like on the final system relative to the directory \verb=/a=

The installer starts to install the packages relatively to the directory /a  
After all, it writes the boot loader to the disks  
Now you can reboot the system.

The Jumpstart process on SPARC is similar, but it uses other protocols like RARP and BOOTP to gather the initial configuration at the network boot.

Jumpstart ServerAlbeit the Jumpstart Server is an important part in the process of installing the operating system on the server, it's a passive part. The Jumpstart process itself ist executed on the client, not on the server. The server is only nescessary to provide information and the installation media. It uses existing protocols for this task as RARP or DHCP for sending configurations to the client about the network or NFS and HTTP for providing access to the installation files.

Further DevelopmentAt the beginning Jumpstart was just able to do the installation. Recent versions include functionality to automatically create boot mirrors (Recent in Solaris context means Solaris 9 and up )

Posted by Joerg Moellenkamp in Solaris at 16:16

### **Less Known Solaris features: Jumpstart Enterprise Toolkit - Part 1: Introduction**

As I wrote at several occassions in my tutorials, Solaris was designed with the enterprise customer in mind. It's quite common, that a company has dozens of servers and you have to put your customized operating system on all of them. A special managment tool, some special admin tools, a certain configuration of the volume manager.

The problem: Let's assume, you have 50 system, all of them will have a slightly different configuration. This is normal, when you install your systems manually. Humans aren't machines, and human work introduces variances in every instance of work. While this is nice for furniture or apparel, this can lead to long searches for needles in haystacks.

On some occasions i had strange problems with systems: One system worked just fine, another made problems. Both were installed with a cookbook. At the end we've used something similar to BART. We've compared both system and found slight differences. On the problematic system, the admin made an typo. I was curious and looked after the working system. Even this system wasn't exactly like demanded by the cookbook.

So: How can you prevent the subtle variances between your systems? The most obvious way is the automatisation of the installation. Let do the computer, what a computer can do at best: Repetitive tasks with the same quality over and over again. Jumpstart and Jet have several ways and means to do an automatic installation, thus making the live of the admin much easier after you've done the initial setup. And in this tutorial I want to show you that this setup is really easy.

Posted by Joerg Moellenkamp in Solaris at 16:11

### **The Caffeine Test**

The problem: I've just had three cups of East Frisian tea ...

Posted by Joerg Moellenkamp in Fundsache at 14:49

### **Springer migriert in Richtung Apple**

Der Axel-Springer-Verlag laesst in einer Pressemitteilung verlauten:  
"Apple steht für Kreativität, Innovation, Ästhetik und Kompetenz und ist damit der ideale Partner für Axel Springer", sagte Dr. Mathias Döpfner, Vorstandsvorsitzender der Axel Springer AG. "Die Umstellung auf Apple ist für uns nicht nur ein technologischer Fortschritt, sondern wichtiger Beschleuniger der kulturellen Modernisierung im Unternehmen."Nun, wenn ich mir die Bild-Zeitung(und die Variante mit mehr Buchstaben ... auch "Welt" genannt) so angucke, koennte dafuer mehr notwendig sein, als der Erwerb von Systemen aus Cupertino ...

(Danke an Local fuer den Hinweis)

Posted by Joerg Moellenkamp in General at 12:07

### **787 problems**

I don't think, that i will fly with the 787 anytime soon ... even in case, this baby will fly next year. There are new delays in the programs. As far as the usual suspects report, a mechanic at Alenia was able to destroy a complete barrel of the fuselage by using wrong fasteners. The damage was so severe, that they wait for a new center fuselage barrel.

Okay, this can happen, but this rises some questions for me as an outsider to this business : How do Boeing expect that an ordinary flight mechanic can repair the 787 on the apron, when a catering truck bumps into the plane or an uncautious loader rams its vehicle into lower deck? When even the people at the factory are unable to repair a damage? And to think it the other way round: Will they throw away such fuselage parts in the future or will they repair this to save some money (Boeing is a public traded company and i think Fight Club was right with it's presentation of the internal logic of companies). The damage must have been severe enough that Boeing examines it complete schedule as Dreamliner Four plays an important role in the flight test schedule. And this just by the mistake of a single mechanic.

Additionally: This damage happended at a part of the plane that was problematic for Boeing in the past as reported for example at Design News: Boeing Stiffens 787 Wing Box Spars: More Schedule Delays Possible.

Some strange: AirportBusiness.com reports in "787 partner halts work briefly after FAA audit":The South Carolina plant that assembles the mid-fuselage of Boeing's 787 Dreamliner halted production for 24 hours beginning Monday night after a Federal Aviation Administration (FAA) audit found lax manufacturing procedures that could result in damage to the aircraft sections.

[...]

He said it wasn't surprising the FAA found issues, given the high number of workers there who "never worked on airplanes before.

[...]

"In the beginning ... they were basically looking the other way to get the plane built," said the contractor. "Now people are paying attention to procedures and everything is being done by the letter of the law."Sorry that i can't resist ... but: The test flight pilots should wear parachutes not only on first flight

Well, i'm not in the aircraft business and i'm just a interested outsider , but at the moment my confidence as an upcoming consumer in this aircraft is waning at best ...

Posted by Joerg Moellenkamp in Aviation at 08:57

Thursday, July 3. 2008

### **A nice toy: Wordle**

A nice toy to create word clouds from websites. For this example i used my blog:

Posted by Joerg Moellenkamp in Fundsache at 19:37

### **links for 2008-07-03**

Platform Independent FMA for sun4v: July 2008 update  
(tags: FMA solaris sun)

Count von Count - Wikipedia, the free encyclopedia  
essential knowledge  
(tags: sesamestreet muppets)

List of Sesame Street characters - Wikipedia, the free encyclopedia  
essential knowledge  
(tags: sesamestreet)

Sun Microsystems and Intel Break Million-Messages-per Second Barrier for Thomson Reuters Market Data System:  
Financial News - Yahoo! Finance  
(tags: finance solaris)

Posted by del.icio.us in del.icio.us at 13:36

### **IBM acquires PSI**

There was a small company called Plattform Solution. This company developed a Mainframe clone based on Itanium procs. And big IBM sued the small PSI. This could be a story of hate, legal dodges and the game of suing and counter suing. But the story ended differently: The big IBM simply simply swallowed the small PSI as reported by Computerworld for example.

Well ... where leave this the mainframe market. IBM took it's nearest competito out of the game. This can lead to two developments in the future. You might think (and some industry analysts indeed do) that IBM will use the technology of PSI to build a new class of cheap mainframes. My personal opinion to that: Dream on and pigs might fly. IBM would canibalize it's high margin mainframe business with such an offer and it needs the high margins to refinance the further development of the mainframe technology for a shrinking group of customers. And additionally: I heard on several occasions, that IBM uses laptops internally with mainframe emulation for test purposes. It would be easy for the large IBM to productize this emulation if they really want to do that. They don't need to buy a company to do that (and more logical targets would be their own pSeries or xSeries and not an a somewhat strange Itanium server)

The more probable possibility is in my opinon a different one: They simply want to take out a competitor to protect their market share. And this would leave this acqusition as an interesting target for antitrust inquiries. But thats only my 2 cents, time will tell.

Posted by Joerg Moellenkamp in The IT Business at 10:52

### **Intel finally admits it ...**

The free lunch of the ever increasing per-core performance will come to an end. Ars Technica writes in Intel: an expensive many-core future is ahead of us: Intel has bad news for software developers. It's been hinted at already, but now the company has stated explicitly: it's not enough for software developers to be targeting dual, quad, or eight cores. No, the future holds tens, hundreds, or thousands of cores, and developers are going to have to bite the bullet and write programs that will scale to such systems. With their article they refer to a blog post on the Research@Sun blog: Unwelcome Advice. So ... please no more questions if CMT is the right solution. It's the only solution when even "throw-tremendous-amount-of-money-to-frequency-increases" Intel doesn't think that the other way is future-proof.

Posted by Joerg Moellenkamp in The IT Business at 09:59

### **Adam Leventhal about flash storage**

Adam Leventhal wrote an excellent article about the advantages of flash based storage and how ZFS circumvents some of the disadvantages (like write latency). Really a must read: Communications of the ACM: Flash Storage Memory. Additionally it explains some of the inner working of the L2 ARC.

(via: Robert Milkowski: RAM->SSD->DISK + ZFS)

Posted by Joerg Moellenkamp in Solaris at 06:59

Wednesday, July 2, 2008

## **Are blades always the densest way to deploy servers?**

Many people think about blades as the best way to put many small servers in a small amount of space. Often this is the case, but not in every case. Last weekend i did some research for a project and found out something interesting in the offering of our competitor HP.

Let's assume you have the requirement to use a quad processor blade. So you need one of the full height blades of HP the two disks of the blade aren't sufficient for you. You need more of them. Our beloved competitor has a blade shelf for this. It's called HP StorageWorks SB40c storage blade. It's a half height blade.

Okay, let's calculate a little bit. You could put 16 half height blades and 8 full height blades into a 10 rack units c7000 blade enclosure. Okay, we need some bays for the storage. You would assume that you can put 5 full height blades into the system and 5 storage blades in the free remaining bays (leaving one free) into the system, but this isn't the case. The storage blades have to be adjacent to the blade. More important, you can put it only in the lower bay. A half-height SB40c storage blade must be installed within the same partition as its partner server blade. If the SB40c is partnered with a full-height server blade, the SB40c must be installed in the lower bay adjacent to the full-height server blade. The HP BladeSystem Enclosure Tech Brief states: You could put a server into the free bay above the storage bay but there is an interesting gem in the documentation. The quick spec document states at the HP website states:NOTE: The lower tape or storage blade cannot be removed without first removing the upper half height blade. The problem, as far as i understand the datasheet of the SB40c it isn't just a bunch of cabling. It has an own RAID-Controller( HP Smart Array P400 with 256MB Battery Backed-Write Cache). It's an active component and there is some likelihood that active components will fail, especially when they contain batteries.

Let's assume you have used the slot above the Storageblade for another system. Let's assume you have to service the storage blade. You have to shut down the perfectly running system on the blade above the storage blade, as you have to remove it first. Doesn't sound reasonable for enterprise usage. This leads me to the conclusion that you use the c7000 enclosure only for 4 full height quad socket blades with 4 storage blades and nothing else.

Okay, that implies that you need 10 rack units (the size of the c7000 enclosure) to implement 4 quad socket systems with 3-8 disks with the HP blades. Now let's look at the problem from a different perspective: The Sun Fire X4450 is a quad socket server with up to 8 harddisks (with more PCIe-slots than a blade). The system takes rack units in you rack. So you could place 5 of them at 10 Rack units. Sometimes the denser solution is at a place where you don't expect it.

Posted by Joerg Moellenkamp in Sun at 19:39

## **Half a Terabyte in a X4600M2**

When i started to work at Sun in 2001, Sun introduced the SunFire 15000. A beast with with up to 574 Gigabytes of RAM. A system with a size of a deep 19" rack. We were able to address this vast amount of memory in a single operating system image. Today DataRam announced a 3rd party (read: unrecommended, unqualified and unsupported by Sun) RAM kit for the X4600 M2 with 8 GB per module. This leads to 512 Gigabyte in a single system with four rack units. That's technological progress.

Posted by Joerg Moellenkamp in Sun at 16:59

## **The market share of Apple notebooks**

I have a personal test for the market share of the Mac OS X platform. I like to sit in the first wagon of the train. To get into the on-board restaurant i have to go to half the train. On this way i count the types of notebooks. 4 years ago, a Mac was a really rare sighting. On my train trip back from Berlin to Hamburg today, i've counted 13 notebooks, 7 of them were from Apple. I was somehow undecided about the move to the Intel architecture, but it seems that this move had paid off for the people at the Infinite Loop.

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

Posted by Joerg Moellenkamp in Apple at 16:51

Tuesday, July 1. 2008

### **Trailer for 007: A quantum of solace**

At last a first trailer for the next Bond:

Posted by Joerg Moellenkamp in General at 22:05

### **links for 2008-07-01**

Spreeblick - Pan: Gegendarstellung oder wie die Bertelsmann-Stiftung Deutschland regieren laesst  
(tags: rap bertelsmann)

Produkte > Remote Buddy > Unterstuetzte Hardware // IOSPIRIT - fueling creative minds worldwide ..  
(tags: mac remotecontrol)

Posted by del.icio.us in del.icio.us at 13:35

### **Out now: MacOS X - 10.5.4**

Apple published an update for Mac OS X yesterday: 10.5.4. This update isn't as large as the last one, but it still contains some important bugfixes. You will find the combo updater and normal updater at the usual locations.

Posted by Joerg Moellenkamp in Apple at 13:31

### **Storagemojo about "The Hitz report"**

Robin from storagemojo wrote an interesting comment about the The Hitz report (Dave Hitz's declaration in the Sun/NetApp lawsuit):The NetApp/Sun patent battle continues. I don't see how NetApp can win this, given the Supreme Court's Teleflex decision, which makes prior art a question that can be appealed all the way to the Supreme Court. But the company is doggedly pursuing the battle, and Dave Hitz's recent declaration - which he hoped would remain private - has been unsealed.

Posted by Joerg Moellenkamp in The IT Business at 10:44