

Friday, August 29. 2008

Final approach ...

Posted by Joerg Moellenkamp in Aviation at 20:58

links for 2008-08-29

Big bucks, bigger hassle: Overbooking is less fun than ever - Upgrade: Travel Better

(tags: aviation businessstravel)

DasErste.de - [Panorama] - Egon Bahr: NATO-Beitritt Georgiens waere "abenteuerlich"

(tags: politics geopolitics)

Neurology
Protective effects of NSAIDs on the development of Alzheimer disease -- Vlad et al. 70 (19): 1672 --
Taking Ibuprofen may keep Alzheimer at bay ...

SPIEGEL Wissen :: MIT DER BOEING ZUR PARTY - Artikel

(tags: music insider)

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

Presentation about Rock

To end all questions and rumours about Rock ... this is the presentation held by my colleagues at the recent HotChips conference. Rock is alive and very healthy.

(click here to view)

Posted by Joerg Moellenkamp in Sun at 13:12

Toulouse revisited

I'm just waiting for the aircraft to Toulouse. I'm looking forward to the better weather, the nice landscape and visiting friends there ... I don't know if i find time to blog there but i start to like the concept behind the word "vacation"

Posted by Joerg Moellenkamp in Travel at 09:37

Thursday, August 28. 2008

A380-900

They don't have their own A380 aircrafts but after Air France/KLM there is a second airline interested in a even bigger A380 as Mr. Gallois disclosed in a newspaper interview. The german Lufthansa wants this plane, too. The A380-900 would be able to carry 963 passengers.

Posted by Joerg Moellenkamp in Aviation at 18:06

links for 2008-08-28

101 Atheist Quotes - The Atheist Blogger

(tags: religion interesting humour science theology skeptic)

TSA: Traveling with Film

(tags: travel security photography film)

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

New multifabric NEM for Sun Blade 6000

You may have noticed, that the Sun Blade Servers don't have Ethernet or FC on board. They just have 4 PCI-E connections to the outside world (2 to the NEM slots, 2 to the PEM slots). Sun announced an really interesting new NEM.

24 ports in total, 10 Gigabit-Ethernet Ports, 10 10-Gigabit-Ethernet-Ports and 4 SAS Ports for outside connections to SAS drive enclosures (albeit not supported in the initial release). You will find further specification on the website for the Sun Blade 6000 10GbE Multi-Fabric Network Express Module By using this module every blade has 1 GBE and one 10GBE port and you can put up to two modules in a Blade 6000. So you have a redundant production network connectivity, redundant admin network connectivity and still two free PEM slots for Infiniband or FC.

Posted by Joerg Moellenkamp in Sun at 07:27

Wednesday, August 27, 2008

MAID, ZFS and some further thoughts ...

I had an interesting discussion with a blog reader and customer via icq (oh wonderful new world of business). It started with a question about the J4500 and it ended in the the discussion, if MAID is really a viable way for storing data of if it's just one of this hypes working their way through the storage business every year.

MAID? What is MAID? MAIDs are Massive Arrays of Idle Disk. The idea is to save power by switching of disks in the case they aren't in use. You can buy such components from companies like Copan.
At first, i don't really think this technology justify companies, that do MAID for a living. It's too easy to integrate this feature in other system. At the end MAID is just hard disk power management in action. When a disk isn't used, switch it off. Just in the case you need data from them, switch them on. As it would be senseless to switch off a disk after two seconds just to reactivate it 2 seconds later (or it wouldn't switch the disk off at all in the case you choosed a longer interval) the concept of cache drives was introduced to the concept MAID. This drives cache the most often requested blocks. In opposition to the MAID drives, this drives are active at any time. As these cache drives should satisfy a large amount of disk requestsm, the MAID drives are not activated, the power management of the hard disk can keep them idle. At the end, this is all.

Well ... then i started to think about it and came to an interesting conclusion. Every OpenSolaris user has the components to build it's own MAID. Let's assume you have an array with 48 disk, the Sun Fire X4540

There are three important functions in Opensolaris that helps to implement a MAID like storage array:the L2 ARCthe separated ZILthe copy-on-write behaviour of ZFSLet's assume the follwing configuration. You take the CF card of the X4540 to boot the system. You take 4 disks to build a 4 TB L2ARC, you take the 2 disk to use them for the seperated ZIL thus. A well primed L2ARC should satisfy most of the requests from the cache. The separated ZIL deflects the writes from the data pool until the ZIL is written to the pool. At this moment the COW behaviour kicks in. The data is never overwritten at it's original location. New blocks are written on a new position. Let's assume the pool is constructed out of a RAID0 of RAID1 the writes should be concentrated to two disks. Thus the standard power management of the hard disk should kick in.

Other systems could use this MAID by Opensolaris functions: CIFS, NFS, iSCSI, FC Target and later iSER and SAS target. But there is one problem at the moment: ZFS is too intelligent. It tries to level the load on the disk by directing the write to the least loaded vdev. At this moment we need a real concat that writes data at the first free block of a disks and just jumps to the next block in the case no free space is on the actual disks (Okay, i assume it's RFE time).

Besides of the last component, everything is already implemented in ZFS and Solaris (and for the last problem i even have an idea for a quick hack: simple and plain just-in-time addition of new disks when the old disks are 95% full). This is the reason why i really think that companies that do MAID for a living are indeed living on borrowed time.

But the discussion ended with a different thought: I have my problems with the MAID concept. I'm now 15 years in this business. And my experience tells me, that electronic devices (especially hard disks) will die in the moment you switch them on ... not while they running. Now MAID implements a mechanism that deliberately switches electronic devices on and off, on and off and so on.

I'm not really convinced about the long time reliability of such a concept. But perhaps i'm too old fashioned with my standpoint that data is only safe when you wrote them to two different tapes written with two different tape drives and stored at two different locations.

Posted by Joerg Moellenkamp in Solaris at 20:46

Fangorn found a new home

There is a new resident in my home office. I'm jumpstarting fangorn at the moment. fangorn is a Sun Ultra 60 workstation (dual-proc UltraSPARC II system with 450 MHz on each proc and 1,5 GB memory). The Ultra60 line had it's GA in 1998. This was shortly after my first contact with Sun systems. I purchased it for a few euros from a colleague to rescue it from his cellar and to have a multi-proc SPARC system for some testing. I think, fangorn is a good name for

such an venerable system.

Posted by Joerg Moellenkamp in Sun at 18:40

Fujitsu develops SPARC64 processor with 8 cores

Sun and Fujitsu have started to deliver systems with the 4 core SPARC64 VII. But Fujitsu already talked about the next generation at the HotChips conference: An 8 core processor in development and they hope to talk about this a little more at the HotChips next year. 8 core SPARC64 would lead to a 1024 threads M9000 ... 4 Teraflops in a single operating system image.

Posted by Joerg Moellenkamp in The IT Business at 17:25

Tuesday, August 26. 2008

More news about the Spanair crash

Relatively short after the crash there were rumours about deployed thrust reversers at one of the engines of the Spanair aircraft and shortly after this you were able to find a photo (login needed) of this engine at the usual discussion areas. Today two spanish newspaper reported according to Aviationweek and the german newssite sueddeutsche.de (based on dpa) that the thrust reverser of one engine were found in the deployed position.

Albeit this would explain some of the reported reactions of the plane, it's much to early to draw any conclusions. The question is now: Were the reversers in this position before the crash or were they moved in this position by the crash. We have to wait for the final accident report ...

PS: Nevertheless this evoked memories to the LaudaAir 004 crash in me: Mid-air deployment of the reversers in 24,700 feet. Nobody in the aircraft survived the crash.

Posted by Joerg Moellenkamp in Aviation at 22:22

Some new fragments for the Kerberos Tutorial

I've uploaded another early pdf of the tutorial. I've added some text and the first to figures. Kerberos is just one of the topics where you should explain only the configuration or you have to explain it almost completly. You can't explain just a few things and still hope your explanation makes sense.

Posted by Joerg Moellenkamp in Solaris at 21:28

Privates Datawarehouse

Kris zeigt in Mein privates Datawarehouse - Sparen mit MySQL wie man mit der CSV-Export-Funktion des Sparkassen-Homebankings, Mysql und einigen Gedanken sich ein eigenes Data Warehouse bauen kann, um die Geldsenken im eigenen Budget zu finden.

Posted by Joerg Moellenkamp in Fundsache at 18:56

Gute Frage ...

Reichen 4 GB fuer eine Seele? Hängt vielleicht von den darin befindlichen Abgründen ab. Und wie gut lässt sich eine Seele komprimieren? Und wie oft muesste man Diffs sichern?

Posted by Joerg Moellenkamp in Braindump at 14:03

links for 2008-08-26

Worktimer / gersh.no - Gerhard Slettens blog

(tags: tools tool osx macosx freeware development productivity)

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

Dell offers Solaris in web-shop

In the Webshop of Dell you can choose Solaris x86 as an option. This is an example from the Dell PowerEdge R905:

But it costs \$231 more than the subscription at Sun

Posted by Joerg Moellenkamp in Solaris at 13:00

iSCSI boot for Solaris approved

The ARC has approved the PSARC/2008/427 iSCSI Boot case. This project is to enable Solaris to boot off iSCSI luns via regular network adapters. Different approaches, iBFT/OBP, are adopted to implement this feature on x86/sparc platforms. For the readers unaware of iBFT: This is a part of the ACPI 3.0 specification to enable a kernel to gather additional data needed to boot via iSCSI.

Posted by Joerg Moellenkamp in Solaris at 08:44

Is here the NY downtown apartment of Mr. and Mrs. Beast?

Posted by Joerg Moellenkamp in Photographie at 06:27

Monday, August 25. 2008

Hello, my name is Joerg and i´m an introvert!

Jonathan Rauch writes in *Caring for Your Introvert*: Do you know someone who needs hours alone every day? Who loves quiet conversations about feelings or ideas, and can give a dynamite presentation to a big audience, but seems awkward in groups and maladroit at small talk? Who has to be dragged to parties and then needs the rest of the day to recuperate? Who growls or scowls or grunts or winces when accosted with pleasantries by people who are just trying to be nice? I´m feeling observed. I can do a lessig-style presentation in front of 200 people without any problems but in a bar i´m really quiet, i hate senseless smalltalking like hell. A text catching my thoughts at the moment.

Posted by Joerg Moellenkamp in *Braindump* at 19:31

Herzlichen Glückwunsch!

An dieser Stelle möchte ich meine Glückwünsche an meine Schwester Solveig und meinen Bruder Jens richten. Solveig hat heute ihren Bachelor in Frühpädagogik mit hervorragender Benotung geschafft. Jens hat am Donnerstag sein Diplom als Wirtschaftsingenieur mit ebenfalls sehr guter Benotung erlangt . Euer grosser Bruder ist stolz auf euch

Posted by Joerg Moellenkamp in *General* at 19:02

Insights to Nehalem

Hothardware.com reports about some internals of Intels upcoming Nehalem architecture. Sounds interesting but i´m sure this turbo mode stuff will lead to many single thread performance benchmarks without practical meaning. Sound a little bit like this benchmark where a well known other chip company deactivated all cores except one, but not the respective cache. Thus they were able to execture large portions of the benchmark in the processor cache. Downclocking all cores except one and upclocking this one sounds like a similiar "benchmark-special".

Posted by Joerg Moellenkamp in *The IT Business* at 15:10

Poor mans Fault Tolerance

I should cook tea the whole day. I had an idea what we could do with all this threads in a 4 socket Victoria Falls system while waiting for the water to boil: Let´s build a poor mans fault tolerant system out of it. It´s an unrefined idea ...

Let´s assume you have to control something: For example a valve in a chemical plant. It´s an important valve, thus you can´t afford a fault (for example by cosmic rays flipping a bit). Thus it´s a common method to compute something multiple times and compare the result. Two computations isn´t enough. When the results differ, you can´t tell what result is the right one. Or better: What´s the result with the highes probability to be correct. The practice suggests to compute it three times and to compare it afterwards. The result with at least two votes of three in this quorum wins.

You could implement every decision making instance in an LDOM. Thus you could even implement different operating system patch levels in each of the systems.

And this fits with the four procs of a fully blown Victoria Falls system . Three procs for computations and one procs for the comparator. To ensure that every process runs on the same core it was tested, you could bind it to a processor.

This would reduce the thread count to 64 effective processors, but you have implemented a poor-mans fault-tolerant system (poor man because of no hardware lockstepping, it would be just application lockstepping) and well, you have more than enough hardware threads in a Victoria Falls system ...

I will think a little bit more about it in the next few days....

Posted by Joerg Moellenkamp in *Sun* at 14:09

links for 2008-08-25

IT-Unternehmen: Erfolgreich mit Autisten - Wirtschaft - SPIEGEL ONLINE - Nachrichten

(tags: autismus asperger)

Posted by del.icio.us at 13:33

Sunday, August 24. 2008

Solaris 7: 1998-2008

On August 15th the support for Solaris 7 has ceased (besides of special agreements) . This is 10 years after the introduction of Solaris 7 in 1998.

Posted by Joerg Moellenkamp in Sun at 07:11

Saturday, August 23. 2008

First fragments of the LKSF-Kerberos tutorial

I started this afternoon to write the Kerberos tutorial for the LKSF series. I've just put the first fragments in the public and will update them regularly. You can find them here. I will integrate them to the LKSF book as soon i've finalized it, but this will take really long time. I didn't even started describing the Solaris implementation so far. Kerberos is such a vast topic ...

Posted by Joerg Moellenkamp in Sun at 21:52

Jahresbericht 2007 der Bundestelle fuer Flugunfalluntersuchung

Wer sich dafür interessiert, was so im Dunstkreis der Bundestelle für Flugunfalluntersuchung so an Unfällen und Störungen passiert ist, kann sich mal den Jahresbericht 2007 durchlesen. Ist aber vielleicht keine Lektüre für Leute mit Flugangst.

Posted by Joerg Moellenkamp in Aviation at 20:51

links for 2008-08-23

Johanna Rothman - What Lifecycle? Selecting the Right Model for Your Project

(tags: toread work project)

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

Friday, August 22. 2008

No thunderstorms

I was really relieved about the fact, that weather forecasting is a rather unprecise science ... no thunderstorms, just turbulences and a hole in the air when the aircraft went through the clouds.

Posted by Joerg Moellenkamp in Business Travel at 18:39

You have viewed too much Battlestar Galactica episodes ...

... when you read "Cylon Tea" instead of "Ceylon Tea".

Posted by Joerg Moellenkamp in Braindump at 15:00

Brooklyn Bridge

Posted by Joerg Moellenkamp in Photographie at 14:53

Thunderstorms

The weather forecast reports thunderstorms in the afternoon for Hamburg . Sh... will be an interesting approach on the flight back from Dresden ...

Posted by Joerg Moellenkamp in Business Travel at 14:30

Spanair crash

Journalists, especially those working for TV news programs are a little bit weird. Every few days they repeat Discovery Channel documentations about aircraft crash analysis (UA811, TWA800, EI-AI Flight 1862, Swiss-Air 111) here in Germany - for example N24 - about the difficult and problematic post crash analysis. Sometime it takes years to find the real reason of an aircraft crash. Nevertheless journalists bash Spanair for giving no further informations about the causes for the crash. They should know better, at least when they look they own news programm. Nobody can tell more than speculative assumptions after just 2 days.

The bad thing: In an interview, a woman blamed Spanair for giving her less informations than the news. The problem: The news in the news are just rumours.

And then the Air Accident report will give an exact answer to the crash cause in a few months, no TV programm will report about it. Well, they will report about it, if Discovery Channel makes a documentation about it ... telling how hard it was to find the reason ... again ...

Posted by Joerg Moellenkamp in Aviation at 14:26

links for 2008-08-22

RIA Novosti - Verteidigung und Sicherheit - Russland weist georgischen Einsatz von Raketenwerfern gegen Zivilisten nach

(tags: russia georgia war)

Posted by del.icio.us in del.icio.us at 13:02

New implementation for NFS over RDMA

With the integration of PSARC 2007/347 NFS/RDMA Solaris got a new implementation for NFS over RDMA: This integration brings the NFS/RDMA Solaris Implementation up to date with the latest IETF Drafts. Specifically, it removes support for the initial NFS/RDMA implementation (version 0) and adds support for rpcrdma [1] and nfsdirect [2] (version 1).

Posted by Joerg Moellenkamp in Solaris at 12:11

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

Thursday, August 21. 2008

links for 2008-08-21

Guide To Painting Computer Case

(tags: tutorial painting howto hardware diy design cool computer)

Column - The State of ZFS - Peter Baer Galvin's Blog

(tags: zfs solaris sun)

Posted by del.icio.us in del.icio.us at 13:32

Wednesday, August 20, 2008

Mindmapping on Mac

Stefan Domanske of domanske.de pointed to a little interesting tool. Maybe you are aware of the technique of mindmapping for gathering ideas. I like this concept, but so far i preferred to use pencil and paper for this task as there was no really focussed and simple tool for it, at least not for the operating system from Cupertino. In the article "Mindmapping am Mac" Stefan talked about Mindnode. It's a tool for MacOS 10.5 and it just does Mindmapping with a really clean interface. It's near to the experience of using the oldfashioned paper because of the simple and clean interface. I really like this tool.

Mindnode is available for free on the website of the project, albeit there will be an Pro version in the future for purchase.

Posted by Joerg Moellenkamp in Apple at 22:23

Militaerisches Wissen

Da sage nochmal man könne mit militärischem Wissen nichts in der Privatwirtschaft werden. Es gibt jetzt von AXE ein Nervengift ... sorry, ein Deo das aus zwei Deodorantsprays kombiniert wird. Sprüht man sich mit beidem ein, erhält man ein drittes Deo. Beim Militär nennt man das einen binären Kampfstoff. Die Aufrüstung hat jetzt also auch chemische Kampfführung im Rahmen der morgentlichen Fahrten mit der S-Bahn erreicht.

Posted by Joerg Moellenkamp in Braindump at 18:27

links for 2008-08-20

At JFK Airport, Denying Basic Rights Is Just Another Day at the Office | Rights and Liberties | AlterNet

(tags: politics culture article)

Posted by del.icio.us at 13:32

Central Park

Posted by Joerg Moellenkamp in Photographie at 11:08

Atom vs. Opteron and an unlikely winner ...

Interestingly the most power-efficient desktop platform isn't the purpose-build Atom CPU from Intel. It's an AMD Opteron. Tom's Hardware found this in tests (i don't like this news outlet but sometimes there are some good articles there) A good example for "the system is key": Even when you CPU runs at two watts, a CPU alone doesn't compute a single command. You need a power-efficient chipset as well. And the chipset is the achilles heel of Atom at the moment.

PS: That's good news, i was afraid of being forced to buy Intel for my next homeserver

Posted by Joerg Moellenkamp in The IT Business at 07:45

Tuesday, August 19, 2008

New Xeon server from Sun: X2250 and X4250

They are somewhat similar, but at the end both represent really different concepts of building a server: Sun announced today the Sun Fire X2250 and the Sun Fire X4250. Both systems are systems in the 2-socket server class, but besides this fact the system are really different.

Sun Fire X2250 Server

The Sun Fire X2250 is a 2 socket XEON system in one rack unit. It can provide room up to 2 disks. It was primarily developed for HPC deployments as a cheap compute node or for highly budget sensitive customers. By the way: It provides an ILOM interface, it's not equipped with the ELOM.

Sun Fire X4250 Server

The Sun Fire X4250 is a two rack unit, dual socket system. The additional rack unit was used in an efficient manner. You can plug up to 16 SAS harddisks into this chassis. At this moment this leads to a total capacity of almost 2.3 TB raw disk capacity. So it's the Xeon based brother to the Opteron based X4240. The website specifies it's an ELOM system, but docs.sun.com already contains the documentation for the X4250 with ILOM.

Posted by Joerg Moellenkamp in Sun at 21:28

Less Known Solaris features: Jumpstart Enterprise Toolkit

I forgot to post the conclusion and the table of content when i published the JET tutorial some weeks ago. Okay ... here they are. This is a real long tutorial. 16 parts in total - the longest tutorial so far:

- Part 1: Introduction
- Part 2: Basic Jumpstart
- Part 3: Files controlling the installation
- Part 4: Jumpstart FLASH
- Part 5: The Toolkit
- Part 6: Prerequisites
- Part 7: Installing the JET Server
- Part 8: Preparing the first installation
- Part 9: Basic installation of a new system
- Part 10: A polished basic installation
- Part 11: Mirrored boot disks
- Part 12: Automatic hardening with SST
- Part 13: Deep dive to post installation
- Part 14: Using flash archives
- Part 15: Using flash archives for recovery
- Part 16: Conclusion

I hope this tutorial is useful for you!

Posted by Joerg Moellenkamp in Solaris at 18:05

Less Known Solaris features: Jumpstart Enterprise Toolkit - Part 16: Conclusion

JET and Jumpstart are incredible powerful tools to ease the installation of your solaris systems. You can install your systems via network, configure and customize them. With tools like Jumpstart FLASH it's easy to clone systems for large computing cluster.

Even when you have only a few systems, it's a good practice to install you systems by a Jumpstart server. A bare metal recovery of your system is much easier, you can use a common template to setup your system and beside the hard disk space the jumpstart server needs only resources at the moment of the installation.

Do you want to learn more?

Documentation

Solaris 10 5/08 Installation Guide: Solaris Flash Archives (Creation and Installation)

Misc.

wikis.sun.com: JET Homepage

jet.maui.co.uk: JET resources

Posted by Joerg Moellenkamp in Solaris at 17:51

Server porn - today: Sun Fire X4450

There is a nice piece of server porn at silicon.de: The dissection of a X4450 in fifty photos. A nice peek into the chassis of the only available Tier1 2RU quadsocket system.

Posted by Joerg Moellenkamp in Sun at 15:06

Unsere politische Kaste auf Reisen

Da wird ein Trip nach Kalifornien mit einem Hauch von Dienstreise getarnt und einige Mitglieder unserer Politischen Kaste sind immer noch nicht zufrieden: So schreibt der Spiegel Online im Artikel "Neger gesucht": Sozialdemokrat Krüger, Potsdam-West ist sein Ortsverein, sann auf Abhilfe. Vor der Stadtführung fuhr er den Generalkonsul nach dessen Erinnerung an: "Wir brauchen einen Neger, der den Rollstuhl schiebt." Einfach mal den Artikel auf der Zunge zergehen lassen und sich dann über nichts mehr wundern.

Posted by Joerg Moellenkamp in policy of ... at 11:38

I found my next car

Next year in February i have to buy or lease a new car. I think, i just found my new car. A Volvo. To be exact: The Volvo XC60. Not because it's a Volvo or because of the fact that it's a really cool car. No, because of a little function integrated into the car. It's called City Assist. The car has the capability to automatically brake the car in such situations like stop-and-go at a traffic light. As accidents in such situations hit me once or twice in the time i own a car, this would really reduce the costs for repairs and this feature would let my car insurance finally drop to a decent level. Sometimes technology is really a heaven-send...

Posted by Joerg Moellenkamp in Technology at 07:57

Tunnel vision

Posted by Joerg Moellenkamp in Photographie at 07:48

Monday, August 18. 2008

Less known Solaris Features book Version 0.28

I've integrated the CacheFS and the JET tutorials in the standard release of the LKSF book, thus now all links referring to the released version of the book should deliver the version 0.28 of the LKSF book (with the date of the 18.08.2008 on the titlepage). You can download it here.

Posted by Joerg Moellenkamp in Solaris at 23:31

Less known Solaris Features: CacheFS

Yet another tutorial finalized - this time about one of the really hidden features of Solaris - CacheFS. CacheFS is something similar to a caching proxy. But this proxy don't cache web page, it caches files from another filesystem. I divided the tutorial in 8 parts.

- Part 1: Introduction
- Part 2: Theory
- Part 3: A basic example
- Part 4: The Cache
- Part 5: On demand consistency checking
- Part 6: A usecase for CacheFS
- Part 7: The future of CacheFS
- Part 8: Conclusion

I hope, this tutorial gives you some new insights into Solaris. Have fun while trying out the stuff in the tutorial

Posted by Joerg Moellenkamp in Solaris at 19:29

Less known Solaris Features: CacheFS - Part 8: Conclusion

CacheFS is one of this features even some experienced admins aren't aware of. But as soon as they try it, most of them can't live without it. You should give it a try.

Do you want to learn more?

Manuals

System Administration Guide: Devices and File Systems - Using The CacheFS File System

man pages

docs.sun.com: cfsadmin

docs.sun.com: cachefsstat

docs.sun.com: mount_cachefs

Posted by Joerg Moellenkamp in Solaris at 18:52

Less known Solaris Features: CacheFS - Part 7: The future of CacheFS

CacheFS is one of this hidden gems in Solaris, it's even that hidden, that it's in sustaining mode since 1999. That means, that only bugfixes are made for this features but no new features found their way into this component since that day. In the recent days there was some discussion about the declaration of the End-of-Feature status for CacheFS which will lead to the announcement of the removal of CacheFS. After a few days of discussion the ARC decided in favour of the removal.

There are several problems with CacheFS (for example it's incompatibility with NFSv4) that mandates a major redesign of CacheFS, but CacheFS solves customer problems today. Thus i'm strongly disagree with the idea of removing a part of Solaris as long there is no equivalent feature (in this case: as long there isn't a new boot-persistent caching for NFS available in Solaris).

As i want to place a RFE to add an equivalent feature it would be very helpful to get some information from you:

Do you already use CacheFS?

In the case you use CacheFS at the moment ...

Is a removal of CacheFS a problem for you?

Would the removal hurt the deployment of Solaris at your site

Would you use an equivalent feature?

In the case you didn't used CacheFS so far ...

Is CacheFS a desirable feature for you after reading the tutorial?

Is CacheFS a feature that would help the deployment and adoption of Solaris at your site?

Please answer this questions by commenting this article. Thank you in advance!

Posted by Joerg Moellenkamp in Solaris at 18:24

Less known Solaris Features: CacheFS - Part 6: A usecase for CacheFS

Let's take an example: Let's assume you have 50 webserver that serve static content or that serves and executes .php files. You may have used scp or rsync to synchronize all this servers.

With CacheFS the workflow is a little bit different: You simply put all you files on a fileserver. But you have noticed, that there was a large load on this system. To reduce load to the fileserver you use your new knowledge to create a cache on every server. After this you saw a lot of requests for directory and file metadata for the documents directory. You know that this directory is provisioned with the changes only at 3 o'clock in the morning. Thus you write a little script, that checks for the file please_check_me and starts cfsadmin -s all on the client in the case of it's existence and it will check at the next access to the files if there's a newer version.

Posted by Joerg Moellenkamp in Solaris at 16:43

Less known Solaris Features: CacheFS - Part 5: On-demand consistency checking

Let's assume you share a filesystem with static content (for example the copy of a cdrom) or a filesystem that changes on a regular schedule (for example at midnight every day). So it would pose unnessesary load to network to check the consistency everytime again.

CacheFS knows a special mode of operation for such an situation. It's called on demand consistency checking. It does exactly what the name says: It checks only the consistency of files in the cache, when you tell the system to do it.

I will demonstrate this with an example:

With normal consistency checkLet's assume, we stay with the normal mode of operation from the example before. We

create a file on the fileserver.[root@theoden:/export/files]# date >> test_with_consistency_check

[root@theoden:/export/files]# cat test_with_consistency_check

Tue Aug 12 14:59:54 CEST 2008When we go to the NFS client and access the directory, this new file is visible

instantaneous. And when we access it, we see the content of the file.[root@gandalf:/files]# cat

test_with_consistency_check

Tue Aug 12 14:59:54 CEST 2008Now we go back to the server, and append additional data to the file:

[root@theoden:/export/files]# date >> test_with_consistency_check

[root@theoden:/export/files]# cat test_with_consistency_check

Tue Aug 12 14:59:54 CEST 2008

Tue Aug 12 15:00:11 CEST 2008And obviously, you will see this change on the client:[root@gandalf:/files]# cat

test_with_consistency_check

Tue Aug 12 14:59:54 CEST 2008

Tue Aug 12 15:00:11 CEST 2008

With on-demand consistency checkNow we unmount it, and remount it:

[root@gandalf:/files]# cd /

[root@gandalf:/]# umount /files

[root@gandalf:/]# mount -F cachefs -o backfstype=nfs,backpath=/var/cachefs/backpaths/files,

cachedir=/var/cachefs/caches/cache1,demandconst theoden:/export/files /files You may have noticed the demandconst option. This option changes everything.Let's assume you created another file on the NFS server:

[root@theoden:/export/files]# date >> test_with_ondemand_consistency_check

[root@theoden:/export/files]# cat test_with_ondemand_consistency_check

Tue Aug 12 15:00:57 CEST 2008 Back on the NFS client you will not even see this file:

```
[root@gandalf:/files]# ls
index.html                pcre_refcount.html
[...]
pcre_info.html           pcretest.html
pcre_maketables.html     test_with_consistency_check
```

You have to trigger a consistency check. This is quite easy.

```
[root@gandalf:/files]# cfsadmin -s all
```

Now you can see the file in the directory.

```
[root@gandalf:/files]# ls
index.html                pcre_study.html
[.]
pcre_info.html           test_with_consistency_check
pcre_maketables.html     test_with_ondemand_consistency_check
pcre_refcount.html
```

Okay, now we can look into the file.

```
[root@gandalf:/files]cat test_with_ondemand_consistency_check
Tue Aug 12 15:00:57 CEST 2008
```

Now we append a new line to the file on the server by executing the following commands on the NFS server

```
[root@theoden:/export/files]date >> test_with_ondemand_consistency_check
```

```
[root@theoden:/export/files]cat test_with_ondemand_consistency_check
```

```
Tue Aug 12 15:00:57 CEST 2008
```

```
Tue Aug 12 15:02:03 CEST 2008
```

When we check this file on our NFS client, we still see the cached version.

```
[root@gandalf:/files]cat test_with_ondemand_consistency_check
```

```
Tue Aug 12 15:00:57 CEST 2008
```

So let's trigger a consistency check:

```
[root@gandalf:/files]cfsadmin -s all
```

Now we can look into the file again, and you will see the new version of the file.

```
[root@gandalf:/files]cat test_with_ondemand_consistency_check
```

```
Tue Aug 12 15:00:57 CEST 2008
```

```
Tue Aug 12 15:02:03 CEST 2008
```

Okay, it's pretty obvious this isn't a feature for a filesystem that change in a constant and fast manner. But it's really useful for situations, where you have control over the changes. As long as a file is cached, the file server will see not a single access for such files. Thus such a file access doesn't add to the load of the server.

There is an important fact here: It doesn't tell CacheFS to check the files right at that moment. It just tells CacheFS to check it at the next access to the file. So you don't have an consistency check storm.

Posted by Joerg Moellenkamp in Solaris at 14:47

Less known Solaris Features: CacheFS - Part 4: The Cache

Okay, we have a working CacheFS mount but where and how is the stuff cached by the system. Let's have a look at the cache.

```
[root@gandalf:/var/cache/fs/cache1]# ls -l
total 6
drwxrwxrwx 5 root root 512 Aug 18 10:54 0000000000044e30
drwx----- 2 root root 512 Aug 11 08:11 lost+found
lrwxrwxrwx 1 root root 16 Aug 11 08:18 theoden:_export_files:_files -> 0000000000044e30
```

To ensure that caches using a single cache directory of the time aren't mixing up their data, they are divided at this place. At first a special directory is generated and at second a more speaking name is linked to this. It's pretty obvious how this name is generated. `=theoden:_export_files:_files` can be easily translated to `theoden:/export/files` mounted at `/files`

Let's assume we've used the cache for another filesystem (e.g. `/export/binaries` on `theoden` mounted to `/binaries`):

```
[root@gandalf:/var/cache/fs/cache1]# ls -l
total 10
drwxrwxrwx 5 root root 512 Aug 18 10:54 0000000000044e30
drwxrwxrwx 3 root root 512 Aug 18 11:18 0000000000044e41
drwx----- 2 root root 512 Aug 11 08:11 lost+found
lrwxrwxrwx 1 root root 16 Aug 18 11:18 theoden:_export_binaries:_binaries -> 0000000000044e41
lrwxrwxrwx 1 root root 16 Aug 11 08:18 theoden:_export_files:_files -> 0000000000044e30
```

With this mechanism, the caches are separated in their respective directories ... no mixing up.

When we dig down a little bit deeper to the directories, we will see an additional layer of directories. This is necessary to prevent a situation where a directory contains too much files and thus slows down.

```
[root@gandalf:/var/cache/fs/cache1/0000000000044e30/0000000000044e00]# ls -l
total 62
```

```
-rw-rw-rw- 1 root root 0 Aug 18 10:54 0000000000044e66
```

```
-rw-rw-rw- 1 root root 1683 Aug 11 08:24 0000000000044eaa
```

```
-rw-rw-rw- 1 root root 29417 Aug 11 08:22 0000000000044eba
```

When you look in on of this file, you will see, that those files are just a copy of the original

```
files:[root@gandalf:/var/cache/fs/cache1/0000000000044e30/0000000000044e00]# cat 0000000000044eaa
```

```
[...]
```

This page is part of the PCRE HTML documentation. It was generated automatically from the original man page. If there is any nonsense in it, please consult the man page, in case the conversion went wrong.

```
[...]
```

```
[root@gandalf:/var/cache/fs/cache1/0000000000044e30/0000000000044e00]#
```

At last, the structure of CacheFS for storing the cache is pretty simple. The metadata of the cache (what file is stored in which location) is stored in the root directory of the cache in some dot-files:[root@gandalf:/var/cache/fs/cache1]# [root@gandalf:/var/cache/fs/cache1]# ls -al

```
*
-rw----- 1 root root 48 Aug 11 08:11 .cfs_label
-rw----- 1 root root 48 Aug 11 08:11 .cfs_label.dup
-rwx----- 1 root root 0 Aug 11 08:11 .cfs_lock
-rw----- 1 root root 109879296 Aug 18 10:54 .cfs_resource
-rw-r--r-- 1 root root 4 Aug 11 08:11 .cfs_unmnt
-rw-r--r-- 1 root root 22 Aug 11 08:11 .nsr
```

```
[root@gandalf:/var/cache/fs/cache1]#
```

All this dot-files in conjunction contain all the metadata necessary to manage and control the cache.

Posted by Joerg Moellenkamp in Solaris at 13:44

Less known Solaris Features: CacheFS - Part 3: A basic example

Okay ... using CacheFS is really easy. Let's assume, you have an fileserver called theoden We use the directory /export/files as the directory shared by NFS. The client in our example is gandalf.

Preparations Let's create a NFS server at first. This is easy. Just share an directory on a Solaris Server. We login onto theoden and execute the following commands with root privileges.

```
[root@theoden:~]# mkdir /export/files
```

```
[root@theoden:~]# share -o rw /export/files
```

```
# share
```

```
- /export/files rw ""
```

Okay, of course it would be nice to have some files to play around in this directory. I will use some files of the Solaris Environment.[root@theoden:~]# cd /export/files

```
[root@theoden:/export/files]# cp -R /usr/share/doc/pcre/html/*
```

Let's do a quick test, if we can mount the directory:[root@gandalf:~]# mkdir /files

```
[root@gandalf:~]# mount theoden:/export/files /files
```

```
[root@gandalf:~]# unmount /files
```

Now you should be able to access the \verb=/export/files= directory on theoden by accessing \verb=/files= on gandalf. There should be no error messages.

Okay, at first we have to create the location for our caching directories. Let's assume we want to place our cache at /var/cache/fs/caches/cache1. At first we create the directories above the cache directory. You don't create the last part of the directory structure manually.[root@gandalf:~]# mkdir -p /var/cache/fs/caches

This directory will be the place where we store our caches for CacheFS. After this step we have to create the cache for the CacheFS.[root@gandalf:/files]# cfsadmin -c -o maxblocks=60,minblocks=40,threshblocks=50 /var/cache/fs/caches/cache1

The directory cache1 is created automatically by the command. In the case the directory already exists, the command will quit without creating the cache.

Additionally you have created the cache and you specified some basic parameters to control the behaviour of the cache. Citing the manpage of cfsadmin:

maxblocks: Maximum amount of storage space that CacheFS can use, expressed as a percentage of the total number of blocks in the front file system.

minblocks Minimum amount of storage space, expressed as a percentage of the total number of blocks in the front file system, that CacheFS is always allowed to use without limitation by its internal control mechanisms.

threshblocks A percentage of the total blocks in the front file system beyond which CacheFS cannot claim resources once its block usage has reached the level specified by minblocks. All this parameter can be tuned to prevent CacheFS from eating away all the storage available in a filesystem, a behaviour that was quite common to early versions of this feature.

```
Mounting a filesystem via CacheFS
We have to mount the original filesystem now.
[root@gandalf:/files]# mkdir -p /var/cache/dfs/backpaths/files
[root@gandalf:/files]# mount -o vers=3 theoden:/export/files /var/cache/dfs/backpaths/files
You may have noticed the parameter that sets the NFS version to 3. This is necessary, as CacheFS isn't supported with NFSv4. Thus you can only use it with NFSv3 and below. The reason of this limitation has its foundation in the different way NFSv4 handles inodes.
```

```
Okay, now we mount the cache filesystem at the old location:
[root@gandalf:/files]# mount -F cache/dfs -o backfstype=nfs,backpath=/var/cache/dfs/backpaths/files,cache/dfs=/var/cache/dfs/caches/cache1 theoden:/export/files /files
The options of the mount controls some basic parameters of the mount:
backfstype specifies what type of filesystem is proxied by the CacheFS filesystem
backpath specifies where this proxied filesystem is currently mounted
cache/dfs specifies the cache directory for this instance of the cache. Multiple CacheFS mounts can use the same cache.
```

```
From now on every access to the /files directory will be cached by CacheFS. Let's have a quick look into the /etc/mnttab. There are two important mounts for us:
[root@gandalf:/etc]# cat mnttab
[...]
theoden:/export/files /var/cache/dfs/backpaths/files nfs vers=3,xattr,dev=4f80001 1219049560
/var/cache/dfs/backpaths/files /files cache/dfs
backfstype=nfs,backpath=/var/cache/dfs/backpaths/files,cache/dfs=/var/cache/dfs/caches/cache1,dev=4fc0001 1219049688
The first mount is our back file system, it's a normal NFS mountpoint. But the second mount is a special one. This one is the consequence of the mount with the \verb=-F cache/dfs= option.
```

Statistics about the cache
While using it, you will see the cache structure at /var/cache/dfs/caches/cache1 filling up with files. I will explain some of the structure in the next section. But how efficient is this cache? Solaris provides an command to gather some statistics about the cache. With cache/dfsstat you print out data like hit rate inclusive the absolute number of cache hits and cache misses:

```
[root@gandalf:/files]# /usr/bin/cache/dfsstat
```

```
/files
  cache hit rate: 60% (3 hits, 2 misses)
  consistency checks: 7 (7 pass, 0 fail)
    modifies: 0
  garbage collection: 0
[root@gandalf:/files]#
```

Posted by Joerg Moellenkamp in Solaris at 12:59

Less known Solaris Features: CacheFS - Part 2: Theory

The history of CacheFS
Sun didn't introduced this feature for web servers. Long ago, admins didn't want to manage dozens of operating system installations. Instead of this they wanted to store all this data on a central file server (you know ... the network is the computer). Thus netbooting Solaris and SunOS was invented. But there was a problem: Swap via Network was a really bad idea that days (it was a bad idea in 10 MBit/s times and it's still a bad idea in 10 GBit/s times). Thus the diskless systems got a disk for a local swap. But there was another problem. All the users started to work at 9 o'clock ... they switched on their workstations ... and the load on the file server and the network got higher and higher. They had a local disk ... local installation again? No ... the central installation had its advantages. Thus the idea of CacheFS was born.

CacheFS is a really old feature of Solaris/SunOS. Its first implementation dates back into the year 1991. I really think you can call this feature matured

CacheFS in theory
The mechanism of CacheFS is pretty simple. As I told you before, CacheFS is somewhat similar to a caching web proxy. The CacheFS is a proxy to the original filesystem and caches files their way through CacheFS. The basic idea is to cache remote files locally on a harddisk, so you can deliver them without using the network when you access them the second time.

Of course the CacheFS has to handle changes to the original files. So CacheFS checks the metadata of the file before delivering the copy. If the metadata has changed, the CacheFS loads the original file from the server. When the metadata hasn't changed it delivers the copy from the cache.

The CacheFS isn't just usable for NFS, you could use it as well for caching optical media like CD or DVD.

Posted by Joerg Moellenkamp in Solaris at 10:21

Less known Solaris Features: CacheFS - Part 1: Introduction

There is a hidden gem in the Solaris Operating Environment called CacheFS, solving a task many admins solve with scripts.

Imagine the following situation. You have a central fileserver, and let's say a 40 webservers. All this webservers deliver static content, this content is stored on the harddisk of the fileserver. Later you recognize, that your fileserver is really loaded by the webservers. Harddisks are cheap, thus most admins will start to use an recursive rcp or an rsync to put a copy of the data to the webserver disks.

Well ... Solaris gives you a tool to solve this problem without scripting, without cron based jobs, just by using NFSv3 and this hidden gem: CacheFS. CacheFS is a really nifty tool. It does exactly what the name says. It's an filesystem, that caches data of an other filesystem. You have to think about it like a layered cake. You mount a CacheFS filesystem with an parameter that tells CacheFS to mount another one in the background.

Posted by Joerg Moellenkamp in Solaris at 10:14

Sunday, August 17, 2008

Nice tricks by merging photographs and video

I assume this demo reel of this project in computerized imaging shows the drain, where the next two or three bigger speed bumps and core multiplications will find their way into consumption. The projects website "Using Photographs to Enhance Videos of a Static Scene" describes:

We present a framework for automatically enhancing videos of a static scene using a few photographs of the same scene. For example, our system can transfer photographic qualities such as high resolution, high dynamic range and better lighting from the photographs to the video. At least you should look at the embedded video presenting the capabilities of this technology.

Posted by Joerg Moellenkamp in Technology at 22:35

links for 2008-08-17

NTSB Report AAR92-02: UNITED AIRLINES FLIGHT 811

(tags: aircraft accident)

Posted by del.icio.us in del.icio.us at 13:32

Georgien und die NATO

Man mag ja von Gerhard Schröder und seiner Nähe zu Putin halten, was man möchte ... aber bezüglich der Ferne des NATO-Beitritts von Georgien hat er recht.

Wir haben auf der einen Seite Russland, wir haben auf der anderen Seite Georgien. Ich denke, auf beiden Seiten ist viel Propaganda dabei. Das momentan eher das Pendel der Zuneigung in Richtung Georgien liegt, dürfte noch mit dem aus dem kalten Krieg stammenden Misstrauen gegenüber Russland zusammenhängen und das man oftmals einfach auf der Seite des Kleineren ist. Ich denke die Nachuntersuchung dieses Konfliktes wird ein differenziertes Bild zeigen.

Deutschland sollte sich ganz ehrlich den Austritt aus der NATO überlegen, wenn diese anfängt Staaten aufzunehmen, die ungeklärte Grenzkonflikte haben, nicht bereits über Jahre hinweg territorial gefestigt sind. Die Beistandspflicht könnte viel zu schnell in einem Schlamassel enden, das wir alle nicht wollen und an dessen Ende viel zu viele Tote stehen. Dafür kann man denke ich auch einen Staat hoffen opfern.

Entgegen aller Rhetorik, die da jetzt aus vielerlei westlichen Hauptstädten kommt, ist man denke ich doch ganz froh, das Paragraph 5 des Bündnisvertrages nicht zieht, weil Georgien noch nicht in der NATO war. Wie man sich aus dem Schlamassel sonst hätte rausreden sollen, ohne gegen Russland in den Krieg zu ziehen wäre sicherlich spannend gewesen.

Man sollte sich kein Land ins Bündnis holen, bei dem klar ist das der Bündnisfall irgendwann eintreten wird, möglicherweise dieses Land die NATO als Schutzschild verwendet, um die Grenzkonflikte im Rahmen von Konflikten niedriger Intensität weiterhin auszufeuchten.

Aus meiner Sicht sind Konstrukte wie Nationen, Regierungen und staatliche Behörden ohnehin dabei ein Konstrukt der Vergangenheit zu werden. Man muss sich nur mal ansehen, wie wenig Einfluss Regierungen wirklich auf wirtschaftliche Zusammenhänge haben und wieviele Steuerungsinstrumente bereits in den Händen überstaatlicher Organisationen liegt.

Damit ich hier nicht falsch verstanden werde: Der russisch-georgische Konflikt war unnötig, war wahrscheinlich auch eine Überreaktion auf russischer Seite. Nur sollte uns das zeigen, das wir nicht jedes Land in die NATO aufnehmen sollten im Interesse von Stabilität im globalen Massstab. Und wir sollten auch daran denken, das auch Russland Sicherheitsinteressen hat.

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

Posted by Joerg Moellenkamp in policy of ... at 10:39

Saturday, August 16. 2008

Strange combination

Some things are even somewhat too cynical even for such cynical minds like me. The Independent writes in NHS should not save patients' lives if it costs too much, says watchdog: The National Institute for Clinical Excellence (Nice) has said the natural impulse to go to the aid of individuals in trouble – as when vast resources are used to save a sailor lost at sea – should not apply to the NHS.[...]The disclosure follows last week's controversial decision by Nice to reject four new drugs for kidney cancer even though they have been shown to extend life by five to six months.[...]The judgement overrules advice from Nice's own citizen's council, which recommended two years ago that it should adopt a "rule of rescue" as an essential mark of a humane society. This happens when pure economic thinking supersedes some of the basic achievements of culture.

This ruling gets an even stronger bad aftertaste when you think about the EPIC FAIL of some parts of the NHS (National Health System) new computer system with at least two big companies that left this project (Accenture and Fujitsu). If they leave the table with the already paid money ... well ... this will be interesting legal disputes.

At the end: As BBC wrote in New blow for NHS e-record system The estimated final overall cost of computerising the NHS in England is currently £12.7bn. With £12.7bn just for a central health record you could give a lot of people a little bit more time by giving them the drugs they need.

Posted by Joerg Moellenkamp in policy of ... at 17:44

Award-winning restroom at Victoria Station

There is no award too improbable to be not awarded. I found this in the gentlemen's restroom at Victoria Station. Sorry for the blur, but I didn't want to use the flash in a restroom

Posted by Joerg Moellenkamp in Fundsache at 16:41

ZFS cache hierarchy controls

An interesting new feature had its flag day. With PSARC 2008/393 a feature called "zfs primarycache and secondarycache prop" found its way into Solaris for build 96. With these properties you can control, what ZFS caches in the ARC and the L2ARC. The introduction of the 'primarycache' property will allow administrators the ability to control what is cached in the ARC. The possible choices are: all, none, and metadata. If 'all' is chosen, then both user data and metadata are cached. If 'none' is chosen, then neither user data nor metadata is cached. If 'metadata' is chosen, then only metadata is cached. The default is 'all'.

The introduction of the 'secondarycache' property will allow administrators the ability to control what is cached in the L2ARC. The possible choices are the exact same as 'primarycache'. The default is also 'all'. This is interesting for situations, when you know that your application is caching data as well (thus it makes no sense to cache it twice) or you are aware of the fact that caching is just waste of time because of the access pattern.

You can find details about this feature in the PSARC discussion for 2008/393 and in the specs.

Posted by Joerg Moellenkamp in Solaris at 16:22

Friday, August 15. 2008

Moellenkamps in London

As there was the question of a reader in my mailbox about the trip to London: It was a private trip. I invited my mother (my father can't use the plane anymore), my brother Florian and my sisters Solveig and Taz for a day in London. Before you think: Does this guy have to much money? The fuel for driving to the airport in Bremen was almost as expensive than the tickets for 5 persons to London and the train tickets from Stansted to London were definitely more expensive. It's ridiclious: A train ticket from Hamburg to Berlin than a aircraft ticket from Bremen to Stansted. Strange.

For Florian and Taz this flight was their first flight. And they had much less fear of flying then i had on my first flight, despite of all the horror fairy tales i've told them before takeoff

Posted by Joerg Moellenkamp in Travel at 09:14

Thursday, August 14. 2008

Immutable files

Darren Moffat shows a nice trick in his blog. He explains how you can make a file in a ZFS filesystem even immutable for the root user. Hmm, i should integrate this to the LKSF book.

Posted by Joerg Moellenkamp in Solaris at 21:18

eLOM to iLOM

There are many people, who don't like the eLOM of our Intel based x86 servers and want the iLOM of the Opteron servers instead. The cure is near: The eLOM to iLOM migration guide already appeared at docs.sun.com

Posted by Joerg Moellenkamp in Sun at 21:04

Zuviel unterwegs

Gestern London, heute Rostock, morgen Fluch ... aeh ... Flug fuer Dresden buchen ... ich bin zuviel unterwegs ..

Posted by Joerg Moellenkamp in Business Travel at 19:14

Nachholen

Es ist irgendwie schon seltsam wenn man 2 Jahre im Leben eines Menschen, den man aus den Augen verloren hat, in 45 Minuten lesenderweis nachholt, weil dieser jetzt auch ein Blog fuehrt (von dem man erst vor kurzer Zeit weiss).

Posted by Joerg Moellenkamp in Blogosphere at 19:07

Tuesday, August 12. 2008

LKSF Sneak Preview: CacheFS

I'm writing a new tutorial for the Less Known Solaris Feature series. This time it's about CacheFS. CacheFS is one of the really hidden gems of Solaris. You can find it in the work-in-progress version of the LKSF book. It starts on page 74.

Posted by Joerg Moellenkamp in Solaris at 18:16

My Canon left me for a repair

Just left my Canon 40D at the repair shop. 140 to 300 Euro for an defect Compact Flash Card Slot. 140 Euro in the case they can solder in some new pins, 300 Euro in the case of a new slot logic board. Fsck. I will buy the biggest and fastest CF card available after this repair and this card will never leave the slot again

Posted by Joerg Moellenkamp in Photographie at 13:29

VMware bug revisited

For my non-german speaking readers a link to an english background article in the Register -Date bug kills VMWare systems:

Reg reader Duncan said VMWare's FAIL represented a "fantastic bug for a company trying to embed itself into the modern computing world." While another reader, Eric, said the "time bomb" contained in the update was causing a lot of panic among businesses.

"Customers were fuming this morning having planned downtime for weeks. VMware have alot of answering to do on this and no doubt share price will take a hit again," he said.

Posted by Joerg Moellenkamp in The IT Business at 12:57

Timebomb in VMware patch

Heise.de reports about a timebomb in the Update 2. The system denies the migration, creation and suspension of virtual machines with the message that the product has expired. That matches my opinion about VMware with all this big vendors introducing virtualisation solutions, but that's a different story Well, Heise writes that the suggest workaround ist to set a date in the past and to disable the time synchronisation. But as virtual clocks are just a rough guess of the actual time, this doesn't look as a viable workaround for me.

BTW: They have another problem. The support website seems to suffer the slashdot effect started by the link on the heise website.

Posted by Joerg Moellenkamp in The IT Business at 10:34

Monday, August 11. 2008

Interesting development at the USAF tanker proposal

According to a report at Aviationweek Boeing thinks about to leave the tanker rebidding process. The new proposal gives more credits to a bigger tanker and thus a 767 based tanker wouldn't be competitive to a A330 based tanker ... and as i wrote before. A 777 based tanker is no alternative for Boeing besides of the short timeframe for this proposal. Now, however, the Pentagon is pushing for the replacement tankers as soon as possible after multiple delays. It remains unclear if a no-bid position from Boeing would drag out the KC-X competition or it would compel the Pentagon to attempt a sole-source of the work to Northrop Grumman/EADS. This is an interesting development. Stopping the former decision with the help of the GOA just do leave the proposal process afterwards will not help to make friends for future proposals

Posted by Joerg Moellenkamp in Aviation at 22:16

Lokalpatriotismus: Dat is Friesland

Irgendwie bekomme ich das "Gülle!Gülle!Gülle!" nicht ausm Kopf

Danke an B-Tina fuer dieses grandiose Video

Posted by Joerg Moellenkamp in Fundsache at 20:40

Hoerempfehlung: Peter Gabriel et al. - Big Blue Ball

The solo studio albums of Peter Gabriel are outstanding. Without any question. And the live performances of Peter Gabriel are not from this world. But i have a somewhat divided relationship to some of it's other works. The "Passion" album for example: I asked other PG fans and it was the album with the "one song". There was only one really good title: A different drum. But not all of his side projects were that mixed: Ovo was excellent.

A few weeks ago, Peter Gabriel released "Big Blue Ball". This release are the results of a collaboration of many WOMAD (a world music label initiated by Mr. Gabriel) musicians round about 10 years ago. And my thoughts are somewhat mixed: I'm not that world music fan. I'm a PG fan. So not convinced about all the tracks. Nevertheless there are really good pieces on it. "Whole Thing"(more like in his US/So phase) for example. Or "Exit Through You" and "Burn you up, burn you down"(both more similar to the music of his Up phase). From the pieces without the master the track "Everything Comes from You" is my favourite.

But at the end: I can't wait for I/O - the next pure Peter Gabriel album.

Posted by Joerg Moellenkamp in Music at 19:35

Faked fireworks

There is some discussion about the faked fireworks sequence of the Olympic opening celebration. Was there really someone really, who believed that this was a real sequence? The flight of the chasing plane was much to stable for a real one and the light wasn't correct. That was pretty obvious.

Posted by Joerg Moellenkamp in General at 18:16

Best x86 server for Java - Sun Fire X4600M2

Good news from the benchmarking department - Sun Fire X4600 M2 Server Is Industry's Fastest x86 Java Server: Sun Microsystems, Inc. (NASDAQ: JAVA) today announced the new SPECjbb2005 world record score, which beats all other x86 systems on the market. The Sun Fire X4600 M2 server, equipped with eight Quad-Core AMD Opteron model 8360SE processors and running the Solaris 10 Operating System (OS), posted an x86 World Record score of 683,542 SPECjbb2005 bops (85,443 SPECjbb2005 bops/JVM). You will find further information at the X4600 Benchmarks page.

Matches my impression that the X4600 is one of the big hits from Andy and his teams.

Posted by Joerg Moellenkamp in Sun at 15:07

Sun Events

Ich habe mal die Liste der Termine in der linken Seitenleiste aufgefrischt. Ich bitte um freundliche Beachtung.

Posted by Joerg Moellenkamp in Sun at 14:55

Rouladen und Chipmultithreading

Was hat die Zubereitung von Rouladen mit Chip Multi Threading zu tun? Sehr viel ...

PS: Das Video ist auch schon beim Indiskretion Ehrensache Blog angekommen:

Andererseits: Ich mag Unternehmen, die genau das wagen. Einfach mal was probieren. Auch auf die Gefahr, dass es schief geht.

Und deshalb: Hut ab vor Sun Microsystems und ihrem Kochstudio.

Posted by Joerg Moellenkamp in Sun at 14:43

Sunday, August 10. 2008

links for 2008-08-10 [delicious.com]

2722267025_e2d7944766_o.png (PNG-Grafik, 750x1000 Pixel) - Skaliert (69%)

More Calvin and Jobs

(tags: jobs funny calvin apple parody)

2723090810_6501c93ae2_o.png (PNG-Grafik, 750x1000 Pixel) - Skaliert (69%)

Calvin and Jobs

(tags: calvin stevejobs satire marketing images jobs tech humour mac parody steve)

Die 100 Besten Schriften

(tags: webdesign typography)

Posted by del.icio.us in del.icio.us at 13:01

Cat content

As there was no "cat content" so far in this blog i fill that gap now. I'm proud to present "Homer" a.k.a Muh (because of his fur) - the Moellenkamp family cat and the most back-stabbing cat of the western hemisphere:

Posted by Joerg Moellenkamp in Photographie at 12:06

Kindergarten revisited

Wer einen Kindergarten in Reinkultur sehen moechte, moege sich bitte mal die Diskussion rund um blastwave auf heise.de angucken. Schily und Fefe prügeln sich da verbal wie die Kesselflicker auf einem niedrigen Niveau ("Lern du mal was ein Integer overflow ist", "Lern du mal programmieren") ... eben Kindergarten revisited.

Posted by Joerg Moellenkamp in General at 11:08

Saturday, August 9. 2008

Synchronized masses

I thought a little bit about the opening ceremony of the Olympic Games yesterday. Those scenes of 2008 people almost in synchronity ... it seems to be a german phenomenon to get remembered to one of our darkest ages by them.

Posted by Joerg Moellenkamp in Braindump at 15:10

links for 2008-08-09 [delicious.com]

Intuitor Insultingly Stupid Movie Physics

(tags: writing video science physics stupid reviews)

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

Friday, August 8, 2008

Kindergarten@Blastwave

Blastwave was essentially dead yesterday. As i'm not really online at the moment, the first news about this occurrence reached me by the twitter feed of Anfalas. Blastwave is one of the important repositories of pre-compiled opensource software in the Solaris community. From one moment to another, the repository was completely missing.

Later on it came apparent, that the repository was deliberately deleted by Dennis Clarke, the owner of the blastwave.org domain and the respective trademark because of an raging and unresolved dispute between him and other members of the freeware community, who fill the blastwave.org site with life by providing packages. At this moment it looks as a dispute between the already named Dennis Clarke and Phil Brown - one of the leading maintainers - as Anfalas writes in a german article on his blog. Well: At the moment blastwave.org is online again. So the biggest technical issues are resolved for the moment.

I will not write about an own position at this issue, because i'm not involved enough to have a well-informed opinion on this. I can't take party for Dennis as well as But will comment on a different problem that arises out of this dispute even for the non-involved observer.

When Sun talked about a way to adopt the Linux-style in some parts, i'm sure we don't have thought about this kindergarten style in sorting out disputes between the stakeholders of a community project. With participation comes responsibility. The more you participate at a community, the more responsibility is on your shoulders. When i close down c0t0d0s0.org without any warning, nobody would really care ... but blastwave.org is an important tool for the community. You can be in anger about something or somebody but you have to keep the people in mind, who a service and who are the source of the importance of a service.

Shutting down a service without warning, shutting down a service without discussing a contingency plan for the time without your involvement seems to be an highly unprofessional move. People may have their disputes, but by providing a service to a community you are responsible. Not only for the service, but for the way you solve your disputes as well. Of course you could withdraw your support in a community project at any time you want, but you have to give other community members the possibility to go forward without you, especially when the community has helped to grow your service by a vast amount of effort.

The Solaris freeware community at large can't afford such a kindergarten behaviour at a central point of it's infrastructure. Perhaps the non-availability of the service will lead to a solution, that's independent from a single site or a single person. So this reaction in anger of Dennis may lead to a waning importance of his blastwave.org project at the long term. The implementation of independent CSW-sites in conjunction with an own infrastructure to provide packages based on the pkg-get mechanism may be first signs of this development.

Posted by Joerg Moellenkamp in Solaris at 20:38

Proof reading

Sorry for writing nothing the last few days, but i'm occupied with the proof reading of the bachelor thesis of my little sister. Arguing about some phrases with her can be quite embarrassing ...

Posted by Joerg Moellenkamp in General at 20:31

Geopolitics

What is advanced geopolitics in practice? Sitting at the opening ceremony of of the Olympic Games while battles between your and georgian soldiers starts to rage in South Ossetia.

Posted by Joerg Moellenkamp in policy of ... at 19:23

The Peking Olympics has started

Whatever you think about the chinese people, the ignition of the olympic flame was a cool stunt. But to hoist the olympic flag by soldiers was a little bit strange.

Posted by Joerg Moellenkamp in Braindump at 18:32

Wednesday, August 6, 2008

Fujitsu-Siemens divorce rumour revisited

I don't have a subscription for the Wall Street Journal, but the free preview says enough about this Fujitsu-Siemens divorce rumour: Siemens Set to Pull Plug On Venture With Fujitsu: In a move that could set the stage for the sale or dismantling of a leading European maker of personal computers, Siemens AG has informed Fujitsu Ltd. that it wants out of their nine-year-old joint venture, people familiar with the matter say. This revisits this rumour reported before from other news outlets. The rumour made it even into a German news broadcasting service today.

Posted by Joerg Moellenkamp in The IT Business at 10:15

Tuesday, August 5. 2008

Solaris Cluster testbed

There is an interesting walkthrough for testing Solaris Cluster with iSCSI and Virtualbox at Building a Solaris Cluster Express cluster in a VirtualBox on OpenSolaris. Whoever wants to play with a mature cluster framework should give this a try.

Posted by Joerg Moellenkamp in Solaris at 16:01

Hoerempfehlung: MGMT - Oracular Spectacular

Auch wieder mal etwas aus der Kategorie "Nichts das man im Radio hoert". Oder wuerde man bei einem Dudelfunksender Lyrics folgenden Inhalts erwarten:

I'm feeling rough, I'm feeling raw, I'm in the prime of my life.
Let's make some music, make some money, find some models for wives.
I'll move to Paris, shoot some heroin, and fuck with the stars.
You man the island and the cocaine and the elegant cars.

Eher nicht, und obiger Text ist eben aus dem Lied "Time to Pretent" von MGNT. So entgeht dem ueblichen Musikhoeerer auf jeden Fall ein gutes Stueck Musik. Wirklich einordnen kann ich diese Musik nicht. Ich wuerde mal sagen elektronischer Independent mit einem massiven 80er Einschlag ... aber ich mag musikalische Schubladen eh nicht. Ich denke wir werden dennoch noch mehr von dieser Gruppe hoeren. Hoerempfehlungen sind Time to pretend, Kids und Electric Feel

Posted by Joerg Moellenkamp in Music at 14:12

links for 2008-08-05 [delicious.com]

Hardening your files with Reed-Solomon codes

(tags: tools sysadmin storage software opensource math linux error)

Posted by del.icio.us in del.icio.us at 13:01

What sportscar are you?

I'm a Porsche 911!

You have a classic style, but you're up-to-date with the latest technology. You're ambitious, competitive, and you love to win. Performance, precision, and prestige - you're one of the elite, and you know it

Take the Which Sports Car Are You? quiz.
(gefunden bei: Desideria)

Posted by Joerg Moellenkamp in Fundsache at 00:25

Stöckchen: FeedReader

Bettina ist zwar nach Stöckchenwurf verreist, aber ich will dann jetzt doch mal antworten:

1. Welchen FeedReader benutzt Du und seit wann?

Also ich nutze NetNewsWire. Hmm, ich denke mal schon so lange wie ich blogge. Also etwa vier Jahre nun. Ich kann mich eigentlich nicht erinnern, etwas anderes verwendet zu haben. Ausser in der mac-losen Zeit, da war es denn Google Reader. Es gibt meines Erachtens keinen mit NNW vergleichbaren Newsreader fuer Windows. Entweder fehlt

die Usability oder die Stabilität. Ich brauche beides.

2. Wieviele Feeds hast Du abonniert, wieviele davon sind Blogs (in %)?

Hmm, NNW sagt 301, davon sind wohl etwa 50% Blogs. Allerdings kann man das nicht so richtig zählen. Ich habe beispielsweise den blogs.sun.com-Feed abonniert. Dahinter stecken 5000 Blogs. Gottseidank sind die nicht alle so schwatzhaft wie mein eigenes Blog.

3. Wieviele Feeds liest Du täglich (wöchentlicher Schnitt)?

Alle 301. Bin Speedreader und versuche "latest News" immer auf 0 zu halten. Aber mal einen Tag nicht gelesen, schon ist man wieder auf 800 unread.

4. Dein meistgelesenster Feed/deutscher Blog-Feed?

International gesehen sicherlich der RSS-Feed von blogs.sun.com.

5. Welche Nachrichtenseiten hast Du abonniert (nenne max. 3)?

The Register

The Inquirer

sueddeutsche.de

6. Dein erstes und Dein letztes Blog im Reader?

Das erste Blog, das ich abonniert habe, war glaube ich das von Kris. Das letzte das ich dazugefügt habe ist jenes von B-Tina und dem Wuki.

7. Lesegewohnheit: alphabetisch, chronologisch, durcheinander?

FIFR ... first in, first read

8. Hast Du Blogs/Feeds, von denen Du nicht zugeben würdest, dass Du sie liest, im Reader?

Ja, habe ich. Nein, ich sag nicht welche. Allerdings habe ich andere Feeds subscribed, fuer die sich andere Leute vielleicht schämen würden. Ich sehe da aber viele Dinge ... aehm ... lockerer. Ich lese zum Beispiel als Einsicht in einen ganz anderen Beruf das Blog des Streetgirls ganz gerne.

Weitergabe des Stöckchens:

An Local, an Kris und an Frau D.

Posted by Joerg Moellenkamp in Blogosphere at 00:05

Monday, August 4. 2008

Zu ehrlich?

Ich glaube, ich bin zu ehrlich fuer diese Welt: Ich bin ohne Schmuggelware aus den USA gekommen, ich warte auch fuer einen 1mm Kratzer mit dem Einkaufswagen. Nur der Zugbegleiter (nebst einigen Mitreisenden) hat mich eben gerade wie ein Auto angeguckt, als ich für die Strecke zwischen Hamburg nach Oldenburg mit Umstieg in Bremen eine Fahrkarte kaufen wollte. Nebenbemerkung: Ich sass schon im Zug zwischen Bremen und Oldenburg.

Uns definiert so wenig als eigenstaendiges Lebenwesen. Integritaet beispielsweise. Ehrlichkeit ist ein Bestandteil der Integritaet. Vielleicht sollte man das nicht wegen 8 Euro aufgeben.

Posted by Joerg Moellenkamp in Braindump at 17:00

links for 2008-08-04 [delicious.com]

Anke Groener - Blog Archive - Grace is Gone
Noch ein Film fuer die Liste der zu guckenden Filme ...

Posted by del.icio.us in del.icio.us at 13:01

Good to be in Hamburg again

Good to be in Hamburg again. At the end it's the most beautiful town in the world

Posted by Joerg Moellenkamp in Braindump at 08:28

Sunday, August 3. 2008

Back from NYC and a first New York photoset at Flickr

I'm finally back in Hamburg. Was a good flight back to Frankfurt and further to Hamburg.

Additionally i've uploaded some unedited photos to a photo set at flickr.

You can view the photoset by clicking on the photo.

Posted by Joerg Moellenkamp in General at 20:43

Saturday, August 2. 2008

Brooklyn Bridge und Downtown New York

I promised a photo to Local. Okay, here it is:

I made this photo with a small compact camera using the Manhattan Bridge as a tripod.

Posted by Joerg Moellenkamp in Photographie at 20:51

Travellog NYC

25.07.2008

No, this isn't the first day of my vacation in NY. But a journey doesn't begin with the pushback at the airport. A journey begins with a decision, with an idea. And this journey began last year.

Normally I thought about travelling to NY end of February. But as you may remember my grandmother was hospitalized at that time and died not much later. I'm glad the trip was postponed at that time, as I weren't able to attend her funeral at that time.

So, somewhere in May or June the idea of visiting New York. But not as a part of a meeting in the US. This time as a real vacation. Taking a week off and doing recreation.

I've booked the flight on short notice end of June during the EM. A day before Germany lost in the European championship.

Believe it or not ... this is my first wide trip for a vacation. I do so much business travel over a year, that I found it more recreating to simply stay at home. But this time I wanted to see something new. The other problem was that my girlfriends didn't like to travel that far or didn't like to fly.

26.07.2008

09:00 MEST I can give you some tips for travelling. At least those tips worked for me. Tip 1: Book a flight in afternoon, but bring your luggage to the airport as early as possible. At the check-in: Take all your smile, be Prince or Princess Charming, ask for an exit row seat. Your legs will thank you for that.

14:00 MEST The journey starts. The plane will be pushbacked in a few minutes. Exit Row Window seat. But window is a great exaggeration in an A321. You sit at a large door with a window two or three inches wide.

16:30 MEST Wow, this idiotic second security screening for the transatlantic flights has ceased. I walked to the gate, but no security controls for transit passengers. Makes sense, I was already screened in Hamburg.

BTW: A big complaint regarding the airport in Frankfurt. Who had the incredible dumb idea to place the Business/FTL lounge at one end of the terminal and to place the aircraft to NYC at the other end. High speed walking to lounge, eat something, high speed walking back to the terminal. This was much better at the old location where the FTL lounge was just upstairs.

18:59 MEST We are now over the Atlantic Ocean.

The old trick still works. So Tip 2: Order a special meal and you get your meal first. There is only one problem. You can't choose your meal. There is no "Chicken/Pasta or Beef" question. When the "low calories meal" contains beef and you don't like beef you are out of luck.

And dear Economy purserette, was it really necessary to call my name out that loud with "Mr. Moellenkamp, did you ordered a special meal?". Head going red guaranteed.

BTW: Didn't miss a thing with the Pasta. Really looked like cooked to death. But nobody really expects gourmet food, even when they call their catering subsidiary "Sky Chefs" or "GateGourmet". The other meal was beef as well.

19:05 MEST or 13:05 EDSTOkay, i switched to New York local time. But why i'm still tired like being busy for the whole day. My inner clock has stil Hamburg time. We don't need pills for erectile dysfunction, we need a pill to adjust the inner clock. "Oh, it's 8 o'clock PM local time ... i have to take the pill with the 8 imprinted". That would really solve problems. Well, perhaps i will think different in 20 years

14:53 EDSTI really hope that the A350 XWB (extra wide body) has some more space for Joerg XWS (extra wide shoulders). It's annoying that all this people bump into my shoulders when they walk through the aisle.

And the amount of people walking by are the biggest disadvantage of exit row in a 747. Row 34 is located near the lavatories. Many people walk here after the meal. At the flight back from SFO to FRA last year in February i had even a bruise from one of these trolleys.

15:12 EDSTOh my goodness. My neighbour took off his shoes. I should try if ear plugs are effective in the nose as well. This guy should put a biohazard sticker on it's socks.

16:22 EDST Landfall. We reached the North American continent between Gander and Goose Bay.

19:00 EDSTI hope that the runway took no damage when the aircraft touched down. No it didn't touched down, it crashed down. Even the stewardess didn't expected such a noise, albeit she expected a rough landing as she admitted later. I knew that New York is located on an island, but i really thought it larger than an aircraft carrier vessel. That was really loud. The complete final approach was strange. Really fast and the aircraft banked a little bit to right. I assume, the pilot asked himself after the landing, if he paid the last rate for the insurance already

Besides of this (okay, and my neighbour Chief Stinking Socks), it was an almost perfect flight. Nice stewardesses. No turbulences.

20:10 EDSTImmigrations, baggage claim and customs in 15 to 20 minutes. That's fast ... sheer luck. Thank you to the one, who put a priority tag on my luggage by mistake. That made my day.

Exit Row 34 in the 747 has another advantage. You are in the first row of economy. Thus you can stay in front of the horde.

DHS guy unfriendly (well, unfriendly is the incorrect word, it was proper behaviour, but you can say at reply on a "Good afternoon" with words besides of "Right index finger, left index finger, look into the camera. Why do you come?"). But i assume you stop to say more after the 1000th visa a day.

The fifth visa in my passport. The fifth visit in the US.

20:25 EDSTLost the saved time while searching for an ATM and waiting for a taxi. But nevertheless ... it was really fast in comparison to my last trips.

21:42 EDSTWell ... the trip Taxi to the hotel Was a good ride ... just a small traffic jam. Friendly taxi driver.

So, i'm in my hotel, took already a first shower. You guys here in the US must be not that tall? I hate this low shower heads.

Hotel looks good so far. I'm in the Marrakesh Hotel. Simple, but clean rooms the usual ridiculous amount of chlorine in the water. Starbucks is left beside the hotel. So the supply of coffee and some of the breakfast standards is assured. The vacation can start.

21:50 EDSTFsch ... they don't have the collector mugs like in Germany (the most important sights and the name of the city in large letters). Only this "Facts" mugs. Bought some bottles of water of Ethos water. I won't make the error again of drinking water from the tap. My stomach isn't accomodated to that large amount of chlorine. Water with an ethos? Well ... i will think about that tomorrow morning.

21:51 EDSTI'm going to bed now.
I'm really tired.
I'm getting too old for this shit ...

27.07.2008

03:39 EDST

Okay, awake for an hour now. Welcome, jetlag, i didn't really missed you. But hey, only six hours time-offset ... could be worser. The time offset from San Francisco to Germany is a lot of harder. Nevertheless, i want a flat earth (carried by four elephants which stand on a big turtle). A round earth proved as ineffective. God (the higher entity of the non-existant kind) wasn't a (business) traveller

I will try to get some additional sleep now.

06:22 EDSTApropos elephants. It must have been a horde of elephants living in the room next to me. Awake now ... well ... looks it makes no sense to try to get an hour more sleep.

09:00 EDSTIt were two hours of additional sleep. Showered and ready for the City!

19:32 EDSTJust arrived back at the hotel. Was quite a walk today. From the 103th to the 53th. Had a nice day at the Museum of the Modern Art and jet another mug for my collection. I'm a proud owner of a MoMA mug

Grand Central Station is a breathtaking building. And that's just a a train station. The hall is huge. NYC, i'm already impressed. It here something not that large ?

22:30 EDST

Good dinner at an Indish restaurant right beside the hotel.

28.07.2008

7:30 EDST

The elephants are back. It's an art to close the door to the shower with such a noise without the according noise of shattering glass falling down.

But heck ... i slept until 07:20. I should try two big glasses of wine everytime i have to accomodate to a new time zone. Seems to help.

Okay, this day will start with Chinatown. Maybe Battery Park later on.

17:30 EDST

Hmm ... my travel guide told me, that Chinatown is the largest chinese communities in the world. I visited Chinatown today and i have one important question: I know that Chinese People are somewhat smaller than the usual North European guy, but where are that much people from China?

Afterwards a nice dinner in Little Italy. I really liked the French Toast. And the calories budget shouldn't be a problem with all this miles and miles of walking. They guy serving this meal got it's tip of the day, because i still didn't really understood the priciples of tipping (he got more than 30%). Later i've learned that you wait for the change and leave a and afterwards i've learned that wait for the change and leave an matching amount of money on the table as a tip.

By the way: Little Italy is pretty much Little China II. Pretty much of the streets designated as Little Italy on my map were filled with shops with neons made of chinese letters on the front.

Afterwards on the way to the Battery Park i walked along the site, were the Twin Towers stood. It's a really disturbing thought that there were the two highest building of New York. That must have been ridiclious forces at work to break down such buildings. The building leaved a giant hole in the earth. A cold feeling ran down my spine. Even more interesting: There is a little church besides the site and the church wasn't damaged at the crash down.

At this moments you think about the forces that protected this building, even when you are an atheist.

00:30

After a short stay in my hotel i had dinner at a restaurant at the Battery Park again. Was excellent. Good wine. Nice salad. Pretty expensive. But i assume you've paid the nice view from almost the southernmost point of Manhattan as well.

29.07.2008

17:00

Yet another point removed from my "Been there - done that" list: Sitting in the Central Park eating lunch and lying in the sun afterwards.

Did some shopping in the afternoon. At 80 degrees fahrenheit and the humidity here in NYC you wear off your t-shirts faster than expected.

Sitting in the park for eating chips, fruit salad with a Corona or two is nice. Especially when you can look onto the Hudson river. But i saw several strange things there. At "Gristedes", a supermarket in the vicinity of the hotel, you was able to buy real beer. Was nice to see, that you could buy Jever when you want to that far from home.

Even more stranger: They play soccer here. Somebody told me that they call this "pussy sports" here in the US but it were definitely mostly male people on the field.

30.07.2008

10:00

Another day for culture. Visited the American Museum of Natural History. Think i've walked at least the distance from the hotel to the museum again in the museum itself. Great museum and didn't even visited the special exhibits. Certainly a goosebumps moment was the oceanic part when you come into hall and see a giant blue whale at the ceiling.

15:30A multitude of new impressions. A multitude of things learned. When this rate stays at the same level my head will explode at the end of the week by new impressions.

16:00I can't be more stereotype-alike. Siting in the Central Park again, eating fruit salad and reading Wired I know it's a wannabee magazine but the gadget section is quite amusing.

17:00Again shopping. That's not interesting in itself. But why to Heck (and it's ruler Phil, the prince of insufficient light) did they named a shopping franchise "Banana Republic"?

19:00Just a remark to the US: I don't know how often i saw the "I'm pissed off my anger management class". Anger Management ... what a ridiclous thing. This "political correctness" thing is utter nonsense. You need a good culture of dispute. Anything oppressed will explode later on. Just with a more devastating effect.

01:34Just back from another long walk. Walked from Canal Street to Brooklyn over the Manhattan Bridge and back to SoHo via the Brooklyn Bridge. Was quite a walk again but it was really nice. Made some spectacular photos of the skyline of New York. It was really good to take this walk. Without it i would have missed a really spectacular and beautiful view to the downtown of new York.

31.07.2008

09:00Ready for the City ... the second last full day of this vacation. Time flies.

16:45You need three and a half hour to walk the Broadway from Battery to the 103rd street. That's round about 12 km when my measurement via USB cable is correct. It was more than 12km in my case because you should look on the numbers of the street after leaving a clothing store by another door than you entered it. You may detect earlier that you walk on the wrong direction on the wrong street because you lost the orientation while shopping in the store.

Now sitting on my bed and giving my feet some well deserved rest ... they carried me well through the city so far.

17:20Fsck wrong setting of the camera. Kept the setting from inside the museum (forget automatics in there. Way to dark.) Well ... no wonder ... my thoughts were somewhere else. Normally i check the settings everytime

18:00I assume, i've breaked the 40 km range this week at least (walked the distance from 103rd street to Columbus Square at least 4 or 5 time). I will add up all this distances with Google Maps when i'm back in the reach of internet.

I really thought about the question, if it was useful to buy a 7 day ticket when you walk any way. Nevertheless. I really think the only way to discover NY is by your own feet. Down in the subway you will see nothing. Walking through NY will give you an perspective about the city, the people, the crowds, the changing architecture and concepts of the buildings.

20:50It's a kind of Sex in the City. Do you remember "The secret of my success" scene, where Mr. Fox laid in the bed and his neighbour were fscking in their bed on the other side of the wall? Well ... the same here in the hotel. Even the wall to the next hotel room bumped like in the movie. Was fun to mimick Mr. Fox in this moment. The moaning of the woman was a little bit high for my taste. That wasn't moaning that was squeaking . And the boy has no condition. From door closing to switching on TV in a time you should invest into the foreplay alone;) But that's just my opinion.

01.08.2008

09:00The last full day. I'm totally overloaded with impressions. But there are some things left to view.

My back is aching, my feet are hurting. I've just asked for a Business Class Upgrade on Miles for my trip back to Frankfurt, as my back would punish me for such an endeavour in a Economy seat.

16:30

Phew ... my feet told me to go to hotel. Walked from Time Square to the United Nations, upwards to the Dag Hammerskjold Plaza. After that i walked to the Intrepid Museum at Pier 68. The Intrepid is an old aircraft carrier used as an aircraft and space museum. There was only a simple problem. I didn't know that they are rebuilding the Pier. No Intrepid there, no museum there either. So if you planed to visit the museum, check if it has reopened. As far as i was able to look at the pier, they already reworked much of the pier. So it couldn't be that long to the reopening.

A short walk to the Central Park and into the Metropolitan Museum of Arts. That museum is a must for everybody visiting NYC. It's a ride through the human history of art. I saw many of the objects in nature, i knew only from books. I was 6 hours in there, and i still think i missed something.

Whoever traveled to NYC without visiting the Metropolitan Museum of Arts missed an important part of New York.

21:00Oh no ... my neighbours again ... the good thing about the situation ... it won't take that long It's time for my earplugs again.

02.08.2007

07:00This is the day to say goodbye from New York. Was a trip full new things. It was worth effort. But now i want back home. When i look at my travel guide, i missed pretty much of the city despite of the city. Think this trip to New York will not be the last time ...

07:50Sh... all this stuff fitted better in the suitcase on the inbound journey to NYC. And i really bought nothing large. Just the minimum amount of stuff to get over a third of the days with fresh clothing.

10:30After importing the last CF card: 1557 Photos in one week. I need a bigger disk in the notebook. 500 GB would be nice

While importing i've read the actual issue of the Wall Street Journal (was the only non-yellow press newspaper available). One sentence made my day. In an article about exchange trades funds, the author wrote: "Many ETF are so overspecialized, that they make a pediatric endocrinologist like a family doctor who makes house calls".

10:45Now it's official, seven days since i've used the internet the last time (besides the small entry in my blog). So: I'm not an addict. Yeah! I'm so proud about this Enough waiting....

Posted by Joerg Moellenkamp in Travel at 17:31