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

Saturday, May 31. 2008

Masters at work

Wow, i think Pixar wants to know it with Wall-E ... Jonathan Ive (you know, the Apple Design Guy) helped with the design of Eve, the female robot in the movie, and a few minutes ago, i´ve read that Peter Gabriel worked on parts of soundtrack.. Stunning.

Posted by Joerg Moellenkamp in Movies at 22:13

Cooooool ....

Big Buck Bunny from Blender Foundation on Vimeo.

Posted by Joerg Moellenkamp in Fundsache at 21:38

Tempokontrollen in der Schweiz

Ich bin ja letztes Jahr ueber Zuerich zu einem Kunden nach Sueddeutschland geflogen. Ich habe hier gerade einen Beitrag auf N.24 ueber die Tempokontrollen in der Schweiz in der Schweiz gesehen. Irgendwie kann ich da nur sagen : Glueck gehabt. Und ich muss mich wohl bei dem bedanken, der mich ausdruecklich vor zu eiliger Fahrweise gewarnt hat... ich frage mich gerade ob die dort die Strafe mit der Tachonadelabweichung nach oben in Bogenminuten mal 50 Rappen errechnen

Posted by Joerg Moellenkamp in Business Travel at 11:33

Friday, May 30. 2008

About the FAA, Sun and bad news.

This is another artitcle of the kind i do not really understand: FAA: Sun box disk failure caused NOTAM database crash. What do you think, when you read this article at first: Sun can´t develop storage? Bad, bad Sun ? Well, the cursory observer might think that.

But at the end it´s a story about the stuff, why Sun develops stuff like ZFS. It´s just doesn´t told this way. It´s a story of undetected silent data corruption, transmission of the data corruption to the backup system, the problem of damages/power outages while updating the data on the harddisks: "What happened was the drive in an end-of-life Sun box failed in the middle of updating the information on the hard drive, so it screwed up the database," Davis said.andAs the technicians were working to fix the database, they decided to go to the backup system. As they did that, they soon realized they had written the error over to the backup system and had corrupted that system as well, Davis said.It sound like a perfect example why the industry needs a filesystem like ZFS. Harddisks fails, that is inevitable, and mostly they do not fail without a last scream. And that is independend of the brand of harddisks or the storage vendor.

So ... dear Computerworld, would you mind to use a less lurid headline next time?

Posted by Joerg Moellenkamp in Sun at 21:19

Twitter

At the moment Twitter is really annoying. It´s like sending SMS in the first minutes of the new year. But i see no rockets and drunken people on the street ...

Posted by Joerg Moellenkamp in Blogosphere at 18:32

Open HA Cluster completely open-sourced

With the last two million lines of code Sun reached the goal of open sourcing the Sun Cluster Framework (some months ahead of the goal). Now all parts of the suite are available under the CDDL license.

Posted by Joerg Moellenkamp in Sun at 13:37

Slideuments

Interesting article about the spreading nuisance to generate Powerpoint Presentations consisting out of several slides for topics easily summarizeable in a few sentences: Toyota chief: refrain from using PowerPoint. The Toyota Motor Corporation CEO Katsuaki Watanabe got to this topic while talking about cost reductions: Watanabe said that (in the good old days?) they used to use one piece of paper to make a clear point or proposal, or to summarize an issue, but now everything is in PowerPoint, he says, which uses many sheets of paper and expensive colors...but it's a waste. In my personal opinion this slideuments waste something more precious than paper or printer colours. They are wasting time!

Posted by Joerg Moellenkamp at 10:45

5 Million Downloads

Virtualbox was downloaded 5 million times in one and a half year. At the moment Virtualbox 1.6 gets downloaded 10000 times a day. Virtualbox isn´t an everyday usage product like Openoffice. So it´s quite a number for a product like a hypervisor.

Posted by Joerg Moellenkamp in Sun at 10:37

Thursday, May 29. 2008

Mission critical

You need a good explanation for the word "mission critical" for you manager and an example, why redundancy is a good thing? Okay ... i have one ... ISS toilet fails to suck. The toilet on the International Space Station is out of order and - to add insult to injury - it´s the only one. Toilets are absolutly mission critical ... well, you can´t open the airlock to take a walk behind the next tree

Posted by Joerg Moellenkamp at 11:11

Less known Solaris features: On passwords

The JET tutorial isn´t complete, but shorten the time a little bit for you, i wrote a tutorial about a small but nevertheless important topic - passwords:

Part 1: Introduction
Part 2: Using strong password hashing
Part 3: Using a password policy
Part 4: Conclusion

Posted by Joerg Moellenkamp in Solaris at 09:59

Less known Solaris features: On passwords - Part 4: Conclusion

These are some simple tricks to make your system more secure, just by ensuring that the keys to your server are well-choosen and not simple ones. But as i stated before there is something you should keep in mind: Don´t make the passwords too hard to remember.

Do you want to learn more?Documentation
docs.sun.com: man passwd(1)
docs.sun.com: Changing the Default Algorithm for Password Encryption

Misc. Links
Learning Solaris 10: Solaris Crypt : better password hashing algorithms

Posted by Joerg Moellenkamp at 08:17

Less known Solaris features: On passwords - Part 3: Using a password policy

User have the habit to break any security policy. At least as long you don't enforce it. One of the most annoying habit from the view of the security people is the tendency to choose weak passwords, the name of the boy or girl friend, the prefered brand of cars, birthdays ... you name it. This passwords are everything but secure. But you can configure Solaris to check the new passwords.

Specifing a password policy
There is a central file in Solaris controling the password policy. In /etc/default/passwd you define what requirements a password must fulfill before Solaris allows the user to set this password. Let´s have a look in the actual file of a standard solaris system. You have to log into your system as root. One important note for trying out this feature. You need to log into your system as a normal user in a different window.root can set any password without a check by the password policy thus it would look like that your configuration changes had no effect:# cat passwd
[... omitted CDDL header ...]
#
MAXWEEKS=
MINWEEKS=
PASSLENGTH=6
#NAMECHECK=NO
#HISTORY=0

#MINDIFF=3
#MINALPHA=2
#MINNONALPHA=1
#MINUPPER=0
#MINLOWER=0
#MAXREPEATS=0
#MINSPECIAL=0
#MINDIGIT=0
#WHITESPACE=YES
#DICTIONLIST=
#DICTIONDBDIR=/var/passwd
You enable the checks by uncommenting it and set a reasonable value to the line. When you enable all the checks, it´s actually harder to find a valid password than a non-valid one. Whenever thinking about a really hard password policy you should take into consideration, that people tend to make notes about their password when they can´t remember it. And a strong password under the keyboard is obviously less secure than a weak password in the head of the user.

Posted by Joerg Moellenkamp in Solaris at 07:47

Less known Solaris features: On passwords - Part 2: Using stronger password hashing

Many people are unaware of the fact, that only the first eight characters of a password are used in the default configuration of Solaris. Don´t believe it? Let´s try it.

Testing the relevant password length for standard cryptOkay, i´ve logged into my test machine and change my password:bash-3.2$ passwd jmoekamp
Enter existing login password: oldpassword
New Password: aa3456789
Re-enter new Password: aa3456789
passwd: password successfully changed for jmoekamp
bash-3.2$ Now let´s try a password that´s different at the ninth character by logging into the Solaris system from remote:mymac:~ joergmoellenkamp$ ssh jmoekamp@10.211.55.200
Password: aa3456780
Last login: Wed May 28 11:24:05 2008 from 10.211.55.2
Sun Microsystems Inc. SunOS 5.11 snv_84 January 2008I´ve told you ... only the first eight characters are relevant.

Stronger hash algorithmsBut it´s not that way, that Solaris can´t do better than that. It´s just the binary compatibility guarantee again. You can´t simply change the mechanism encrypting the password. There may be scripts that still need the old unix crypt variant. But in case you are sure, that you haven´t such an application you can change it, and it´s really simple to do.

When you look into the file /etc/security/crypt.conf you will find the additional modules for password encryption.# The algorithm name _unix_ is reserved.

1 crypt_bsdmd5.so.1
2a crypt_bsdbf.so.1
md5 crypt_sunmd5.so.1The hashing mechanisms are loaded as libraries in the so-called Solaris Pluggable Crypt Framework. It´s even possible to develop your own crypting mechanism in the case you don´t trust the implementations delivered by Sun.

Short
Algorithm
Description

unix
standard unix crypt

The standard unix crypt mechanism. It´s not loaded as a module, as it´s part of the libc

1
BSD alike,md5 based
The crypt_bsdmd5 module is a one-way password hashing module for use with crypt(3C) that uses the MD5 message hash algorithm. The output is compatible with md5crypt on BSD and Linux systems.

2a
BSD alike, blowfish based
The crypt_bsdbf module is a one-way password hashing module for use with crypt(3C) that uses the Blowfish cryptographic algorithm.

md5
Sun, md5 based
The \verb=crypt_sunmd5= module is a one-way password hashing module for use with \verb=crypt(3C)= that uses the MD5 message hash algorithm. This module is designed to make it difficult to crack passwords that use brute force attacks based on high speed MD5 implementations that use code inlining, unrolled loops, and table lookup

Each of the last three mechanisms support passwords with up to 255 characters. It´s important to know, that the different hashing algorithm can coexist in your password databases. The password hashing for a password will be changed when user change his or her password.

Changing the default hash mechanismLet´s use the md5 algorithm in our example. But before that, we should look into the actual \verb=/etc/shadow=# grep "jmoekamp" /etc/shadow
jmoekamp:nM2/fPrCTe3F6:14027::::::It´s simple to enable a different encryption algorithm for password. You have just to change a single line in /etc/security/policy.conf. To edit this file you have to login as
root:CRYPT_DEFAULT=md5Okay, now let´s change the password.

```
# passwd jmoekamp
New Password: aa1234567890
Re-enter new Password:  aa1234567890
passwd: password successfully changed for jmoekampWhen you look in the /etc/shadow for the user, you will see a slighly modified password field. It´s much longer and between the first and the second $ you seee the used encryption mechanism:# grep "jmoekamp" /etc/shadow
jmoekamp:$md5$vyy8.OVF$$FY4TWzuauRl4.VQNobqMY.:14027::::::Now let´s try the login:mymac:~ joergmoellenkamp$ ssh jmoekamp@10.211.55.200
Password: aa1234567890
Last login: Wed May 28 11:38:24 2008 from 10.211.55.2
Sun Microsystems Inc.   SunOS 5.11      snv_84  January 2008
$ exit
Connection to 10.211.55.200 closed.
mymac:~ joergmoellenkamp$ ssh jmoekamp@10.211.55.200
Password: aa1234567891
Password: aa1234567892
Password: aa1234567893
Permission denied (gssapi-keyex,gssapi-with-mic,publickey,keyboard-interactive).
mymac:~ joergmoellenkamp$You see, the correctness of the complete password is tested, not just the first 8 characters.
```

Posted by Joerg Moellenkamp in Solaris at 06:52

Less known Solaris features: On passwords - Part 1: Introduction

The password is the key to the system. When you know the username and the password, you can use the system. If not ... well ... go away. You can´t overemphasize the value of good passwords. There is something you can do as the admin at the system level to ensure such passwords: At first you can use stronger mechanisms to hash the passwords. And you can force the users of your system to use good password. This tutorial will explain both tasks.

Posted by Joerg Moellenkamp in Solaris at 06:49

Wednesday, May 28. 2008

Mac OS X 10.5.3 out now

There is a new update for MacOS X 10.5. You find the release notes for 10.5.3 in the Knowledge Base: About the Mac OS X 10.5.3 Update. And the release notes are looong ... Okay, you can download the combo update at Apples Download site. You just have to download 536 MB. In the case you already use 10.5.2 you can use the smaller Updater - just 420 MB

Posted by Joerg Moellenkamp in Apple at 22:21

Tuesday, May 27. 2008

Solaris Preso in MUC

200 Slides in 45 minutes. And i was only 4 minutes over the time at the end. I wasn´t aware of the fact, that i´m able to speak that fast, but at the end this was really exhausting.

Posted by Joerg Moellenkamp in Solaris at 17:48

Offlining a 3D server modell from the Sun homepage

Maybe you´ve noticed, that we have 3D models of many servers at our home page. You can look at the system from every side, look into the system, take out some components. I found this quite practical in presentations at a customer site instead of showing some pictures, not showing exactly that part, that the customer wanted to see. So i searched for a way to take the complete modell to the customer and finally found one.
1. Go to the Webpage of the server an click for the 360 degrees model of the server. I will use the 3D model of the Sun Fire x4140.
2. Use the context menu to display the source code of the page.
3. Search for the tag IFRAME. This should yield something like
The highlighted part is the directory for the model on the Sun webserver. You will need this directory in the following steps.
4. Create a directory for the files and change to this directory.
mkdir -p server/4140
#cd server/4140
5. At first gather with your prefered tool the index.html. I will use curl on MacOS X for this task, but other tools like wget will do it as well.
curl http://frsun.downloads.edgesuite.net/sun/08C12320/index.html > index.html
6. Grep for kaonApp in the index.html:
grep "kaonApp" index.html
kaonApp(535, 413, 'SplashMessageText=""; SplashMessageImage=""; maxWaitTime=10;
Data.Program.model="model_4140.jar"; Data.Program.modelUI="model_4140_ui.jar"; shadowDist=0.1;
customizeUI={panLock={false};}; customizeModel={};');
7. The both highlighted files are the interesting ones. Download both to your local machine.
curl http://frsun.downloads.edgesuite.net/sun/08C12320/model_4140.jar > model_4140.jar
#curl http://frsun.downloads.edgesuite.net/sun/08C12320/model_4140_ui.jar > model_4140_ui.jar
8. Now you have to download serveral other files from this directory:
#curl http://frsun.downloads.edgesuite.net/sun/08C12320/mesonApplet.jar > mesonApplet.jar
curl http://frsun.downloads.edgesuite.net/sun/08C12320/sceneGluon.jar > sceneGluon.jar
curl http://frsun.downloads.edgesuite.net/sun/08C12320/rasterGluon.jar > rasterGluon.jar
curl http://frsun.downloads.edgesuite.net/sun/08C12320/ui.jar > ui.jar
curl http://frsun.downloads.edgesuite.net/sun/08C12320/kaonapp.js > kaonapp.js
9. At the end all files should found their way to your directory:
dhcp-eham02-157-71:4140 joergmoellenkamp$ ls -l
total 2168
-rw-r--r-- 1 joergmoellenkamp staff 611 17 Apr 18:11 index.html
-rw-r--r-- 1 joergmoellenkamp staff 1558 17 Apr 18:14 kaonapp.js
-rw-r--r-- 1 joergmoellenkamp staff 32753 17 Apr 18:16 mesonApplet.jar
-rw-r--r-- 1 joergmoellenkamp staff 887335 17 Apr 18:14 model_4140.jar
-rw-r--r-- 1 joergmoellenkamp staff 2641 17 Apr 18:14 model_4140_ui.jar
-rw-r--r-- 1 joergmoellenkamp staff 67924 17 Apr 18:17 rasterGluon.jar
-rw-r--r-- 1 joergmoellenkamp staff 62282 17 Apr 18:17 sceneGluon.jar
-rw-r--r-- 1 joergmoellenkamp staff 37328 17 Apr 18:17 ui.jar
10. Now open the index.html as a file in your browser, in my example on my system open
file:///Users/joergmoellenkamp/4140/index.html

Tips:
1. You can resize the window by editing the index.hmtl:
kaonApp(535, 413, 'SplashMessageText=""; SplashMessageImage=""; maxWaitTime=10;

Data.Program.model="model_4140.jar"; Data.Program.modelUI="model_4140_ui.jar"; shadowDist=0.1; customizeUI={panLock={false};}; customizeModel={};');By changing the both highlighted values, you can resize the visualisation to every size you want.

Posted by Joerg Moellenkamp in Sun at 08:09

Monday, May 26. 2008

No blogging or LKSF today ...

... as i´m celebrating my "half of Joergs expected MTBF" birthday this evening...

Posted by Joerg Moellenkamp at 15:14

Elektrolux

Jetzt weiss ich endlich, was damit gemeint ist, wenn gesagt wird, das Geräte von Elektrolux für den A.... wären:

Wie Logos wirklich entstehen, könnt ihr hier sehen: Logólogos

Posted by Joerg Moellenkamp at 13:20

Sunday, May 25. 2008

On GPUs and CPUs

I assume many of us played a while with raytracing programs at a point of our contact with computers. But i think almost all of us stopped to play with it out of the same two reason like me: At first ray tracing was really really slow on a i486 and because of a lack in artistic capabilties.
Okay, the second reason can´t be solved by technology, but i found this older blog entry really interesting: At the Research@Intel blog you will find Real Time Ray-Tracing: The End of Rasterization? written by Jeffrey Howard. The key point of the computing power of ordinary systems reaches fastly the level of being sufficient to do raytracing in real time.

I think this leads to an interesting question: There are two groups of components in the industry. Components with and without a steady rise in performance. When you look at hard disks , cdroms and to an certain part memory, these components got faster over time, but at a modest rate and sometimes these technologies reached there maximum capacity without breaking physics or their adequateness for home use . (You simply don´t want a 15k harddisk in your deskside pc. I´ve worked in hearing distance of a RAID array with 12 15k disks and learned that my noise cancelation earphones are not just good for flying). The other group of components, GPUs and CPUs made giant steps in the last let´s say 10 years.

I get ware this fact everytime when i talk about Kilofractals per seconds as a metric to measure processor speed, remembering that my first fractal code written in C64 assembler needed virtually days to complete a single fractal.

Some years ago there was a gap to fill. Processors weren´t powerful enough to render graphics for computer games. Some companies stepped in and developed coprocessor for this task (anybody remember the Voodoo cards?) and over time this coprocs got really huge beasts (look at the transistor counts, clock frequencies and power usage of modern highend graphic card). But at the end these procs are specialized coprocs.

Now an old technology got new clothes: Modern general purpose processors are fast enough to calculate the mathematics behind ray tracing at high speed and even better, ray tracing is a task scaling exceptionally well over many cores, as the caluclations for the pixels are independent from each other.

Okay, when you are able to do photorealisic (do you need more than photorealistic?) ray tracing at 50fps (good enough) with a eight core proc at 1080p (more than enough), where is the need of using specialized coprocs with a specialized programming model. At the end: Where is the need for companies producing this specialized coprocs?

Integrating FPUs into the CPU obliterated the market for external FPU (anybody remembering IIT coprocs? Cyrix FasMath? Weitek Abacus? C&T SuperMath ?) besides a highly specialized market for FPGA based coprocs. The hunt for more cores in a single socket of a general purpose CPU may do the same for GPUs. Think only about the 12 core proc announced by AMD. Interesting times are lying in front of us.

Posted by Joerg Moellenkamp in The IT Business at 22:25

747 breaks at start

I didn´t believed, that something like that happens: A freighter 747 starts to take off, pilot rotates the machine, hears a strange sound, aborts the take off and the whole airframe breaks. You don´t believe it? Okay, here is a photo of N704CK at the end of the runway 20 in Brussels. The most unbelievable: No casualties despite of having enough fuel in the wings to fly to Bahrain.

Posted by Joerg Moellenkamp in Aviation at 19:39

links for 2008-05-25

animals on the underground
(tags: art cool design funny map web)

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

Kris erklaert: "Warum root logins boese sind?"

So möchte ich auch mal erklären können: Warum eigentlich ist "PermitRootLogin yes" in sshd_config evil?. Bei Solaris gibt es für solche Zwecke Privileges und RBAC, erklaert in diesem Tutorial: Less known Solaris Features: RBAC and Privileges

Posted by Joerg Moellenkamp at 12:56

Saturday, May 24. 2008

Debian openssl problem

I´ve suspected in a coffee kitchen discussion with a customer after a presentation a few days ago, that the Debian openssl vulnerability will haunt us for a while and hit us from behind while showing us many problemswith our existing internet security infrastructure. I assume i was correct. Fefe explains some of the uprising problems in his blog.

Posted by Joerg Moellenkamp in Security at 13:56

Infoworld reviews X4150

Paul Venezia of Infoworld tested the Sun Fire X4150 and was quite pleased with the performance - Lab test: Sun's X4150 shows the beef:
All in all, it's a very real competitor to HP's ProLiant DL360 G5, IBM's x3550, and Dell's PowerEdge 1950 III. In fact, the fit and finish of the X4150 are generally better than the rest of the field's and far ahead in local storage.
[...]
The X4150 is priced similarly to other servers of comparable spec, yet offers storage and I/O capabilities generally found in 2U and larger systems – this not only saves on rack space, but given the relatively low power consumption of the X4150, it can save on power as well.
[...]
The most suitable roles for the Sun Fire X4150 will be in database, Web serving, and virtualization tasks. In these spaces, the X4150 provides significant bang for the buck in nearly all respects

Posted by Joerg Moellenkamp in Sun at 13:31

Really early preview of the JET tutorial.

Just to give you an impression of the upcoming JET tutorial i´ve uploaded a really early version of the tutorial. You can download it here. The Jumpstart tutorial starts at page 215. All graphics are missing, some sections are not merged, but it should give you a first view.

Update: The link provided an earlier version of this article was incorrect. I´ve corrected it, you should be able to download the new version. Note to myself: Don´t write an article while waiting for TeX

Posted by Joerg Moellenkamp in Solaris at 11:48

Jonathan is a Mac user ...

White case, a black spot above the LC display, the typical monitor dongle at the side. This looks like a MacBook. Jonathan is a Mac user...

Posted by Joerg Moellenkamp in Apple at 10:03

Friday, May 23. 2008

aviation.c0t0d0s0.org

I will kill of aviation.c0t0d0s0.org in the next few days. I simply didn´t had enough time for a second blogging workflow. I will integrate that special interest blog into the main blog in the next few days.

Posted by Joerg Moellenkamp in Blogosphere at 11:05

Breaking news: Hell freezes over ... or: Microsoft supports ODF

This are good news: Microsoft announced to deliver ODF support in service pack 2:The 2007 Microsoft Office system already provides support for 20 different document formats within Microsoft Office Word, Office Excel and Office PowerPoint. With the release of Microsoft Office 2007 Service Pack 2 (SP2) scheduled for the first half of 2009, the list will grow to include support for XML Paper Specification (XPS), Portable Document Format (PDF) 1.5, PDF/A and Open Document Format (ODF) v1.1.

Posted by Joerg Moellenkamp in The IT Business at 08:17

Thursday, May 22. 2008

links for 2008-05-22

Extreme Networks X650
... neat ...
(tags: 10GbE ethernet network)

Posted by del.icio.us in del.icio.us at 13:37

Wednesday, May 21. 2008

Lux

I want this at home:

Multi-touch is definitely one of the "next big things" in user interfaces. Work for a few minutes with an iPhone ... or: simply look at the video and you will understand. And Lux looks promising to bring that to your Mac soon. My personal favorite was the moment, when they controlled Google Earth with gestures

Posted by Joerg Moellenkamp in Apple at 23:23

CIFS at home with Solaris

There is a nice walkthrough for configuring Opensolaris as an CIFS fileserver for home use at the Sun Developer Network . It´s really simple: Developer Recipes: Setting Up an OpenSolaris NAS Box. And NFS ist just a zfs set sharenfs=on examplepool/examplefilesystem away.

Posted by Joerg Moellenkamp at 22:54

Seltsames Weltbild

Ist jemanden bei der UnionInvestment Werbung im Fernsehen schon das unrealistische Weltbild aufgefallen. Die mit dem von der Wiege bis zur Bahre. Da gibt es ja auch das "der erste Kuss" Szenario. Nur, der Typ sitzt auf dem Fahrersitz eines Autos. Der Herr muss demnach 18 sein. Wer glaubt denn wirklich, das heute mit dem ersten Kuss noch bis 18 gewartet wird?

Posted by Joerg Moellenkamp in Braindump at 19:25

Tuesday, May 20. 2008

links for 2008-05-20

A Mini-ITX board with an Atom CPU
(tags: computer Intel diy hardware)

Posted by del.icio.us in del.icio.us at 13:37

Dual use

Bruce Schneier wrote a really interesting article about the implications of dual-use in information technlogy:Dual-Use Technologies and the Equities Issue:You know you've got a problem when you can't tell a hostile attack by another nation from bored kids with an axe to grind.

Posted by Joerg Moellenkamp at 13:08

Thermal problems

Your old notebook has severe thermal problems because of a defunct fan or something like that and you need a last boot. This may be an effective solution for you: How to rescue data from an iBook with thermal problems

Posted by Joerg Moellenkamp in Apple at 11:47

SunSTAR Erik Fischer

Wondered, it took this long for the Divas to name him a SunStar. This guy has forgotten more about processors than many others know about processors. Heard a really good presentation about CPU design on the CEC2006 and his internal nomac newsletter is excellent.

Posted by Joerg Moellenkamp in Sun at 11:05

State of the Jumpstart Tutorial

Pheew ... i shouldn´t have started to write the JET tutorial. 14 Pages and i barely covered the basics.

Posted by Joerg Moellenkamp in Solaris at 10:04

Great ...

Posted by Joerg Moellenkamp in Fundsache at 08:25

Monday, May 19. 2008

Und dann war da noch ...

... die ältere, distinguierte Dame im schwarzen Businesskostuem, die eine Sitzreihe vor mir ins Telephon sprach: "... dann ist die Terminplanung für den Arsch" ...

Posted by Joerg Moellenkamp in Bahn at 19:00

Acrobat Reader on Solaris x86

Albeit there are some good alternatives, the missing original Acrobat Reader was a pain point for many people in Solaris x86. Adobe announced in the Acroread blog a port for Solaris x86. At last.

Posted by Joerg Moellenkamp in Solaris at 14:33

links for 2008-05-19

Review: Casio EX-F1 Is a Speed-Demon Snapper | Gadget Lab from Wired.com
(tags: cool hardware photography)

Posted by del.icio.us in del.icio.us at 13:37

Albträume

Spät abends noch Tee trinken, mit knurrendem Magen ins Bett gehen und das Ganze kurz vor Vollmond. Oder anders gesagt: Ich habe mal wieder die Landebahnbefeuerung für Albträume eingeschaltet. Hätte den Traum gleich aufschreiben sollen, hätte einen interessanten Verschwörungsroman abgegeben V for Vendetta meets 1984 meets Matrix meets Twelve Monkeys meets eXistenZ meets Welt am Draht. Kann ich nicht mal einen vernünftigen Albtraum habe. So mit Fallschirmspringen ... ohne Fallschirm ... oder mit Spinnen oder Ratten. Halt was normales. Naja, jetzt erinnere ich mich auch nur noch an die Hälfte. Anderen scheints aber ähnlich zu gehen, im Haus gegenüber brennt seit etwa einer halben Stunde auch wieder Licht.

Posted by Joerg Moellenkamp in Braindump at 03:56

Sunday, May 18. 2008

Hoerempfehlung: Three Doors Down - Three Doors Down

Vor einigen Jahren bin ich ueber eine Band gestolpert, die mich damaligerzeit entfernt an an die von mir in Dauerrotation gehoerte "Throwing Copper" von Live erinnert hat, also jene CD von dieser Gruppe, die noch richtig gut war. Ich finde, das "I alone" bis heute eines der wirklich wichtigen Lieder der 90er Jahre war. Wird wohl wieder einer Zeit mit Liebeskummer gewesen sein in der ich die CD wieder rausgekramt habe, kein Lied passt fuer mich dann besser. Muesste ich alle Musikstuecke der 90er bis auf 10 loeschen, haette ich nur noch neun weitere Titel, die ich behalten wuerde.

Ich schweife ab: Das Lied, das ich damals gehoert habe , war Loser von 3 Doors Down. Seit dem höre ich ziemlich gerne die Musik dieser Band, insbesondere Landing in London hat es mir angetan. Vorspulen in die Jetztzeit:3 Doors Down haben ein neues Album herausgebracht. Titel: Three doors down. Gefällt mir sehr. Sehr hörbar. Hörempfehlung: Your arms feel like home, Train und Give it to me. Neuer Kandidat fuer Dauerrotation im iPhone, ob sie allerdings Styrofoam abloesen in meinen persoenlichen Charts abloesen wird, weiss ich noch nicht.

Posted by Joerg Moellenkamp in Music at 19:05

Saturday, May 17. 2008

Strange

Hmm, strange ... Solaris Express b87 don't like Apples Mighty Mouse (the one with cable) when it boots into the installation when used with the ASUS M2N-VM DVI board. Logitech connected to ASUS ... everything is okay, Mighty Mouse connected to Asrock ... everything is okay ... just Mighty Mouse connected to Asus ... deeeep freeeze ...

Posted by Joerg Moellenkamp in Solaris at 15:28

links for 2008-05-17

randomness
(tags: debian fun)

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

Jumpstart Tutorial postponed

I´ve slowed dow the work on my Jumpstart tutorial. PXEgrub isn´t able to initialize the network cards at the moment (i even tried an old realtek card) and i haven´t the least inkling of an idea, why this is the case.

Posted by Joerg Moellenkamp in Solaris at 11:44

Friday, May 16. 2008

Kernel code quality

Interesting artcle about the quality of the kernel source code. Diomidis Spinelli compared the FreeBSD, GNU/Linux, Solaris, and Windows kernels for his article A Tale of Four Kernels:Despite various claims regarding the efficacy of particular open or close-source development methods, we can see from the table that there is no clear winner (or loser). The two systems with a commercial pedigree (Solaris and WRK) have slightly more positive than negative marks. However, WRK also has the largest number of negative marks, while Solaris has the second lowest number of positive marks. Therefore, the most we can read from the overall balance of marks is that open source development approaches do not produce software of markedly higher quality than proprietary software development.An interesting read.

Posted by Joerg Moellenkamp in Solaris at 22:42

Problems

Well ... i finally know why i like OBP and my old UltraSPARC 10 workstations. Network installs are vastly easier with them. The process of writing the JET tutorial was stopped by an mediocre implementation of PXE ... fsck .... will buy an Intel network card tomorrow morning.

Posted by Joerg Moellenkamp at 22:07

links for 2008-05-16

The 25 Year Old BSD Bug
(tags: bsd geek history programming unix)

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

Thursday, May 15. 2008

Sorry

Sorry for writing nothing today, i'm preparing my new test system at the moment. As this system doesn't even have a optical drive, it's the basis for my next tutorial. This will be a revised version of my wiki article about JET, as i want to jumpstart the new machine.

Posted by Joerg Moellenkamp in Solaris at 20:59

links for 2008-05-15

NachDenkSeiten - Die kritische Website Â» Was Sie beachten sollten, wenn Sie (als junger Akademiker/Student) einem Finanzdienstleister gegenÃ¼bersitzen?
(tags: anlage geld)

Market Ticker: Tall Tale Tuesday
(tags: economics government statistics US)

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

Wednesday, May 14. 2008

Will Firefox go bad?

The large user community of Firefox results from a distrust in regard of the good intentions of Firefox and a the strong believe, that open source products are inherently good. But what browser should you use, when the CEO of Mozilla thinks about massively collecting data from the users. Going back to Microsoft, using Opera or should i say goodbye to all this multimedia stuff and start to use lynx. Will this end in telnet c0t0d0s0.org 80 as the last way to browse without being subject to data collecting?

Posted by Joerg Moellenkamp in Privacy at 20:33

New write throttling in ZFS

Roch reports in The new ZFS write throttle about a new algorithm in ZFS for throttling writes appearing in upcoming versions of Opensolaris. Every five seconds ZFS commits changes to it´s pools to the disk, even when you are not forcing this syncing. Whenever your harddrives doesn´t keep up with the amount of writes from the transaction group, it was possible, that ZFS trottled writing application to ensure that the you can write all the outstanding data in the transaction groups. Depending of the amount of the data and the speed of your harddisks this state could last a few seconds.

The new behaviour is much better from my point of view. Roch writes in his article:So the new steady state behavior of write intensive workloads is that, starting with an empty TXG, all threads will be allowed to dirty memory at full speed until a first threshold of bytes in the TXG is reached. At that time, every write system call will be delayed by 1 tick thus significantly slowing down the pace of writes. If the previous TXG completes it's I/Os, then the current TXG will then be allowed to resume at full speed. But in the unlikely event that a workload, despite the per write 1-tick delay, manages to fill up the TXG up to the full threshold we will be forced to throttle all writes in order to allow the storage to catch up.This new mechanism should keep away this complete write throttling pauses really effectively.

Posted by Joerg Moellenkamp in Solaris at 17:11

Spring

Posted by Joerg Moellenkamp in Photographie at 08:17

Less known Solaris features: About crashes and cores - Appendix A: Crash Dump Analysis revisited

I have forgotten to describe a really handy function when you analyse a crash dump. Solaris has an in-memory buffer for the console messages. In the case you write a crash dump, obviously this messages are written into the crash dump as well.
With the ::msgbuf command of mdb you can read this message buffer.

```
# mdb unix.0 vmcore.0
Loading modules: [ unix genunix specfs cpu.generic uppc scsi_vhci ufs ip hook neti sctp arp usba nca lofs zfs random nsctl sdbc rdc sppp crypto ptm ]
> ::msgbuf
MESSAGE
SunOS Release 5.11 Version snv_84 32-bit
Copyright 1983-2008 Sun Microsystems, Inc.  All rights reserved.
Use is subject to license terms.
features: 10474df
mem = 331388K (0x1439f000)
root nexus = i86pc
pseudo0 at root
pseudo0 is /pseudo
```

[...]
devinfo0 is /pseudo/devinfo@0

panic[cpu0]/thread=db3aea00:
forced crash dump initiated at user request

d5efcf4c genunix:kadmin+10c (5, 0, 0, db5c8a98)
d5efcf84 genunix:uadmin+8e (5, 0, 0, d5efcfac,)

syncing file systems...
 done
dumping to /dev/dsk/c0d0s1, offset 108593152, content: all
>
So it´s really easy to get this last messages of a dying system with mdb from the crash dump alone.

Posted by Joerg Moellenkamp in Solaris at 07:14

Tuesday, May 13. 2008

New server from Sun: X4240

Sun announced some new systems today: At first the X4140 and X440 are available with quadcore AMD Opteron processors . Additionally we´ve announce a completely new system, the

Posted by Joerg Moellenkamp in Sun at 15:26

links for 2008-05-13

Moonwatcher: Why Doesn't Apple Face The Innovator's Dilemma?
(tags: apple article blog business technology)

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

Ars.Technica about Opensolaris 2008.05

Ars.technica tested Opensolaris 2008.05 in "First look: OpenSolaris 2008.05 a work in progress":Although the OpenSolaris development community still has a lot of work to do before the operating system is ready to take on Linux on the desktop, the progress so far indicates that the project deserves further attention. We will keep an eye on future releases to see how the platform evolve.

Posted by Joerg Moellenkamp in Solaris at 08:00

At last a decent OS for the Eee PC ;)

Rolf Kersten managed to install Opensolaris 2008.05 on one of this small EeePC devices. You need some workarounds, but overall the installation is fairly easy with the help of Rolfs article: OpenSolaris 2008.05 on the Eee PC.

Posted by Joerg Moellenkamp at 07:09

First livesign of iSER in Opensolaris

The development team for iSER (iSCSI over RDMA over Infiniband) announced, that their prototype is up and running : [iser-dev] iSER prototype implementation is working!!!

Posted by Joerg Moellenkamp in Solaris at 07:03

Zones in Opensolaris 2008.05

Containers in Opensolaris 2008.05 work a little bit different like the ones in the other versions of Solaris. They are an own brand at the moment. Dan Price reports in his article A field guide to Zones in OpenSolaris 2008.05 about the differences.

Posted by Joerg Moellenkamp in Solaris at 06:53

Monday, May 12. 2008

Vollkontaktbahnfahren

Was passiert, wenn ein Zug ausfaellt (okay, unbestimmt verspätet, das ist das gleiche wie ein Ausfall, war wohl trotz Feiertag wieder Suicide Monday) und sich dann die Passagiere, die eigentlich eine Fahrt im IC geplant haben, auch noch in einen ohnehin regelmaessig sehr vollen Regionalexpress quaelen muessen. Genau ... Vollkontaktbahnfahren. Wäre dieser Zug an einer Kontrolle für Viehtransporte vorbeigefahren, waere wahrscheinlich der Zug ausser Verkehr gezogen worden.

Oben stehen, unten stehen, auf den Treppen stehen, im Eingang stehen, im Ausgang stehen, vorm Klo stehen und dazwischen die gluecklichen die eine Stufe, eine Lehne oder gar einen Sitz ergattert habe. Dann dazwischen die Leute, die gezwungenermassen ihr Gepack auf einen Sitz stellen muessen, weil die Bahn in ihrer unendlichen Weissheit Zuege bestellt hat, die ueber keine nennenswerten Kofferablagen verfuegt. Die heutige Fahrt war echt zum Abgewoehnen. Ich weiss schon, warum ich versuche zu den Zeiten nach Oldenburg zu fahren, bei denen beide Streckensegmente mit IC befahren werden. Wenn sie denn fahren ...

Posted by Joerg Moellenkamp in Bahn at 21:31

links for 2008-05-12

Sun exec ponders OpenSolaris, Linux | InfoWorld | News | 2008-05-09 | By Paul Krill
(tags: debian interview linux opensolaris opensource)

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

Hawking

This is the kind of program i would expect from the public broadcasting system in Germany, not this ever-increasing insult consisting out of game show derviates, daily soaps and half-assed documentaries. When you look at the primary programs in TV you will see the tenth telecast about making a special sausage or how hard it was to transport somebing big from A to B but nothing really interesting.

Okay, back to this movie: This is an really interesting documentary produced by the BBC about Stephen Hawking, the information paradoxon and his actual research.

Posted by Joerg Moellenkamp at 11:44

Magnetic dust

A hard disk goes down in flames through the atmosphere and specialists were still able to recover 90 percent of the data: Data recovered from Seagate drive in Columbia shuttle disaster. At the end, data seems to be more resistant then most of us believe, when you put a vast amount of effort behind the recovery.

This leads me to annother thoughtgame: Let´s assume, you shredder a disk to pieces as tiny as 1 square millimeter. This would just look like dust to you. A modern harddisk stores 200 gigabit per square inch. One square inch are 645.16 square milimeters. Thus a harddisk would store 310 megabits on a square milimeter. Let´s assume 10 bits per byte (for error correction and similar things) and you have 31 megabytes worth of data on one of these pieces of dust.

It´s just a question of effort to recover the data, when you can yield 10 million euros out of the data (trade secrets, credit card data) it would give you a nice profit when you spend for example 9 million to recover the data. Yet another reason

for cryptography everywhere or you may end up with degaussing, shreddering and remelt your old harddisks just to be safe.

Posted by Joerg Moellenkamp in Security at 09:57

Load Average

It´s a common practice between us admins to look after the load averages. Interestingly the correct math behind this numbers is not really simple. Roman wrote about it at How many kernel engineers does it take to solve one differential equation?. Interesting read.

Posted by Joerg Moellenkamp in Technology at 08:37

Sunday, May 11. 2008

Spring

Posted by Joerg Moellenkamp in Photographie at 21:03

Pfingsten in Barßel

Der 80igste Geburtstagsfeier meines Grossvaters mütterlicherseits ist hinter mich gebracht, elterliche Grillage im Magen versenkt und jetzt einfach nur noch aufm Balkon oder auf der Terasse sitzen und sich den Rest der Sonne auf den Schaedel scheinen lassen und dann morgen sich in den Zug quaelen.

Posted by Joerg Moellenkamp at 19:17

COMSTAR in build 90

The new Common SCSI Target Framework made it into build 90 of Opensolaris. In this first integration, the target mode driver for qlogic cards ins included. It´s now really easy to create a FC storage array out of a X4500 and opensolaris. You will find further informations on the project page at opensolaris.org.

Posted by Joerg Moellenkamp in Solaris at 17:45

links for 2008-05-11

Linux File Systems: You Get What You Pay For
(tags: linux filesystems)

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

Saturday, May 10. 2008

links for 2008-05-10

Designing and implementing malicious hardware
(tags: CPU hardware research security)

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

2000 Downloads of the LKSF book

2000 downloads of the "Less known Solaris features" book in four days. Pretty cool. I have already some plans for the future. It will stay free under the Creative Commons license. I have no plans to make a real book out of it. One of the next larger projects will be the rewrite of my tutorial regarding the Jumpstart Enterprise Toolkit.

Okay, as the fourth person just asked me about a Amazon Wishlist: Okay, you find it here. This is just for the ones, who really can´t resist. The tutorial is my way to give something back to the community. I didn´t wrote it to get something for. Okay, one important remark: I´m normally i use it as an public "to buy" list for myself. Thus do not wonder, when there are expensive items on it. I´ve tagged them with ignore, so just do that: ignore them.

Posted by Joerg Moellenkamp at 12:36

Friday, May 9. 2008

VomitOne

Alec Muffet coined this phrase for the final day of the JavaOne. Somehow the Norovirus was found at the JavaOne. The symptoms of norovirus illness are nausea, vomiting, diarrhea, and some stomach cramping. The Valleywag reports about this, too. In earlier times, it was Montezuma's revenge, the beer too much or unfamiliar meals. Today it has to be an virus ...

BTW: I find stomach problems in the US quite normal. I don´t like this chlorinated water, and my stomach hates it. This is something i like in Germany. Pure water out of the tap without the feeling of being in a swimming pool.

Posted by Joerg Moellenkamp in Sun at 19:16

Less known Solaris Features: About crashes and cores

It´s getting warmer, so this tutorial is a little bit shorter. It´s about the a really hated topic: Panics, crash dumps, core dumps and the tools to manage and leverage this the informations resulting from this functions:

Part 1: Introduction
Part 2: Forcing a dump
Part 3: Controling the behaviour of the dump facilities
Part 4: Crash dump analysis for beginners
Part 5: Conclusion
Appendix A: Crashdump analysis revisited
Appendix B: Live Crashdumps

I´ve also updated the pdf book version of this tutorial. So the releases beginning with the 9th May 2008 contain this tutorial.

Posted by Joerg Moellenkamp in General at 14:32

Less known Solaris Features: About crashes and cores - Part 5: Conclusion

The Solaris Operating Environment has several functions to enable the user or the support engineer in the case something went wrong. Crash and core dumps are an invaluable resource to find the root cause of a problem. Don´t throw them away without looking at them.
Do you want to learn more?
Documentation
coreadm(1M)
dumpadm(1M)
uadmin(1M)
mdb(1)
Solaris Modular Debugger Guide

Books
Solaris Internals: Solaris 10 and Open Solaris Kernel Architecture - Richard McDougall and Jim Mauro
Solaris Performance and Tools: DTrace and MDB Techniques for Solaris 10 and OpenSolaris - Richard McDougall , Jim Mauro and Brendan Gregg

Posted by Joerg Moellenkamp in Solaris at 13:39

Less known Solaris features: About crashes and cores - Part 4: Crashdump analysis for beginners

Okay, now you have all this crash and core dumps, it would be nice to do something useful with it. Okay, i show you just some basic tricks to get some insight into the state of a system when it wrote a crash dump.
Basic analysis of a crash dump with mdbAt first we load the dump into the mdb:# mdb unix.0 vmcore.0

Loading modules: [unix genunix specfs cpu.generic uppc scsi_vhci ufs ip hook neti sctp arp usba nca lofs zfs random nsctl sdbc rdc sppp]
>A nice information is the backtrace. This helps you to find out, what triggered the crash dump. In this case it´s easy. It´s was the uadmin syscall.> $c
vpanic(fea6388c)
kadmin+0x10c(5, 0, 0, db39e550)
uadmin+0x8e()
sys_sysenter+0x106()But it would be nice, to know more of the state of the system, at the moment of the crash. For example we can print out the process table of the system like we would do it with ps> ::ps
S PID PPID PGID SID UID FLAGS ADDR NAME
R 0 0 0 0 0 0x00000001 fec1d3d0 sched
[...]
R 586 1 586 586 0 0x42000000 d55f58a8 sshd
R 545 1 545 545 0 0x42000000 d5601230 fmd
R 559 1 559 559 0 0x42000000 d55fb128 syslogd
[...]
R 533 494 494 494 0 0x4a014000 d55f19c0 ttymonWe can even lookup, which files or sockets where opened at the moment of the crash dump. For example: We want to know the open files of the ssh daemon. To get this information, we have to use the address of the process from the process table (the eigth column) and extend it with "::pfiles":> d55f58a8::pfiles
FD TYPE VNODE INFO
0 CHR d597d540 /devices/pseudo/mm@0:null
1 CHR d597d540 /devices/pseudo/mm@0:null
2 CHR d597d540 /devices/pseudo/mm@0:null
3 SOCK db688300 socket: AF_INET6 :: 22 And here we look into the open files of the syslogd,> d55fb128::pfiles
FD TYPE VNODE INFO
0 DIR d5082a80 /
1 DIR d5082a80 /
2 DIR d5082a80 /
3 DOOR d699b300 /var/run/name_service_door [door to 'nscd' (proc=d5604890)]
4 CHR db522cc0 /devices/pseudo/sysmsg@0:sysmsg
5 REG db643840 /var/adm/messages
6 REG db6839c0 /var/log/syslog
7 CHR db522840 /devices/pseudo/log@0:log
8 DOOR db6eb300 [door to 'syslogd' (proc=d55fb128)]
As the core dump contains all the pages of the kernel (or more, in the case you configure it) you have a frozen state of your system to investigate everything you want.

And to get back to my security example: With the core dump and mdb you can gather really interesting informations. For example, you can see that an ssh connection was open at the time of the crash dump.
> ::netstat
TCPv4 St Local Address Remote Address Stack Zone
db35f980 0 10.211.55.200.22 10.211.55.2.53811 0 0
[...]
An example from the fieldYou can do it like the pros and and look at source code and crash dump side by side to finde the root cause for an error. Or like some colleagues at the Sun Mission Critical Support Center who wouldn´t surprise me, when they find the error by laying their hand on a system).

For all others, there is a more simple way to analyse your crash dump to have at least a little bit more informations to search in a bug database.

I will use a crash i´ve analysed a long time ago to show you the trick. Okay, you have to start a debugger. I used mdb in this example:bash-3.00# mdb -k unix.4 vmcore.4
Loading modules: [unix krtld genunix specfs dtrace cpu.AuthenticAMD.15 uppc pcplusmp ufs md ip sctp usba fcp fctl nca lofs cpc fcip random crypto zfs logindmux ptm sppp nfs ipc]A prompt appears, just type in $C to get a stack trace.
> $C
fffffe80000b9650 vpanic()
fffffe80000b9670 0xfffffffffb840459()
fffffe80000b96e0 segmap_unlock+0xe5()
fffffe80000b97a0 segmap_fault+0x2db()
fffffe80000b97c0 snf_smap_desbfree+0x76()

fffffe80000b97e0 dblk_lastfree_desb+0x17()
fffffe80000b9800 dblk_decref+0x66()
fffffe80000b9830 freeb+0x7b()
fffffe80000b99b0 tcp_rput_data+0x1986()
fffffe80000b99d0 tcp_input+0x38()
fffffe80000b9a10 squeue_enter_chain+0x16e()
fffffe80000b9ac0 ip_input+0x18c()
fffffe80000b9b50 i_dls_link_ether_rx+0x153()
fffffe80000b9b80 mac_rx+0x46()
fffffe80000b9bd0 bge_receive+0x98()
fffffe80000b9c10 bge_intr+0xaf()
fffffe80000b9c60 av_dispatch_autovect+0x78()
fffffe80000b9c70 intr_thread+0x50()Okay, now start at the beginning of the trace to strip all lines from the operating system infrastructure for error cases. Okay, vpanic() generates the panic. The second line is useless for our purposes to. The next both lines with segmap are generated by the error but not the root cause. The interesting line ist snf_smap_desbfree

With this name you can go to Sunsolve or bugs.opensolaris.org. Et voila : System panic due to recursive mutex_enter in snf_smap_desbfree trying to re-aquire Tx mutex. When you type this error into the PatchFinder, you will find a patch fixing this bug: 124255-03

Two hints:

It´s a good practice to know mdb. It´s very useful at compiling open source software in the case your compiled code throw cores, but you don´t know why. core files are not just for deleting them.
Error reports with a stack trace are more usefull than an error report just with "The system paniced when i did this"

Posted by Joerg Moellenkamp in Solaris at 13:23

Less known Solaris features: About crashes and cores - Part 3: Controlling the behaviour of the dump facilities

Solaris has mechanisms to control the behaviour of dump mechanisms. These mechanisms are different for crash and core dumps.
Crash dumpsYou can configure the content of crashdumps, where they are located and what you do with them after the boot. You control this behaviour with the dumpadm command. When you use this command without any further option, it prints out the actual state.# dumpadm
Dump content: kernel pages
Dump device: /dev/dsk/c0d0s1 (swap)
Savecore directory: /var/crash/incubator
Savecore enabled: yesThis is the default setting: A crash dump contains only the memory pages of the kernel and uses /dev/dsk/c0d0s1 (the swap device) to store the crash dump in the case of a kernel panic. savecore is a special process, that runs at the next boot of the system. In the case of an crash dump at the dump device, it copies the dump to the configured directory to keep it for analysis before it´s used for swapping again.
Let´s change the behaviour. At first we want to configure, that the complete memory is saved to the crash dump in case of a panic. This is easy:# dumpadm -c all
Dump content: all pages
Dump device: /dev/dsk/c0d0s1 (swap)
Savecore directory: /var/crash/incubator
Savecore enabled: yes
Okay, now let´s change the location for the crash dump. The actual name is an artefact of my original VM image called incubator. To get a new test machine i clone this image. I want to use the directory /var/crash/theoden for this purpose.# mkdir /var/crash/theoden
chmod 700 /var/crash/theoden
dumpadm -s /var/crash/theoden
Dump content: all pages
Dump device: /dev/dsk/c0d0s1 (swap)
Savecore directory: /var/crash/theoden
Savecore enabled: yesNow the system will use the new directory to store the crash dumps. Setting the rights of the

directory to 700 is important. The crash dump may contain sensitive information, thus it could be dangerous to make them readable by anyone else than root.

Core dumpsA similar facility exists for the core dumps. You can control the behaviour of the core dumps with the coreadm command. Like with dumpadm you can get the actual configuration by using coreadm without any option.# coreadm
```
     global core file pattern:
     global core file content: default
       init core file pattern: core
       init core file content: default
            global core dumps: disabled
      per-process core dumps: enabled
     global setid core dumps: disabled
per-process setid core dumps: disabled
    global core dump logging: disabled
```
This programm has more options than dumpadm. I won´t go through all options, but some important ones.

From my view the file patterns are the most interesting ones. You can control, where core dumps are stored. The default is to store the core dumps in the working directory of a process. But this may lead to core dumps dispersed over the filesystem.

With core adm you can configure a central location for all your coredumps.
coreadm -i /var/core/core.%n.%f.%u.%p
coreadm -uWith -i you tell coreadm to set the location for the per-process core dumps. The parameter for this option is the filename for new core dumps. You can use variables in this filename. For example %n will be translated to the machine name, %f to name of the file, %u to the effective user id of the process and %p will be substituted with the process id. The coreadm -u forces the instant reload the configuration. Otherwise, this setting would get active at the next boot or the next refresh of the coreadm service. Okay, let´s try our configuration.# ps -ef | grep "bash" | grep "jmoekamp"
```
jmoekamp   681   675   0 20:59:39 pts/1       0:00 bash
```
Now we trigger a core dump for a running process.
gcore -p 681
gcore: /var/core/core.theoden.bash.100.681 dumped
As you see, the core dump isn´t written at the current working directory of the process, it´s written at the configured position.

Core dump configuration for the normal userThe both configuration described so far are global ones, so you can do this configuration only with root privileges. But a normal user can manipulate the core dump configuration as well, albeit only for processes owned by her or him.

Let´s login as a normal user. Now we check one of our processes for it´s coreadm configuration:$ ps -ef | grep "jmoekamp"
```
jmoekamp   712   670   0 01:27:38 pts/1       0:00 -sh
jmoekamp   669   666   0 22:29:15 ?           0:00 /usr/lib/ssh/sshd
jmoekamp   670   669   0 22:29:16 pts/1       0:00 -sh
jmoekamp   713   712   0 01:27:38 pts/1       0:00 ps -ef
```
$ coreadm 669
669:    /var/core/core.%n.%f.%u.%p   defaultNow let´s check a process owned by root.
$ ps -ef | grep "cron"
```
jmoekamp   716   670   0 01:28:13 pts/1       0:00 grep cron
   root   322     1   0 22:25:24 ?           0:00 /usr/sbin/cron
```
$ coreadm 322
322: Not ownerThe system denies the access to this information. Now we change the setting for the process 669 from the first example. It´s quite simple:$ coreadm -p /export/home/jmoekamp/cores/core.%n.%f.%u.%p 669
$ coreadm 669
669:    /export/home/jmoekamp/cores/core.%n.%f.%u.%p   default
The per-process core file name pattern is inherited by future child processes of the affected processes.

Why should you set an own path and filename for an application or an user? There are several reasons. For example to ensure that you have the correct rights to an directory for the cores, when the process starts to dump the core or to seperate the cores from certain applications a different locations.

Posted by Joerg Moellenkamp in Solaris at 13:11

Less known Solaris features: About crashes and cores - Part 2: Forcing a dump

Okay, a dumps are not only a consequence of errors. You can force the generation of both kinds. This is really useful when you want to freeze the current state of the system or an application for further examination.
Forcing a core dumpLet´s assume you want to have an core dump of a process running on your system:# ps -ef | grep "bash" | grep "jmoekamp"
jmoekamp 681 675 0 20:59:39 pts/1 0:00 bashOkay, now we can trigger the core dump by using the process id of the process.# gcore 681
gcore: core.681 dumpedOkay, but the kicker is the fact, that the process still runs afterwards. So you can get an core dump of your process for analysis without interrupting it.# ps -ef | grep "bash" | grep "jmoekamp"
jmoekamp 681 675 0 20:59:39 pts/1 0:00 bashNeat isn´t it. Now you can use the mdb to analyse it, for example to print out the backtrace:# mdb core.681
Loading modules: [libc.so.1 ld.so.1]
> $c
libc.so.1`__waitid+0x15(0, 2a9, 8047ca0, 83)
libc.so.1`waitpid+0x63(2a9, 8047d4c, 80)
waitjob+0x51(8077098)
postjob+0xcd(2a9, 1)
execute+0x77d(80771c4, 0, 0)
exfile+0x170(0)
main+0x4d2(1, 8047e48, 8047e50)
_start+0x7a(1, 8047eec, 0, 8047ef0, 8047efe, 8047f0f)Forcing a crash dumpOkay, you can force a crash dump, too. It´s quite easy. You can trigger it with the uadmin command.bash-3.2# uadmin 5 0

panic[cpu0]/thread=db47700: forced crash dump initiated at user request

d50a2f4c genunix:kadmin+10c (5, 0, 0, db325400)
d50a2f84 genunix:uadmin+8e (5, 0, 0, d50a2fac,)

syncing file systems... 2 1 done
dumping to /dev/dsk/c0d0s1, offset 108593152, content kernel
100% done: 31255 pages dumped, compression ratio 5.31, dump succeeded
Press any key to reboot.
Why should you do something like that? Well, there are several reasons. For example, when you want to stop a system right at this moment. There is an effect in clusters called "split brain" . This happens, when both nodes of a cluster believe they are the surviving one, because they´ve lost the cluster interconnect(simplification warning). Sun Cluster can prevent this situation by something called quorum. In a high availability situation the nodes of a cluster try to get this quorum. Whoever gets the quorum, runs the service. But you have to ensure that the other nodes don´t even try to write something to disks. The simplest method: Panic the machine.

Another use case would be the detection of an security breach. Let´s assume, your developer integrated a security hole as large as the Rhine into a web applicaiton by accident and now someone else owns your machine. The false reaction would be: Switch the system off or trigger a normal reboot. Both would lead to the loss of the memory content and perhaps the hacker had integrated a tool in the shutdown procedure to erase logs. A more feasible possibility: Trigger a crash dump. You keep the content of the memory and you can analyse it for traces to the attacker.

Posted by Joerg Moellenkamp in Solaris at 13:01

Less known Solaris features: About crashes and cores - Part 1: Introduction

No software is without errors. This is a basic law of computer science. And when there is no bug in the software (by a strange kind of luck) your hardware has bugs. And when there are no bugs in the hardware, cosmic rays are flipping bits. Thus an operating system needs some mechanisms to stop a process or the complete kernel at once without allowing the system to write anything back to disk and thus manifesting the corrupted state. This tutorial will cover the most important concepts surrounding the last life signs of a system or an application.
A plea for the panicThe panic isn´t the bug, it´s the reaction of the system to a bug in the system. Many people think of panics as the result of an instability and something bad like the bogey man. But: Panics and crash dumps are your

friend. Whenever the system detects a inconsistency in it´s structures, it does the best what it could to: protect your data. And the best way to do this, is to give the system a fresh start, don´t try to modify the data on the disk and write some status information to a special device to enable analysis of the problem. The concepts of panic and crash dump were developed to give the admin exactly such tools.

A good example: Imagine a problem in the UFS, a bit has flipped. The operating environment detects an inconsistence in the data structures. You can´t work with this error. It would unpredictably alter the data on your your disk. You can´t shutdown the system by a normal reboot. The flushing of your disks would alter the data on the disk. The only way to get out of the system: Stop everything and restart the system, ergo panic the system and write a crash dump.

Furthermore: Some people look at the core and crash dumps and think about their analysis as an arcane art and see them as a wast of disk space. But it´s really easy to get some basic data and hints out of this large heaps of data.

Difference between Crash Dumps and Core DumpsMany people use this words synonymously ("The system paniced and wrote a core dump"). Every now and then i´m doing this as well. But this isn´t correct. The scope of the dump is quite different:

crash dump A crash dump is the dump of the kernel. It is done in case of a crash(kernel panic) of the system
core dump The core dump is the dump of the memory of a single process

Posted by Joerg Moellenkamp in Security at 12:45

Thursday, May 8. 2008

ZDnet about Linux and Solaris.

A really interesting article written by Jason Perlow at ZDNet: Unixfication II. It´s an insightful article about the different ways to define scalability in the Linux and Unix community and about the effects an GPLv3 Solaris may have.
So these implementations have been tested, but they are not exactly what you would call mainstream systems. And if I'm not correct, you don't currently see the level of geometric performance increases on Linux above 16 cores like you do with UNIX. The maturity in the Linux kernel for this level of enterprise performance and stability on this type of hardware just isn't there yet. The problem for Linux is, that systems with more than 16 cores are available soon (think about an Sun Fire 4600 with Quadcores) and will be quite common when AMD delivers a 12 core version of their processors.

Think about other multicore procs. The announcement of UltraSPARC T2 sound on the hardware side relatively small. Dual Socket or Quad socket isn´t such a big numerical jump. For Solaris this developement means: Jumping from 64 to 128 respectively 256 strands to schedule. And we won´t stop there.

The chance for Linux: Systems with more than 8 cores were out of reach for many developers, and systems with more than 32 cores even for many software developers. Now such machines are relatively cheap and within reach of a community backed by companies, and this may fuel the community effort to optimize Linux for larger SMP systems.

Posted by Joerg Moellenkamp in Solaris at 21:37

links for 2008-05-08

TeX Tutorial
Ein TeX Tutorial das unter anderem von einem Kollegen von mir geschrieben worden ist
(tags: TeX)

Futurama - Vergleich zwischen Original- und Synchronfassung - satre synchron
(tags: english translation)

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

Wednesday, May 7. 2008

OpenSolaris is here to stay

Sometimes you read articles like this and you think "Dear author, whatever you did smoke, don´t smoke it again". I was temped to write a harsh comment about this, but Bill Beebe already wrote a good answer in "OpenSolaris is here to stay": And that's why I now wonder if sjvn hasn't gleefully stepped into the same role that that Dan Lyons and Rob Enderle have held for so long, the Paided [sic] Shill. Except this time, instead of shilling on behalf of tSCOg against Linux, sjvn is shilling on behalf of the Linux majors against Sun and OpenSolaris.

Posted by Joerg Moellenkamp at 13:43

links for 2008-05-07

Uperf - A network performance tool
(tags: benchmark network performance tools)

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

Less known Solaris Features - the book

As my time allows it, i´m continue to work at the "Less known Solaris Features" book. You can download it here.. There will be no further announcements about the document until it reaches 1.0 . Nevertheless i will announce new versions at my Twitter page.

PS: A little bit short 500 downloads so far ...

Posted by Joerg Moellenkamp in Solaris at 12:52

Tuesday, May 6. 2008

petergabriel.com

WTF? petergabriel.com was stolen last weekend. Not the domain ... the server ....

Posted by Joerg Moellenkamp at 07:40

Monday, May 5. 2008

Sneak preview: LKSF pdfbook

It´s far from completion, it needs correction reading, some small tutorials are still missing and almost everything may change until i call this document final. But nevertheless i want you give a preview of the "Less known Solaris Features" pdfbook. 200 pages long. You can download it here.

Posted by Joerg Moellenkamp in Solaris at 22:00

links for 2008-05-05

LaTeX-Tutorial
(tags: blog tex howto german)

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

Jonathan about our results - or: The joys of GAAP

Jonathan comments the results of Q3FY2008 in his blog - Our Q3:Well, without dipping into GAAP accounting, we generated a lot of cash in the quarter (more than $320m), and getting from cash to GAAP income involves a fair number of line items associated with acquisition accounting, amortization of goodwill, tax provisions, stock option expensing - all of which, on a non-cash basis, added up to 22 cents worth of charges.
PS: Dad, i think you are correct with your opinion about our numbers ...

Posted by Joerg Moellenkamp in Sun at 10:17

opensolaris.com and Opensolaris 2008.05

There is a new website out there opensolaris.org and with this website the final release of Opensolaris 2008.05 is available. Download and try it. Additionally you will find the the torrents at the Sun Website as well.

Posted by Joerg Moellenkamp in Solaris at 09:21

Sunday, May 4. 2008

The blue screen of death

(Thanks to Kris for the link)

Posted by Joerg Moellenkamp in Fundsache at 07:41

Saturday, May 3. 2008

New features of Solaris: Alternate boot environments based on snapshots

One of the limitations of Opensolaris 2008.05 will be the missing LiveUpgrade. But ... well ... you have something better. The whole concept of LifeUpgrade was transformed into the future by using the capabilities of ZFS.
Using snapshots for boot environmentsOne of the nice features of ZFS is the fact, that you get snapshots for free. The reason lies in the copy-on-write nature of ZFS. You can freeze the filesystem by not freeing the old blocks. as new data is written is written to new blocks, you don't even have to copy the blocks (in this sense the COW of ZFS is more like a ROW ... redirect on write).

ZFS boot enables the system to work with such snapshots, as you can use one of these to boot from. You can establish multiple boot environments just by snapshoting the bootfilesystems, clonimg them and promoting them to real filesystems. This are features inherent to ZFS.

A practical exampleA warning at first: Don´t try this example without a backup of your system. Or use a test system or test VM. We will fsck up the system during this example. Okay....

I've updated my system, so i have alread two boot environments on my system:jmoekamp@glamdring:~# beadm list

```
BE           Active Active on Mountpoint Space
Name               reboot             Used
----         ------ --------- ---------- -----
opensolaris-1 yes    yes       legacy     2.31G
opensolaris   no     no        -          62.72MThis mirrors the actual state in your ZFS pools. You will find filesystems with
accordings names.NAME                                                USED  AVAIL  REFER  MOUNTPOINT
rpool                                              2.39G   142G  56.5K  /rpool
rpool@install                                         18.5K      -    55K  -
rpool/ROOT                                          2.37G   142G    18K  /rpool/ROOT
rpool/ROOT@install                                     0      -    18K  -
rpool/ROOT/opensolaris                             62.7M   142G  2.23G  legacy
rpool/ROOT/opensolaris-1                            2.31G   142G  2.24G  legacy
rpool/ROOT/opensolaris-1@install                       4.66M      -  2.22G  -
rpool/ROOT/opensolaris-1@static:-:2008-04-29-17:59:13      5.49M      -  2.23G  -
rpool/ROOT/opensolaris-1/opt                        3.60M   142G  3.60M  /opt
rpool/ROOT/opensolaris-1/opt@install                   0      -  3.60M  -
rpool/ROOT/opensolaris-1/opt@static:-:2008-04-29-17:59:13      0      -  3.60M  -
rpool/ROOT/opensolaris/opt                             0   142G  3.60M  /opt
rpool/export                                       18.9M   142G    19K  /export
rpool/export@install                                  15K      -    19K  -
rpool/export/home                                  18.9M   142G  18.9M  /export/home
rpool/export/home@install                             18K      -    21K  -After doing some configuration, you can create an
boot environment called opensolaris-baseline:
It's really easy. You just have to create a new boot environment:# beadm create -e opensolaris-1
opensolaris-baselineBut we will not work with this environment. We use it as a baseline, as a last resort when we
destroy our running environment. To run the system we will create another snapshot: # beadm create -e opensolaris-1
opensolaris-workNow let´s look into the list of our boot environments.jmoekamp@glamdring:~# beadm list

BE                 Active Active on Mountpoint Space
Name                     reboot             Used
----               ------ --------- ---------- -----
opensolaris-baseline no     no        -          53.5K
opensolaris-1        yes    yes       legacy     2.31G
opensolaris          no     no        -          62.72M
opensolaris-work     no     no        -          53.5KOkay, now we activate the opensolaris-work boot
environment:jmoekamp@glamdring:~# beadm activate opensolaris-workOkay, let´s look at the list of boot environments
again.jmoekamp@glamdring:~# beadm list
```

BE Active Active on Mountpoint Space
Name reboot Used
---- ------ --------- ---------- -----
opensolaris-baseline no no - 53.5K
opensolaris-1 yes no legacy 24.5K
opensolaris no no - 62.72M
opensolaris-work no yes - 2.31G
jmoekamp@glamdring:~#You will see that the opensolaris-1 snapshot is still active, but that the opensolaris-work will be active at the next reboot. Okay, now reboot:jmoekamp@glamdring:~# beadm list

BE Active Active on Mountpoint Space
Name reboot Used
---- ------ --------- ---------- -----
opensolaris-baseline no no - 53.5K
opensolaris-1 no no - 54.39M
opensolaris no no - 62.72M
opensolaris-work yes yes legacy 2.36GOkay, you see ... the boot environment opensolaris-work is now active and it´s activated for the next reboot (until you activate another boot environment).

Now we can reboot the system. The GRUB comes up and it will default to the opensolaris-work environment. Please remember on whicht position you find opensolaris-baseline in the boot menu. You need this position in a few moments. After a few seconds, you can log into the system and work with it.

Okay ... now let's drop the atomic bomb of administrative mishaps to your system. Log into your system, assume the root role and do the following stuff:# cd /
rm -rf *You known what happens. Depending from how fast you are able to interrupt this run to get an slightly damaged system up to a system fscked up beyond any recognition. Normaly the system would send you to the tapes now. But remember. You have some alternate boot environments.

Reboot the system, wait for the grub. You may have an garbeled output, so it's hard to read the output from the grub. Choose opensolaris-baseline. The system will boot up quite normaly.

You need a terminal window now. How you get such a terminal window depends from incurred damage. The boot environment snapshots doesn't cover the home directories. So you may have no home directory any longer. I will assume this for this example: You can get a terminal window by clicking on "Options", then "Change Session" and choose "Failsafe Terminal" there.

Okay, login via the graphical login manager, a xterm will appear. At first we delete the defunct boot environment:# beadm destroy opensolaris-work1
Are you sure you want to destroy opensolaris-work1? This action cannot be undone (y/[n]):
yOkay, now we clone the opensolaris-baseline environment to form a new opensolaris-work environment.# beadm create -e opensolaris-baseline opensolaris-workWe reactivate the opensolaris-work boot environment:# beadm activate opensolaris-workNow check, if you still have a homedirectory for your user:# ls -l /export/home/jmoekamp
/export/home/jmoekamp: No such file or directoryIf your home directory doesn't exist any longer, create a new one:# mkdir -p /export/home/jmoekamp
chown jmoekamp:staff /export/home/jmoekampNow reboot the system:# rebootWait a few moments. The system starts up. The GRUB defaults to opensolaris-work and the system starts up normaly without any problem in that condition the system had, when you create the opensolaris-baseline boot environment.# beadm list

BE Active Active on Mountpoint Space
Name reboot Used
---- ------ --------- ---------- -----
opensolaris-baseline no no - 3.18M
opensolaris-1 no no - 54.42M
opensolaris no no - 62.72M
opensolaris-work yes yes legacy 2.36GObviously you may have to recover your directory with data. It's a best practice to make snapshots of this directories on a regular schedule. So you can simply promote a snapshot to your actual version of the directory.

ConclusionYou see, this is a really neat feature. Recovering from a disaster in a minute or two. Snapshotting opens a completely new way to recover from errors. Unlike with Liveupgrade you don't need extra disks or extra partitions, and

as ZFS snapshots are really fast, creating alternate boot environments on zfs are extremly fast as well.

At the moment this feature is available on Opensolaris 2008.05 only. But with future updates it will find it´s way into Solaris as well.

Posted by Joerg Moellenkamp in Solaris at 13:40

links for 2008-05-03

Summer will bring a GPU war - The INQUIRER
(tags: GPU ATI AMD)

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

Friday, May 2. 2008

Oouch ... this hurts ...

Posted by Joerg Moellenkamp in Sun at 21:21

Hagenbecks Tropical Aquarium

Hamburg has a new object of interest since last year. The Tropical Aquarium at the Hagenbeck zoo. Really interesting. It´s fscking hot and sticky in there, but ... well ... it´s named tropical for a reason. Whenever you visit Hamburg you should visit the Zoo and the tropical aquarium as well. I made some photos there to give you some impressions.

Posted by Joerg Moellenkamp at 19:42

sudo explained

from xkcd:

(found via Juergen)

Posted by Joerg Moellenkamp in Fundsache at 11:35

Thursday, May 1. 2008

Results for Q3FY2008

Sh.. we published the numbers for the third quarter:Revenues for the third quarter of fiscal 2008 were $3.266 billion, a decrease of 0.5 percent as compared with $3.283 billion for the third quarter of fiscal 2007. Total gross margin as a percent of revenues was 44.9, an increase of 0.4 percentage points, as compared with the third quarter of fiscal 2007.

Net loss for the third quarter of fiscal 2008 on a GAAP basis was $34 million, or ($0.04) per share, as compared with net income of $67 million, or $0.07 per share, for the third quarter of fiscal 2007. In the third quarter of fiscal 2008, the company recorded a $52 million dollar tax provision, as compared to a tax benefit of $3 million in the third quarter of fiscal 2007. Net loss for the third quarter included charges related to the acquisition of MySQL, which reduced earnings per share by approximately $0.04.Not unexpected when one of your most important markets has a slight economic problem.

Posted by Joerg Moellenkamp in Sun at 22:41

Hoerempfehlung: Nils Petter Molvaer - Re-Vision

Ich hatte ja schon im Podcast gesagt, das ich teilweise einen recht merkwuerdigen Musikgeschmack habe. Nunja ... und Nils Petter Molvaer (NPM) ist so ein Fall: Industrialambientelectrojazz. Vor einigen Jahren ... ich glaube 1998 ... hatte mir eine Freundin, Aynur, eine Kassette mit seinen Erstling geschenkt (okay, auf der einen Seite war Juno Reactor und auf der anderen Herr Molvaer).

Seit dem höre ich seine Musik echt gerne ... mittlerweile ist das für mich Musik, um einfach in Träumen wegzutauchen. Oder auch um Gäste rauszuwerfen. Wie soll man die Musik beschreiben ... hmmm ... vielleicht nach der Art norwegischer Trent Reznor, der Trompete spielt und Jazz statt Industrial macht. Denn die Pfade sind ähnlich wenig ausgetreten, die Herr Molvaer beschreitet. Stellt euch Jazz vor, den man auch im Club spielen kann und zu dem man vernuenftig tanzen kann.

NPM hat nun ein neues Album mit dem Namen re-vision herausgebracht. Gut wie immer. Aber genauso gewöhnungsbeduerftig wie immer. Also wieder ungewoehnlich. Wieder die unverkennbare Handschrift (oder sollte ich sagen Trompete) des Musikers. Ich moechte hier eine Empfehlung für die CD aussprechen. Aber mehr als sonst gilt hier vorher reinhören. Die Musik ist wirklich nicht jedermanns und jederfraus sache. Zum Reinhören sind Leaps and Bounds, Violent Trip und Arctic Dub empfehlenswert. Für MPM-Einsteiger empfehle ich allerdings Kauf von Khmer. Das ist verdaubarer

Posted by Joerg Moellenkamp in Music at 19:56

A380 MSN11 in Finkenwerder

Posted by Joerg Moellenkamp in Photographie at 15:39