

Thursday, March 4, 2010

**send a break signal to a M5000 from the XSCF prompt**

reset -d 0xir

Posted by rdircio in Solaris at 19:33

Tuesday, January 19. 2010

## Solaris password expiry explained with words

If we check this "passwd -sa" output:

```
bash-3.00# passwd -sa | egrep 'rdircio|root'  
root    PS  
rdircio PS 12/21/09 14 56 14
```

It means that root's password never expires since it doesn't have any expiry rule. rdircio's password was last changed 12/21/09, the system won't let me change it but 14 days after 12/21/09, it will expire and need be changed 56 days after 12/21/09, i will be warned 14 days before it expires that i need to change it

Posted by rdircio in Solaris at 21:13

Saturday, October 17. 2009

### **quick and dirty script to obtain info from your server before you reboot**

Use the outputs of this script to compare basic health before and after you reboot, transfer the results to another host just in case the one you reboot doesn't come up :)

```
B=`uname -n`  
mkdir $B  
cd $B  
df -h > df-h.txt  
metastat -p > metastat-p.txt  
metastat -t > metastat-t.txt  
netstat -nr > netsat-nrv.txt  
ifconfig -a > ifconfig-a.txt  
zpool status -v > zpool-status-v.txt  
zfs list > zfs-list.txt  
format < /dev/null > format.txt  
prtconf -vp | grep -i wwn > wwns.txt  
zoneadm list > zoneadm.txt  
zoneadm list | while read z;do zonecfg -z $z info> ${z}.txt; done  
cd ..  
chmod 777 $B  
cp -fR $B /var/tmp
```

Posted by rdircio in Solaris at 12:53

Sunday, September 27, 2009

### **c queue max run limit reached**

If you get the message:

*c queue max run limit reached*

in your `/var/adm/messages` file and some cron jobs just don't run, you need to increase the number of cronjobs that can be run simultaeously in the `/etc/cron.d/queuedefs` file. On the line that starts with "b".

for instance:

*b.2j2n90w*

means that you can only run 2 cron jobs at the same time. Do a `man queuedefs` for details

Posted by `rdircio` in Solaris at 16:58

## Suntrunking Quick setup notes

This is a quick setup guide for suntrunking 1.3

- Install the Suntrunking package. Install the most recent suntrunking patch.
- grab our ce instance numbers, we can get these from path\_to\_inst.

```
phoenix# grep ce /etc/path_to_inst
"/sbus@3,0/SUNW,ce@0,8c00000" 0 "ce"
"/sbus@3,0/SUNW,ce@0,8c10000" 1 "ce"
"/sbus@3,0/SUNW,ce@0,8c20000" 2 "ce"
"/sbus@3,0/SUNW,ce@0,8c30000" 3 "ce"
```

- make sure the local-mac-address? setting is set to true

```
phoenix# eeprom local-mac-address?
local-mac-address?=false
phoenix# eeprom local-mac-address?=true
```

In this case we are going to choose ce0 as the head of the trunk, so we set /etc/hostname.ce0 to the hostname and make sure that

it has an entry in our host file.

There are four possible policies:

1 - MAC Default policy, at its simplest will divide the traffic equally, ie each interface in this example will handle 25% of the traffic

2 - Round Robin Use each interface in turn

3 - IP Destination Address Use IP Destination address to get the transmission path

4 - IP Source Address/IP Destination Address Deterministic connections

In our case we are going to choose round robin as our policy. So to enable this permanently we edit

/etc/opt/SUNWconn/bin/nettr.sh and add the following line

```
nettr -setup 0 device=ce policy=2 members=0,1,2,3
```

We now reboot the server

```
phoenix# init 6
INIT: New run level: 6
```

The system is coming down. Please wait.

System services are now being stopped.

Print services already stopped.

```
Sept 3 09:41:12 phoenix syslogd: going down on signal 15
```

The system is down.

syncing file systems... done

Program terminated

.....

Configuring Sun Trunking devices

ce trunk members:

ce0 (head)

ce1

ce2

ce3

configuring IPv4 interfaces: hme0 ce0.

Hostname: phoenix

.....

Now let's check the configuration

```
phoenix# ./nettr -conf
```

Key: 0; Policy: 1;

Aggr MAC address: 0:14:4f:41:f4:34

Name Original-Mac-Addr Speed Duplex Link Status

```
----
ce0 0:14:4f:41:f4:34 1000 full up enb
ce1 0:14:4f:41:f4:35 1000 full up enb
ce2 0:14:4f:41:f4:36 1000 full up enb
ce3 0:14:4f:41:f4:37 1000 full up enb
```

And let's check the load for each interface, this server is restoring data from tape and sending it to other servers, so it has mainly outbound traffic. It is very important to push traffic through the trunk in order to see how is it really operating.

phoenix# ./nettr -stats 0 interval=5

Sep 16 09:16:33 2009

Name lpkts lerrs Opkts Oerrs Collis Crc %lpkts %Opkts

```
----
ce0 6 0 26873 0 0 0 0.01 25.01
ce1 6 0 26877 0 0 0 0.01 25.01
ce2 56764 0 26862 0 0 0 99.98 25.00
ce3 0 0 26848 0 0 0 0.00 24.98
```

(Aggregate Throughput(Mb/sec): 252.00(New Peak) 282.00(Past Peak) 89.36%(N))

Posted by rdircio in Solaris at 10:50

Sunday, August 30, 2009

### changing a label from EFI to SMI

If your slice 2 is empty and you got partitions 0-9 in your disk when checkign it in "format", you have an EFI label.

If you want to change it to a SMI label, you need to use "format -e" and when prompted for the label type, choose SMI

```
partition> label
[0] SMI Label
[1] EFI Label
Specify Label type[1]: 0
Warning: This disk has an EFI label. Changing to SMI label will erase all
current partitions.
Continue? y
Auto configuration via format.dat[no]?
Auto configuration via generic SCSI-2[no]?
partition>
```

Posted by rdircio in Solaris at 12:18

## offline and online a cpu

If you want to disable a cpu in opensolaris:

```
root@cluster1:~# psradm -f 1
root@cluster1:~# psrinfo -v
Status of virtual processor 0 as of: 08/13/2009 10:10:45
on-line since 08/12/2009 04:00:42.
The i386 processor operates at 2993 MHz,
and has an i387 compatible floating point processor.
Status of virtual processor 1 as of: 08/13/2009 10:10:45
off-line since 08/13/2009 10:10:42.
The i386 processor operates at 2993 MHz,
and has an i387 compatible floating point processor.
```

And if you want to reenale it

```
root@cluster1:~# psradm -n 1
root@cluster1:~# psrinfo -v
Status of virtual processor 0 as of: 08/13/2009 10:11:35
on-line since 08/12/2009 04:00:42.
The i386 processor operates at 2993 MHz,
and has an i387 compatible floating point processor.
Status of virtual processor 1 as of: 08/13/2009 10:11:35
on-line since 08/13/2009 10:11:32.
The i386 processor operates at 2993 MHz,
and has an i387 compatible floating point processor.
```

Posted by rdircio in Solaris at 05:22

Sunday, May 31, 2009

## Extended Diagnostic Boot

1. Break the system to the open boot prom by issuing a Stop+A

2. Set the following parameters from the ok prompt:

```
setenv auto-boot? false
setenv diag-switch? true
setenv diag-level max
setenv diag-device disk
```

3. Alternately, you could set the parameters from the system level:

```
eeeprom "auto-boot?=false"
eeeprom "diag-switch?=true"
eeeprom "diag-level=max"
eeeprom "diag-device=disk"
```

4. Turn the system off. Wait about a minute, then turn it back on.

Posted by rdircio in Solaris at 23:27

Monday, May 25. 2009

## **change the mountpoint of a zfs filesystem**

```
zfs set mountpoint=/mount/point/dir mypool/myfs  
Posted by rdircio in Solaris at 10:29
```

Tuesday, February 17. 2009

## last failed/successful "sulog" entries for N days ago

You want to know for the latest N days who was able/unable to su, in solaris. Sulog is in /var/adm/sulog, and we don't have the niceties of GNU date, so we use perl inside a shell... "suentries.ksh"

```
#!/bin/ksh
#-- n has how many days ago we want
n=$1
today=`usr/bin/perl -e 'printf "%d\n", time;`
x=$n
while [ $x -gt -1 ];do
ago=$(( $today-86400*$x))
export ago
DAY=`perl -e '($a,$b,$c,$d,$e,$f,$g,$h,$i)=localtime($ENV{ago});printf"%02d/%02d\n", $e+1,$d`
echo "---- $DAY"
cat /var/adm/sulog | grep $DAY
x=$(( $x-1))
done
```

This would print, for 2 days ago:

```
bash-3.00# date
Wednesday, February 18, 2009 5:45:42 AM GMT
```

```
bash-3.00# ./suentries.ksh 2
---- 02/16
SU 02/16 14:19 + pts/1 someuser-oracle
---- 02/17
---- 02/18
SU 02/18 03:30 + pts/1 rdircio-root
SU 02/18 04:58 + pts/1 rdircio-root
```

Posted by rdircio in Solaris at 23:42

Friday, February 13. 2009

## swap in solaris

You want to know how much swap is configured in your server in gb.

swap -l tells us 512byte blocks, if you have more than one swap chunk then you need to sum and convert

```
bash-3.00# swap -l
swapfile      dev swaplo blocks free
/dev/md/dsk/d1 85,1   16 16780208 16627776
/dev/md/dsk/d2 86,1   16 26780208 16111111
```

You can do a ksh to do that, swap.ksh:

```
#!/bin/ksh
blockstotal=0 ; swapentries=$((`swap -l | wc -l`-1)) ; while [ $swapentries -ne 0 ] ; do blocks=`swap -l | awk
{'print $4'} | head -${(swapentries+1)} | tail -1` ; blockstotal=$((blockstotal+blocks)) ;
swapentries=$((swapentries - 1)) ; done ; echo $((blockstotal*512/1024/1024/1024))Gb
```

Thx to Hans Beiskjaer

Posted by rdircio in Solaris at 12:36

## socktest

Have you ever wanted to test a connection to only one port, and found nmap overkill, and you like to do it in C?

Normally I'd recommend something like

```
echo "~." | telnet -r $H $X 2>&1 | grep -i connected | wc -l | sed 's//g'
```

But, if the host doesn't answer the default tcp/ip timeout for a socket is 3 minutes!!!! You have to wait long time...

You don't want to wait too long for that, so you can use a program to test a socket connection, like socktest.c, it tries to connect and if it can't within 5 seconds it will bail out.

you should compile with gcc, if using solaris "gcc -o socktest -lsocket -lnsl socktest.c" will suffice.

So, for instance you want to test port 80 and port 44 in google.com:

```
bash-3.00$ ./socktest google.com 80 ; echo $?
Connection to port successful
bash-3.00$ ./socktest google.com 44 ; echo $?
ERROR: Timeout or error() 4 - Interrupted system call
1
```

You can get the exit code 0 as successful, 1 as not successful.

Source here:

```
#include <stdio.h>
#include <sys/types.h>
#include <sys/socket.h>
#include <netinet/in.h>
#include <netdb.h>
#include <errno.h>
#include <sys/fcntl.h>
//-----
// Roberto Dircio Palacios-Macedo
// socktest - program to test a connection
// to a host in a specific port
// usage - ./socktest <host> <port>
//-----
void error(char *msg)
{
    perror(msg);
    exit(0);
}
int main(int argc, char *argv[])
```

```

{
    int res, valopt;
    struct timeval tv;
    fd_set myset;
    socklen_t lon;
    int soc, portno, n;
    struct sockaddr_in serv_addr;
    struct hostent *server;
    char buffer[256];
    if (argc < 3) {
        fprintf(stderr, "usage %s hostname port\n", argv[0]);
        exit(1);
    }
    portno = atoi(argv[2]);
    server = gethostbyname(argv[1]);
    soc = socket(AF_INET, SOCK_STREAM, 0);
    if (soc < 0) {
        error("ERROR: opening socket");
        exit(1);
    }
    if (server == NULL) {
        fprintf(stderr, "ERROR: no such host\n");
        exit(1);
    }
    bzero((char *) &serv_addr, sizeof(serv_addr));
    serv_addr.sin_family = AF_INET;
    bcopy((char *)server->h_addr,
        (char *)&serv_addr.sin_addr.s_addr,
        server->h_length);
    serv_addr.sin_port = htons(portno);
    long arg = fcntl(soc, F_GETFL, NULL);
    arg |= O_NONBLOCK;
    fcntl(soc, F_SETFL, arg);
    res=connect(soc,&serv_addr,sizeof(serv_addr));
    if (res < 0) {
        if (errno == EINPROGRESS) {
            tv.tv_sec = 5;
            tv.tv_usec = 0;
            FD_ZERO(&myset);
            FD_SET(soc, &myset);
            if (select(soc+1, NULL, &myset, NULL, &tv) > 0) {
                lon = sizeof(int);
                getsockopt(soc, SOL_SOCKET, SO_ERROR, (void*)&valopt, &lon);
                if (valopt) {
                    fprintf(stderr, "ERROR: Error in connection() %d - %s\n", valopt, strerror(valopt));
                    exit(1);
                }
            }
        }
        else {
            fprintf(stderr, "ERROR: Timeout or error() %d - %s\n", valopt, strerror(valopt));
            exit(1);
        }
    }
}

```

```
else {
    fprintf(stderr, "ERROR: Error connecting %d - %s\n", errno, strerror(errno));
    exit(1);
}
}
arg = fcntl(soc, F_GETFL, NULL);
arg &= (~O_NONBLOCK);
fcntl(soc, F_SETFL, arg);
fprintf(stdout, "Connection to port successful\n");
exit(0);
}
```

Enjoy!

Posted by rdircio in Solaris at 22:04

Thursday, November 6, 2008

**sshd[419]: error: Failed to allocate internet-domain X11 display socket.**

If you cannot get ssh X11 forwarding to work even when you configured your sshd\_config properly, and you get the message "sshd[419]: error: Failed to allocate internet-domain X11 display socket." in your console... you need to disable openssh's ipv6 mode.

add this to your sshd\_config file:

```
ListenAddress 0.0.0.0
```

And change your init script to pass the "-4" option to sshd, as:

```
case $1 in
'start')
    create_key $SSHDIR/ssh_host_rsa_key rsa
    create_key $SSHDIR/ssh_host_dsa_key dsa
    /opt/acs/sbin/sshd -4 &
    ;;
```

Restart sshd and you'll have X11Forwarding working

Posted by rdircio in Solaris at 21:04

Monday, October 13. 2008

## **Xerox WorkCentre Pro 7345 on latest cups (1.2.7) from blastwave**

Get the xerox driver bundle.

Copy all the ppd files to the cups filter directory(/opt/csw/share/cups/model/), reconfigure cups to use the new ppd (xrx7345.ppd)

The xerox workcentre pro 7345 is basically a Postscript printer, but we used this PPD as it came directly from Xerox

Posted by rdircio in Solaris at 15:00

### **remove filesystem from Zone**

First do a dump of your config, in case you want to get it back. On the global zone:

```
# zonecfg -z zzzpit120b export -f zzzpit120b.txt
```

Now remove it from the config, and umount it, all this in the global zone.

```
# zonecfg -z zzzpit120b remove fs dir=/oracle/GYT/saptrace  
# umount /zones/zzzpit120b/root/oracle/GYT/saptrace
```

Thanks to Alejandro Siller

Posted by rdircio in Solaris at 13:58

Wednesday, October 1. 2008

**find files that have grown in the latest 5 days, and print them sorted by size**

```
# find . -mount -mtime -5 -exec du -sk {} \; | sort -n
```

Posted by rdircio in Solaris at 09:48

Tuesday, September 30. 2008

## Add a network printer in solaris 9

```
bash-2.05# lpadmin -p chubbchkprt -v /dev/null  
bash-2.05# lpadmin -p chubbchkprt -m netstandard  
bash-2.05# lpadmin -p chubbchkprt -o dest=10.12.7.250:9100 -o protocol=TCP  
bash-2.05# lpadmin -p chubbchkprt -l any  
bash-2.05# accept chubbchkprt
```

```
destination "chubbchkprt" now accepting requests  
bash-2.05# /usr/bin/enable chubbchkprt  
printer "chubbchkprt" now enabled  
new printer  
bash-2.05# lpstat -p chubbchkprt  
printer chubbchkprt is idle. enabled since Tue Sep 30 19:57:42 2008. available.
```

Posted by rdircio in Solaris at 18:55

Wednesday, September 17. 2008

### **latest freenx+nx for OpenSolaris Nevada x86**

Get it at [http://www.opensolaris.org/os/project/awards/awards\\_land/Entries/](http://www.opensolaris.org/os/project/awards/awards_land/Entries/)

I would love it to be part of opensolaris.

:)

UPDATE: you can also get it [here](#).

Posted by rdircio in Solaris at 08:33

Monday, September 15, 2008

## **Disable NIS client in solaris**

```
mv /a/usr/lib/netsvc/yp/ypbind /a/usr/lib/netsvc/yp/ypbind.orig
```

```
mv /a/usr/lib/netsvc/yp/ypserv /a/usr/lib/netsvc/yp/ypserv.orig
```

```
init 6
```

Posted by rdircio in Solaris at 07:34

Sunday, July 13. 2008

## Calculate ETA of a file transfer

You are transferring a 7.2Gb file using scp on a wan, you want to know at some moment how much more time you will have to wait until it is transferred, i wrote a script to give an average, it must be run on the file receiving server.

```
#!/bin/ksh
FS=7343108 #--- the filesize
f=/app/db/oracle/u01/oracle/rmanbackupsets/v9jlbbu9_1_3 #--- the file that is growing..
s1=`du -s $f |awk '{ print $1 }'`
echo "...sleeping 10"
sleep 10
s2=`du -s $f | awk '{ print $1 }'`
delta=`echo "${s2} - ${s1}" | bc`
echo "Bytes transferred in 10 secs : $delta"
REMAINING=`echo "$FS - $s2" | bc`
echo "Original Filesize : $FS"
echo "Remaining : $REMAINING"
ETA=`echo "scale=3;(($REMAINING / $delta) / 6)/60" | bc`
echo "Hours remaining for transfer $ETA"
du -sh $f
```

Posted by rdircio in Solaris at 20:16

Friday, June 20. 2008

## Add a couple of disks to your zpool and grow some filesystems

First, check the status of your pool

```
bash-3.00# zpool status -v
```

```
pool: amx002zpool1
```

```
state: ONLINE
```

```
scrub: none requested
```

```
config:
```

NAME	STATE	READ	WRITE	CKSUM
amx002zpool1	ONLINE	0	0	0
c6t60060480000190101353533030393543d0	ONLINE	0	0	0
c6t60060480000190101353533030393536d0	ONLINE	0	0	0
c6t60060480000190101353533030393530d0	ONLINE	0	0	0
c6t60060480000190101353533030393441d0	ONLINE	0	0	0
c6t60060480000190101353533030393434d0	ONLINE	0	0	0
c6t60060480000190101353533030433435d0	ONLINE	0	0	0

```
errors: No known data errors
```

Then add a couple of disks to it:

```
bash-3.00# zpool add amx002zpool1 c6t60060480000190101353533031343236d0
```

```
bash-3.00# zpool add amx002zpool1 c6t60060480000190101353533031343243d0
```

Now check if they're added:

```
bash-3.00# zpool status -v
```

```
pool: amx002zpool1
```

```
state: ONLINE
```

```
scrub: none requested
```

```
config:
```

NAME	STATE	READ	WRITE	CKSUM
amx002zpool1	ONLINE	0	0	0
c6t60060480000190101353533030393543d0	ONLINE	0	0	0
c6t60060480000190101353533030393536d0	ONLINE	0	0	0
c6t60060480000190101353533030393530d0	ONLINE	0	0	0
c6t60060480000190101353533030393441d0	ONLINE	0	0	0
c6t60060480000190101353533030393434d0	ONLINE	0	0	0
c6t60060480000190101353533030433435d0	ONLINE	0	0	0
c6t60060480000190101353533031343236d0	ONLINE	0	0	0
c6t60060480000190101353533031343243d0	ONLINE	0	0	0

```
errors: No known data errors
```

List free space on the pool:

```
bash-3.00# zpool list
```

NAME	SIZE	USED	AVAIL	CAP	HEALTH	ALTROOT
------	------	------	-------	-----	--------	---------

```

gskhr009zpool1      69.5G  5.87G  63.6G   8% ONLINE  -
gskhr009zpool2      139G   73.6G  65.4G  52% ONLINE  -
gskhr009zpool3      1.70T  1.24T  468G   73% ONLINE  -

```

List all the filesystems in the pool:

```

bash-3.00# zfs list
NAME                USED AVAIL REFER MOUNTPOINT
amx002zpool1        345G  157G  24.5K  /amx002zpool1
amx002zpool1/amx21   64.6G  6.37G  64.6G  /amx21
amx002zpool1/amx22   32.9G  4.12G  32.9G  /amx22
amx002zpool1/amx23   26.1G  4.95G  26.1G  /amx23
amx002zpool1/amx24   58.9G  6.15G  58.9G  /amx24
amx002zpool1/amx25   49.6G  7.39G  49.6G  /amx25
amx002zpool1/oracle   6.50G  1.50G  6.50G  /oracle
amx002zpool1/oracle-export 21.6G  23.4G  21.6G  /oracle/export
amx002zpool1/oracle2  24.5K  8.00G  24.5K  /opt/oracle
amx002zpool1/psreports 24.5K  3.00G  24.5K  /opt/reports
amx002zpool1/vendor   690K  20.0G  690K  /opt/vendor

```

Get a quota and reservation for the filesystem to grow:

```

bash-3.00# zfs get quota amx002zpool1/amx21
NAME                PROPERTY VALUE          SOURCE
amx002zpool1/amx21 quota      71G             local
bash-3.00# zfs get reservation amx002zpool1/amx21
NAME                PROPERTY VALUE          SOURCE
amx002zpool1/amx21 reservation 71G             local

```

Grow the filesystem:

```

bash-3.00# zfs set quota=100G amx002zpool1/amx21
bash-3.00# zfs set reservation=100G amx002zpool1/amx21

```

That's it, have fun

Sunday, June 8. 2008

## **Add a VCS hagui admin user to your cluster**

To add an hagui admin user you do:

```
# haconf -makerw
```

```
# hauser -add rdircio
```

```
# hauser -addpriv rdircio Administrator
```

```
# haconf -dump -makero
```

Posted by rdircio in Solaris at 13:56

Saturday, June 7. 2008

### **Jumpstart nic comes up on half duplex**

if your jumpstart nic comes up on half duplex you can send it

```
ok% boot net:speed=1000,duplex=full
```

And that will force the speed

Thanks to Ranganathan Kuppusamy

Posted by rdircio in Solaris at 12:49

### **Using dlpiping to test an llt link on a veritas cluster**

if an llt link shows down, you can always inject some traffic through the wire and see if you can catch it on the other side.

We first, on one side of the cluster find out the mac address of our interface, then start a "dlpiping server" process on it.

```
# /opt/VRTSllt/getmac /dev/ce:0  
/dev/ce:0 08:00:10:A7:1E:DC  
# /opt/VRTSllt/dlpiping -s /dev/ce:0
```

Now, on the other side of the link, on the other server, we query for the mac address in the wire.

```
# /opt/VRTSllt/dlpiping -c /dev/ce:0 08:00:10:A7:1E:DC  
08:00:10:A7:1E:DC is alive
```

if it says that the mac is alive then we have a good link.

Posted by rdircio in Solaris at 06:40

Monday, May 19, 2008

## **Rename powerpath devices**

Perhaps you want to use your own nomenclature, so the device names mean something to you.

You can do it this way:

```
# emcpadm renamepseudo -s emcpowerh -t emcpowerrh
```

This renames the device called emcpowerh, to be now called emcpowerrh

Posted by rdircio in Solaris at 14:09

Friday, May 9. 2008

## To wipe out all disks

Sometimes you need to delete all info in your disks, perhaps because it contains private data. You can effectively do this by erasing the first 1024 blocks on the disk. After doing this the disk will lose its label and the partition schema. You can run the next script to obtain the list of commands that will do it on your server:

```
#!/bin/ksh
format < /dev/null | egrep "cyl" | while read l;do
D=`echo $l | awk '{ print "dd if=/dev/zero of=/dev/dsk/"$2"s2 count=1024" }'`
echo $D
done
```

Thanks to Alejandro Siller

Posted by rdircio in Solaris at 22:32

Monday, April 28. 2008

### **Get tape drive information from NetBackup**

`/usr/opensv/volmgr/bin/scan`

will tell you detailed info about your tape drives

Thanks to Alberto Vazquez

Posted by rdircio in Solaris at 11:34

## Calculate all your interfaces' netmasks

This script for Solaris prints these values,

- 1- the interface name
- 2- MAC Address
- 3- IPMP Group (if any)
- 4- IP
- 5- netmask in hex
- 6- netmask in decimal
- 7- netmask length

```
bash-3.00$ ./netmask.ksh
lo0,,,127.0.0.1,ff000000,255.0.0.0,8
ce0,0:3:ba:84:dd:a5,management,198.148.129.96,ffffffc0,255.255.255.192,26
ce0:1,,,198.148.129.110,ffffffc0,255.255.255.192,26
ce0:2,,,198.148.129.98,ffffffc0,255.255.255.192,26
ce3,0:3:ba:84:dd:a8,,192.168.5.11,fffff00,255.255.255.0,24
ce3:1,,,192.168.130.254,ffff000,255.255.240.0,20
qfe0,0:3:ba:4:45:4,management,198.148.129.97,ffffffc0,255.255.255.192,26
qfe3,0:3:ba:4:45:7,,192.168.132.254,ffff000,255.255.240.0,20
```

This is the code for netmask.ksh, notice that we also have an unused function, Len2Mask, which is very useful to calculate a netmask in decimal when you're given the netmask length, like "24", or "16":

```
#!/bin/ksh
IFCONFIG=/usr/sbin/ifconfig
ECHO=/usr/ucb/echo
integer G=0
raiseP(){
  x=$1
  y=$2
  integer total=1
  integer j=0
  while ((j < y));
  do
    (( total*=x ))
    (( j = j + 1 ))
  done
  $ECHO $total
  return $total
}
L2M(){
  integer nmask=$1
  if [ nmask -lt 1 ];then
    $ECHO -n "0"
  return "0"
fi
```

```

integer ncalc=0
integer x=7
while [ $x -ge 0 ];do
integer P=`raiseP 2 $x`
(( ncalc= ncalc + P ));
(( nmask=nmask - 1 ));
G=$nmask
if [ $nmask -lt 1 ];then
$ECHO -n $ncalc
return $ncalc
fi
(( x= x - 1 ));
done
$ECHO -n $ncalc
return $ncalc
}
Len2Mask(){
L=$1
if [ $L -lt 0 -o $L -gt 32 ];then
$ECHO "Your mask length can only be 0 - 32"
break
fi
L2M $L
$ECHO -n "."
L2M $G
$ECHO -n "."
L2M $G
$ECHO -n "."
L2M $G
}
Mask2Len(){
M=$1
m[0]=`$ECHO $M | awk -F"." '{ print $1 }`
m[1]=`$ECHO $M | awk -F"." '{ print $2 }`
m[2]=`$ECHO $M | awk -F"." '{ print $3 }`
m[3]=`$ECHO $M | awk -F"." '{ print $4 }`
loop=0
mask=0
while [ $loop -lt 4 ];do
div=256
while [ $div -gt 1 ];do
let "div=$div / 2"
let "test=${m[$loop]} - div"
if [ ! $test -lt 0 ];then
let "mask=$mask + 1"
m[$loop]=$test
else
break
fi
done
let "loop=$loop+1"
done
$ECHO "$mask"

```

```
}  
$IFCONFIG -a | grep ":" | egrep -v "ether" | awk '{ print $1 }' | sed 's/^(.*):/\1/' | while read i; do  
F=/tmp/${i}.$$  
$IFCONFIG $i > $F  
I=`cat $F | egrep -v ":|group" | awk '{ print $2 }`  
E=`cat $F | egrep "ether" | awk '{ print $2 }`  
GP=`cat $F | egrep "groupname" | awk '{ print $2 }`  
NMH=`cat $F | grep -v ":" | awk '{ print $4 }`  
rm $F  
TMP=`$ECHO $NMH | sed 's/./ 0x&/g`  
NMD=`printf "%d.%d.%d.%d\n" $TMP`  
M2L=`Mask2Len $NMD`  
$ECHO "$i,$E,$GP,$I,$NMH,$NMD,$M2L"  
done
```

Posted by rdircio in Solaris at 02:53

Friday, March 28. 2008

### **Veritas vxvm version**

In order to know your veritas volume manager version do a:

```
# pkginfo -l VRTSvxvm | grep VERSION  
VERSION: 3.5,REV=06.21.2002.23.14
```

Posted by rdircio in Solaris at 00:17

Thursday, March 27. 2008

## Physical PCI card position in sunfire E12k/E15k/E20k/E25k

if you need to know where, for instance is "qfe0" in your e25k, so someone can connect a cable there... you can use sun doc #202571, which i have transformed to a script called "pos.ksh":

```
#!/bin/ksh
echo "Expander Slot Device"
grep pci /etc/path_to_inst |while read nic;do
A=`echo $nic | awk -F"@ " '{ print $2 }' | awk -F"," '{print $1}'`
O=`echo $nic | awk -F"," '{print $2}'| awk -F"/" '{ print $1 }'`
T1=`grep " $A" t1.txt`
E=`echo $T1 | awk '{ print $2 }`
C=`echo $T1 | awk '{ print $1 }`
S=`grep "^$C...$O" t2.txt | awk '{ print $2 }`
echo "$E $S $nic"
done
```

You will need t1.txt:

```
IOC Expander AgentID
0 0 1c(28)
1 0 1d(29)
0 1 3c(60)
1 1 3d(61)
0 2 5c(92)
1 2 5d(93)
0 3 7c(124)
1 3 7d(125)
0 4 9c(156)
1 4 9d(157)
0 5 bc(188)
1 5 bd(189)
0 6 dc(220)
1 6 dd(221)
0 7 fc(252)
1 7 fd(253)
0 8 11c(284)
1 8 11d(285)
0 9 13c(316)
1 9 13d(317)
0 10 15c(348)
1 10 15d(349)
0 11 17c(380)
1 11 17d(381)
0 12 19c(412)
1 12 19d(413)
0 13 1bc(444)
```

```

1 13 1bd(445)
0 14 1dc(476)
1 14 1dd(477)
0 15 1fc(508)
1 15 1fd(509)
0 16 21c(540)
1 16 21d(541)
0 17 23c(572)
1 17 23d(573)

```

and t2.txt

```

IOC Slot Offset
0 0 600000
0 1 700000
1 2 600000
1 3 700000

```

So basically, inside any given domain you can do:

```

#./pos.ksh > slotmap.txt

```

then, if you want to know where your Quad Fast Ethernet "qfe" cards are, you can do:


```

bash-2.03# egrep "qfe|Expander" slotmap.txt
Expander Slot Device
2 3 "/pci@5d,700000/pci@1/SUNW,qfe@0,1" 0 "qfe"
2 3 "/pci@5d,700000/pci@1/SUNW,qfe@1,1" 1 "qfe"
2 3 "/pci@5d,700000/pci@1/SUNW,qfe@2,1" 2 "qfe"
2 3 "/pci@5d,700000/pci@1/SUNW,qfe@3,1" 3 "qfe"
17 3 "/pci@23d,700000/pci@1/SUNW,qfe@0,1" 4 "qfe"
17 3 "/pci@23d,700000/pci@1/SUNW,qfe@1,1" 5 "qfe"
17 3 "/pci@23d,700000/pci@1/SUNW,qfe@2,1" 6 "qfe"
17 3 "/pci@23d,700000/pci@1/SUNW,qfe@3,1" 7 "qfe"

```

And you'll find out you have 2 quads, one is at IOBoard 2 slot 3, the other is at IOBoard 17 slot 3

A typical IOBoard has 4 slots, numbered from 0-3,



So Both cards are at the upper left, one on IOBoard 2 the other on IOBoard 17

Qfe0 is at IOBoard 2 slot 3

Posted by rdircio in Solaris at 21:11

Wednesday, March 26. 2008

## Kernel Async IO

In order to know if your OS has async io enabled do a

```
bash-2.05# modinfo | grep -i aio
141 7828e000 4916 178 1 kaio (kernel Async I/O)
141 7828e000 4916 178 1 kaio (kernel Async I/O for 32 bit com)
```

if the kaio modules are loaded, you have async io enabled

Posted by rdircio in Solaris at 14:02

Tuesday, March 25. 2008

## Quick and dirty script to query NIC speed

I already have a script like this in "miniexplorer", but it generates html and that may be overkill, so this is a quickie one to query all network card ports in your solaris server about NIC speed.

```
#!/bin/ksh
ifconfig -a | grep flags | grep -v 'lo' |awk -F: '{ print $1 }' |uniq|while read IF;do
  D=`echo $IF| sed s/[0-9,:]*$//`
  IN=`echo $IF| sed 's/^[a-z]*//`
  echo "$D $IN -----"
  ndd -set /dev/$D instance $IN
  (ndd -get /dev/$D \?|awk '$1!~/\?/{print$1}'|while read p;do
    /usr/ucb/echo -n $p;;ndd -get /dev/$D $p;
  done
  kstat ${D}:${IN}) | egrep 'adv|link|speed' |sed 's/ //g'
done
```

Posted by rdircio in Solaris at 20:09

Thursday, March 6, 2008

## A Netra T1 SPARC running Windows98 using Qemu

I've been running Microsoft Windows98 inside a Sparc system, using QEMU and running Solaris 10 as the host system. The speed is actually usable.

Get the package from

<http://www.martux.org/qemu/RELEASES/sparc/SUNWqemu-0.8.2,REV=2006.10.14-sol8-sparc-opt.pkg.gz>

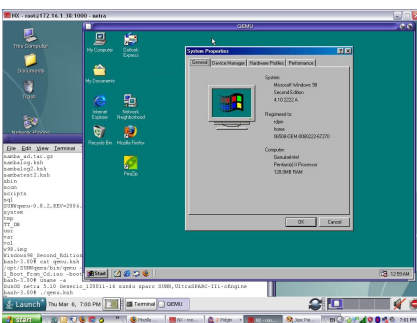
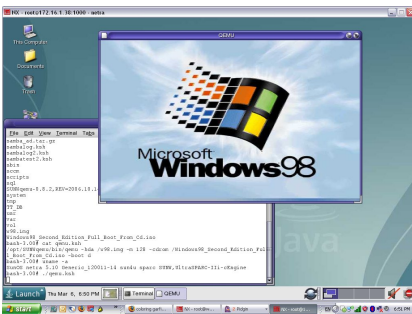
it'll end up at /opt/SUNWqemu, to install windows you can use a iso image of your Windows98 cd and create the virtual disk that will contain the OS as:

```
# /opt/SUNWqemu/bin/qemu-img create /w98.img 1000M
```

Then start the qemu emulator as:

```
# /opt/SUNWqemu/bin/qemu -hda /w98.img -m 128 -cdrom  
/Windows98_Second_Edition_Full_Boot_From_Cd.iso -boot d
```

This will boot it to the Windows98 install iso image so you can install the OS.



Posted by rdircio in Solaris at 23:35

Friday, February 29. 2008

## Disk contention

To see if we have disk contention problems we'd take a look at an iostat. if asvc\_t >50 for sustained periods, the system has disk contention problems.

```
# iostat -nx 3 1000 | egrep -v 'rmt|0.0 0.0 0.0 0.0 0.0 0.0 0.0 0.0'
```

### *extended device statistics*

<i>r/s</i>	<i>w/s</i>	<i>kr/s</i>	<i>kw/s</i>	<i>wait</i>	<i>actv</i>	<i>wsvc_t</i>	<i>asvc_t</i>	<i>%w</i>	<i>%b</i>	<i>device</i>
0.0	1.3	0.0	5.7	0.0	0.0	0.0	1.5	0	0	d0
0.0	10.3	0.0	47.7	0.0	0.0	0.0	0.6	0	1	d1
0.0	2.3	0.0	11.2	0.0	0.0	0.0	0.9	0	0	d2
0.0	0.3	0.0	0.3	0.0	0.0	0.0	0.9	0	0	d4
0.0	0.3	0.0	0.3	0.0	0.0	0.0	0.9	0	0	c2t0d15
0.0	0.3	0.0	2.7	0.0	0.0	0.0	0.7	0	0	c2t0d12
0.0	1.0	0.0	0.5	0.0	0.0	0.0	0.6	0	0	c2t0d10
0.0	5.7	0.0	45.3	0.0	0.0	0.0	0.7	0	0	c2t0d7
0.0	1.3	0.0	5.7	0.0	0.0	0.0	1.4	0	0	c2t0d3
0.0	1.0	0.0	8.0	0.0	0.0	0.0	1.3	0	0	c2t0d2
0.0	4.7	0.0	2.3	0.0	0.0	0.0	0.5	0	0	c2t0d1

### *extended device statistics*

<i>r/s</i>	<i>w/s</i>	<i>kr/s</i>	<i>kw/s</i>	<i>wait</i>	<i>actv</i>	<i>wsvc_t</i>	<i>asvc_t</i>	<i>%w</i>	<i>%b</i>	<i>device</i>
0.0	1.3	0.0	5.7	0.0	0.0	0.0	0.9	0	0	d0
0.0	6.3	0.0	30.7	0.0	0.0	0.0	0.8	0	0	d1
0.0	1.3	0.0	5.7	0.0	0.0	0.0	0.7	0	0	d2
0.0	0.7	0.0	0.3	0.0	0.0	0.0	0.6	0	0	c2t0d10
0.0	3.7	0.0	29.3	0.0	0.0	0.0	1.0	0	0	c2t0d7
0.0	1.3	0.0	5.7	0.0	0.0	0.0	0.9	0	0	c2t0d3
0.0	0.7	0.0	5.3	0.0	0.0	0.0	0.8	0	0	c2t0d2
0.0	2.7	0.0	1.3	0.0	0.0	0.0	0.5	0	0	c2t0d1

This system doesn't seem to.

Posted by rdircio in Solaris at 16:56

Wednesday, February 27. 2008

## **gawk construct to build a line from each 3 lines**

We got the output from "sar -ruq" in a file, we stored in sar2.txt, and it looks like this:

```
11:44:51 633773 22954346
4 1 3 93
0.0 0 0.0 0
11:44:56 633773 22954288
2 2 1 94
0.0 0 0.0 0
11:45:01 632206 22786120
13 18 13 56
9.0 20 0.0 0
11:45:06 625509 22462685
53 47 0 0
5.0 99 0.0 0
```

So we want to plot it in excel, we need one line per sample, much like this:

```
11:45:06 625509 22462685 53 47 0 0 5.0 99 0.0 0
11:45:11 628749 22770298 59 41 0 0 4.8 100 0.0 0
11:45:16 630432 22733837 55 45 0 0 6.8 100 0.0 0
```

So we wrote a small gawk expression to do it:

```
gawk 'ORS=NR%3?" ":"\n"' sar2.txt
```

So simple, we're using the modulus function here to do the job. Coming to a blog near you!

Posted by rdircio in Solaris at 19:58

## **kill orphans, defuncts or zombies**

Mhhh, it is not halloween, and i'm not a sadist.

When you see things like this on a "ps -ef":

```
root 29988 29973 0 17:13 pts/1 00:00:00 <defunct>
```

You can kill them on solaris >=9 with the command "preap"

So in this example you would issue:

```
# preap 29988
```

And it will be dead. Note this only works in solaris ;)

Posted by rdircio in Solaris at 17:12

Tuesday, February 26. 2008

## Enable sar system accounting on solaris8 and 9

To enable sysstat's sar to collect stats in the background:

```
# ln -s /etc/init.d/perf /etc/rc2.d/S21perf
# /etc/rc2.d/S21perf start
```

Use "crontab -e sys" to make sure sys crontab looks like:

```
0 * * * 0-6 /usr/lib/sa/sa1
20,40 8-17 * * 1-5 /usr/lib/sa/sa1
5 18 * * 1-5 /usr/lib/sa/sa2 -s 8:00 -e 18:01 -i 1200 -A
```

To verify, issue "sar -A" or just "sar"

Posted by rdircio in Solaris at 12:29

Friday, February 15. 2008

## Sar -ruq

Still on the performance+sysstat topic, I found useful to use "sar -ruq" to obtain memory, cpu and runqueue stats in one command

```
root@kraftek:~# sar -ruq 3 10000
Linux 2.6.21.5-smp (kraftek) 02/15/2008
06:51:28 PM CPU %user %nice %system %iowait %steal %idle
06:51:31 PM all 97.01 0.00 2.99 0.00 0.00 0.00
06:51:28 PM kbmemfree kbmemused %memused kbbuffers kbcached kbswpfree kbswpused %swpused
kbswpcad
06:51:31 PM 32528 866560 96.38 42564 178940 0 0 0.00 0
06:51:28 PM runq-sz plist-sz ldavg-1 ldavg-5 ldavg-15
06:51:31 PM 1 256 1.85 2.01 1.37
06:51:31 PM CPU %user %nice %system %iowait %steal %idle
06:51:34 PM all 96.67 0.00 3.33 0.00 0.00 0.00
06:51:31 PM kbmemfree kbmemused %memused kbbuffers kbcached kbswpfree kbswpused %swpused
kbswpcad
06:51:34 PM 32156 866932 96.42 42624 178940 0 0 0.00 0
06:51:31 PM runq-sz plist-sz ldavg-1 ldavg-5 ldavg-15
06:51:34 PM 1 256 1.85 2.01 1.37
```

Posted by rdircio in Solaris at 20:41

Wednesday, February 6. 2008

## get a table with HBA info without fcinfo

We want a table that maps cX to its wwn and to its device path, like

Controller	WW port name	WW node name	Device Path
c1	p:210100e08b27e1c7	n:200100e08b27e1c7	/pci@9,700000/pci@4/SUNW,qlc@4

So we wrote a script for it, hbainfo.ksh:

```
#!/bin/ksh
prtpicl -v -c scsi-fcp | egrep 'ww|devfs' | sed 's/:node-wwn//g;s/:port-wwn//g;s/:devfs-path//g;s/ //g' | nawk
'ORS=NR%3?" ":"\n" > /tmp/piclmap.$$
(for c in `cfgadm -al | grep "fc-" | awk '{ print $1 }';do
  /usr/ucb/echo -n $c;
  luxadm -e dump_map /dev/cfg/${c} | grep Adap |while read l;do
    for w in `echo $l`;do
      WC=`echo $w|wc -c`
      if [ $WC -eq 17 ];then
        /usr/ucb/echo -n "$w "
      fi
    done
  done
  echo ""
done
done
) | while read l;do
  W=`echo $l |awk '{ print $2 }'`
  L=`echo "$l " | awk '{ print $1 " p:" $2 " n:"$3 }'`
  /usr/ucb/echo -n "$L "
  grep $W /tmp/piclmap.$$ | awk '{ print $3 }'
done
rm /tmp/piclmap.$$
```

Thanks to Hiram Ruiz for the commands

Posted by rdircio in Solaris at 12:46

## Single Sign On + AD Authenticated Samba on Solaris 10 Sparc to a Windows 2003 Active Directory PDC

To add to the previous post, we'll show now a way to ssh to your Solaris 10 SPARC server using an ActiveDirectory user+pass, additional to use samba shares anywhere, with your AD credentials.

### BUILDING

Samba from blastwave does not include pam\_winbind.so, to plug it to PAM, and use our AD user+pass to login to solaris, so we need to build our own samba. The release i got is samba 3.0.28, from <http://us4.samba.org/samba/ftp/stable/samba-3.0.28.tar.gz>

We need to make sure all required packages to build and use SAMBA are in the system, so please obtain [this list](#) using pkg-get from blastwave.

Our SAMBA sources like GNU ld, so we are not going to use solaris native ld to build it, we do so by doing:

```
# mv /usr/ccs/bin/ld /usr/ccs/bin/ld.old
# ln -s /opt/csw/bin/gld /opt/csw/bin/ld
# export PATH=$PATH:/opt/csw/bin
```

We uncompress SAMBA in a directory called "/builds" and then:

```
# cd /builds/samba-3.0.28/source
# ./configure --with-ldap --with-ads --with-krb5=/opt/csw --with-winbind --with-pam --prefix=/usr/samba_ad
```

This will use the openldap client libs in /opt/csw, same as the krb5 in /opt/csw. If all reqs are passed do a

```
# make && make install
```

Then, it will all end at /usr/samba\_ad

After building these you can always remove the gcc suite, automake, autoconf, flex and bison.

We now need to copy the nsswitch and pam extensions to their locations:

```
# cp /builds/samba-3.0.28/source/nsswitch/libnss_win* /usr/lib
# ln -s /usr/lib/libnss_winbind.so /usr/lib/libnss_winbind.so.1
# ln -s /usr/lib/libnss_wins.so /usr/lib/libnss_wins.so.1
```

Now we copy the pam extension:

```
# cp /builds/samba-3.0.28/source/bin/pam_winbind.so /usr/lib/security
# cp /builds/samba-3.0.28/source/bin/pam_winbind.so /lib/security
# ln -s /usr/lib/security/pam_winbind.so /usr/lib/security/pam_winbind.so.1
# ln -s /lib/security/pam_winbind.so /lib/security/pam_winbind.so.1
```

### CONFIGURING

Now we configure `/etc/krb5.conf` and make an exact copy at [/opt/csw/etc/krb5.conf](#):

And we configure samba at [/usr/samba\\_ad/lib/smb.conf](#):

After this, we configure [/etc/nsswitch.conf](#), to add winbind to users and groups:

```
passwd: files winbind
group: files winbind
```

Our Windows 2003 ActiveDirectory Primary Domain controller is "fwa-dc01.aww.com", our AD domain is "AWW.COM", and our solaris 10 sparc server is "netra.aww.com", so we add these entries to `/etc/hosts`, taking care the FQDN appears first.

```
172.16.1.38 netra.aww.com netra loghost
172.16.1.55 fwa-dc01.aww.com fwa-dc01
```

Now we configure `pam.conf`, we need to make sure home directories are created whenever a user from the windows world logs in, so we get `pam_mkhome` from [http://mega.ist.utl.pt/~filipe/pam\\_mkhome-sol/](http://mega.ist.utl.pt/~filipe/pam_mkhome-sol/), at [http://mega.ist.utl.pt/~filipe/pam\\_mkhome-sol/pam\\_mkhome-sol-0.0.1.tgz](http://mega.ist.utl.pt/~filipe/pam_mkhome-sol/pam_mkhome-sol-0.0.1.tgz)

We build `pam_mkhome` by uncompressing it at `/builds/pam_mkhome` and editing its Makefile to use the solaris native `ld`:

```
LD=/usr/ccs/bin/ld.old
```

now we just build it

```
# cd /builds/pam_mkhome
# make
```

And then we copy it to the system default for PAM modules:

```
# cp /builds/pam_mkhome/pam_mkhome.so /usr/lib/security
# ln -s /usr/lib/security/pam_mkhome.so /usr/lib/security/pam_mkhome.so.1
```

We now configure [/etc/pam.conf](#), to use both `pam_winbind` and `pam_mkhome`, adding 4 lines:

```
login auth sufficient pam_winbind.so.1
other auth sufficient pam_winbind.so.1
other account sufficient pam_winbind.so.1
other session optional pam_mkhome.so.1 skel=/etc/skel umask=0022
```

## STARTING DAEMONS

After this, we can start samba, and see the effects:

```
# export PATH=/usr/samba_ad/sbin:/usr/samba_ad/bin:/opt/csw/bin:/opt/csw/sbin:$PATH
# pkill winbindd; pkill smbd;pkill nmbd, /etc/init.d/cswsamba stop
# /opt/csw/bin/kdestroy
# net ads leave -U "Administrator%M1cr050ft123"
# /opt/csw/bin/kinit Administrator@AWW.COM
# net ads join -U "Administrator%M1cr050ft123"
```

```
# winbindd -B  
# smbd -D  
# nmbd -D
```

## TESTING

I wrote an easy test script called sambatest.ksh:

```
#!/bin/bash  
export
```

```
PATH=/usr/samba_ad/sbin:/usr/samba_ad/bin:/opt/csw/bin:/opt/csw/sbin:/usr/sbin:/usr/bin:/usr/openwin/bin:/opt/csw/  
export LD_LIBRARY_PATH=/usr/samba_ad/lib  
set -x  
which net  
which smbd  
which nmbd  
which winbindd  
which klist  
which kinit
```

```
sleep 5
smbclient -L netra -U%
wbinfo -m
wbinfo -g
wbinfo -u
getent passwd
getent group
net ads status
net ads info
```

The test log file should look like [this](#). Please note that wbinfo -m shows your AD domain, wbinfo -g and wbinfo -u show users and groups, getent passwd should show both your local and AD users. Now we use ssh with an ActiveDirectory user+pass, it should look like this:

```
root@kraftek:~# ssh bmahock@netra
Password:
Creating directory '/export/home/AWW.bmahock'.
Last login: Sat Feb 2 19:53:08 2008 from localhost
/usr/openwin/bin/xauth: creating new authority file /export/home/AWW.bmahock/.Xauthority
Sun Microsystems Inc. SunOS 5.10 Generic January 2005
-bash-3.00$ id
uid=10004(bmahock) gid=10008(domain users)
-bash-3.00$ pwd
/export/home/AWW.bmahock
```

Then it's just a question of enjoying. :)

## **BINARIES**

If you don't want to compile much, just install the CSW packages mentioned above, get these binaries: [samba\\_ad\\_solaris\\_10\\_sparc.tar.gz](#), put them at / and untar them, be careful, these contain the libraries needed at /usr/lib/security and /usr/lib, then just configure /etc/nsswitch.conf, /etc/pam.conf, /usr/samba\_ad/lib/smb.conf, /etc/krb5.conf and /opt/csw/etc/krb5.conf, verify for each program and library that no symbol is missing using ldd, and start your engines!

## **SCREENSHOTS**

These are some screenshots on the ActiveDirectory PDC: [Computers](#) and [Users](#)

A screenshot of a [freenx solaris gnome session](#) for an ActiveDirectory user, yes he can also use sudo!

Our ActiveDirectory user can login to [windows xp](#) using his AD user+pass, and passwordless-ly view and create files at NETRA's samba shares, since he is the only one that can do it, according to [smb.conf](#).

Posted by rdircio in Solaris at 20:46

## Solaris 10 Sparc ActiveDirectory integration with Samba 3.0.23 from blastwave

This is how to integrate the blastwave-provided samba 3.0.23 into a Windows 2003 Active Directory Primary Domain Controller, using kerberos, ldap, nss, samba, winbind.

This gets you to the stage where you can access your samba shares passwordless, when logged in as an authorized user to access them. This does not get you to login to your solaris server (ssh,telnet) with an ActiveDirectory user+pass.

First, set your pkg-get repo by editing /opt/csw/etc/pkg-get.conf:

```
url=http://ibiblio.org/pub/packages/solaris/csw/unstable
```

Now, we must get these pkgs:

```
CSWkrb5user, CSWkrb5libdev, CSWsamba, CSWsambaclient, CSWsambacommon,CSWsambalib, CSWsambalibdev, CSWsambawb
```

We do so by issuing:

```
# pkg-get -i krb5_user krb5_lib_dev samba samba_client samba_common samba_lib samba_libdev samba_wb
```

Several dependencies will be installed automagically:

```
CSWsasl, CSWoldaprt, CSWlibpopt, CSWlibnet, CSWlibcups, CSWkrb5lib, CSWiconv, CSWgettext, CSWfam, CSWcommon
```

After getting all packages make sure you have samba 3.0.23 by issuing:

```
# pkginfo -l CSWsambalibdev CSWsamba CSWsambalib CSWsambacommon CSWsambaclient | grep VERSION
VERSION: 3.0.23,REV=2006.08.09b
VERSION: 3.0.23,REV=2006.08.09b
VERSION: 3.0.23,REV=2006.08.09b
VERSION: 3.0.23,REV=2006.08.09b
VERSION: 3.0.23,REV=2006.08.09b
```

It is very important that all versions above are equal and that you have 3.0.23.

We can now configure, we start by configuring kerberos at /etc/krb5.conf, and make an exact copy at /opt/csw/etc/krb5.conf

```
[libdefaults]
clockskew = 300
default_realm = AWW.COM
[realms]
AWW.COM = {
```

```

kdc = fwa-dc01.aww.com
default_domain = AWW
kpasswd_server = fwa-dc01.aww.com
}
[domain_realm]
.AWW = AWW.COM
[logging]
default = SYSLOG:NOTICE:DAEMON
kdc = FILE:/var/log/kdc.log
kadmind = FILE:/var/log/kadmind.log
[appdefaults]
pam = {
ticket_lifetime = 1d
renew_lifetime = 1d
forwardable = true
proxiabile = false
retain_after_close = false
minimum_uid = 0
debug = false
}

```

We now configure samba, the config file is at /opt/csw/etc/samba/smb.conf

```

[global]
realm = AWW.COM
workgroup = AWW
password server = fwa-dc01.aww.com
security = ADS
encrypt passwords = yes
map to guest = never
client use spnego = yes
idmap uid = 10000-60000
winbind gid = 10000-60000
winbind use default domain = yes
netbios name = netra
loglevel = 10
interfaces = 172.16.1.38
bind interfaces only = yes
[200gb]
comment = 200gb disk
path = /200gb
read only = no
public = yes
browseable = yes
writeable = yes
create mode = 644
valid users = AWW\bmahock
[htdocs]
comment = Apache htdocs
path = /usr/local/apache2/htdocs
read only = no
public = yes
browseable = yes

```

```
writeable = yes
create mode = 644
valid users = AWW\bmahock
```

Our Windows 2003 ActiveDirectory PDC is called fwa-dc01. Our Solaris 10 server is called "netra" and it's fqdn is "netra.aww.com", so we need to add that to /etc/hosts, please make sure the FQDN appears FIRST!!!!!!!!!!!!!!

```
172.16.1.38 netra.aww.com netra loghost
172.16.1.55 fwa-dc01.aww.com fwa-dc01
```

We don't need to edit /etc/resolv.conf to add any particular dns settings.

We need to add winbind to /etc/nsswitch.conf:

```
passwd: files winbind
group: files winbind
```

We need to get libgroups in order for winbind to work, check <http://www.blastwave.org/~fredrik/>, and get <http://www.blastwave.org/~fredrik/libgroups.SPARC.so>

I placed that at /usr/lib/libgroups.SPARC.so, and as stated at that url, added a line to the top of /etc/init.d/cswsamba with:

```
LD_PRELOAD=/usr/lib/libgroups.SPARC.so
```

Now we can join samba to the AD PDC and start rolling!

If not already there, set /opt/csw/\* in the PATH

```
# export PATH=$PATH:/opt/csw/sbin:/opt/csw/bin
```

Kill any already running samba:

```
# pkill winbindd
# pkill smbd
# pkill nmbd
# /etc/init.d/cswsamba stop
```

Destroy any kerberos ticket we might have:

```
# /opt/csw/bin/kdestroy
```

Test ticket creation to the AD PDC, this will test if ur /etc/krb5.conf and /opt/csw/etc/krb5.conf is ok

```
# /opt/csw/bin/kinit Administrator@AWW.COM
```

Password for Administrator@AWW.COM:

Delete any previous instance of our server in the AD

```
# /opt/csw/bin/net ads leave -U Administrator
Administrator's password:
```

Join our samba server to the AD

```
# /opt/csw/bin/net ads join -U Administrator
Administrator's password:
Using short domain name -- AWW
Joined 'NETRA' to realm 'AWW.COM'
```

```
Start all samba daemons
# /etc/init.d/cswsamba start
```

Now test if all is ok:

Display samba status:

```
bash-3.00# /opt/csw/bin/smbclient -L netra -U%
Domain=[AWW] OS=[Unix] Server=[Samba 3.0.23b]
Sharename Type Comment
-----
200gb Disk 200gb disk
htdocs Disk Apache htdocs
IPC$ IPC IPC Service (Samba 3.0.23b)
Domain=[AWW] OS=[Unix] Server=[Samba 3.0.23b]
Server Comment
-----
NETRA Samba 3.0.23b
Workgroup Master
-----
AWW
```

Check for Domains we belong to, this is very important!!! our domain (AWW) must be there!

```
bash-3.00# wbinfo -m
AWW
```

Check for users in the domain:

```
bash-3.00# wbinfo -u
Administrator
Guest
SUPPORT_388945a0
krbtgt
bmahock
tpsmith
quser
```

Check for groups in the domain:

```
bash-3.00# wbinfo -g
BUILTIN\administrators
BUILTIN\
Domain Computers
Domain Controllers
Schema Admins
Enterprise Admins
Domain Admins
```

Domain Users  
Domain Guests  
Group Policy Creator Owners  
DnsUpdateProxy

Get information for a user, this would test that you configured /etc/nsswitch.conf ok with winbind, and that it is using /lib/libnss\_winbind.so

```
bash-3.00# getent passwd "bmahock"  
bmahock:*:10004:10002:Brian Mahocker:/home/AWW/bmahock:/bin/false
```

Now test if you did really join through ActiveDirectory, using "net ads" queries:

```
bash-3.00# net ads testjoin  
Join is OK  
bash-3.00# net ads info  
LDAP server: 172.16.1.55  
LDAP server name: fwa-dc01.aww.com  
Realm: AWW.COM  
Bind Path: dc=AWW,dc=COM  
LDAP port: 389  
Server time: Sat, 02 Feb 2008 11:13:52 CST  
KDC server: 172.16.1.55  
Server time offset: 76
```

Now get all ActiveDirectory properties for your computer:

```
bash-3.00# net ads status  
objectClass: top  
objectClass: person  
objectClass: organizationalPerson  
objectClass: user  
objectClass: computer  
cn: netra  
distinguishedName: CN=netra,CN=Computers,DC=aww,DC=com  
instanceType: 4  
whenCreated: 20080202171033.0Z  
whenChanged: 20080202171208.0Z  
uSNCreated: 45102  
uSNChanged: 45137  
name: netra  
objectGUID: bea32ea2-8208-4119-9d4c-87186a3866c4  
userAccountControl: 69632  
badPwdCount: 0  
codePage: 0  
countryCode: 0  
badPasswordTime: 0  
lastLogoff: 0  
lastLogon: 128464460293750000  
localPolicyFlags: 0  
pwdLastSet: 0  
primaryGroupID: 515  
objectSid: S-1-5-21-875145104-1252384033-126309636-1125
```

accountExpires: 9223372036854775807  
logonCount: 13  
sAMAccountName: netra\$  
sAMAccountType: 805306369  
dNSHostName: netra.aww.com  
servicePrincipalName: HOST/netra.aww.com  
servicePrincipalName: HOST/NETRA  
objectCategory: CN=Computer,CN=Schema,CN=Configuration,DC=aww,DC=com  
isCriticalSystemObject: FALSE  
----- Security Descriptor (revision: 1, type: 0x8c14)  
owner SID: S-1-5-21-875145104-1252384033-126309636-512  
group SID: S-1-5-21-875145104-1252384033-126309636-513  
----- (system) ACL (revision: 4, size: 120, number of ACEs: 2)  
----- ACE (type: 0x07, flags: 0x5a, size: 0x38, mask: 0x20, object flags: 0x3)  
access SID: S-1-1-0  
access type: AUDIT OBJECT  
Permissions:  
[Write All Properties]  
----- ACE (type: 0x07, flags: 0x5a, size: 0x38, mask: 0x20, object flags: 0x3)  
access SID: S-1-1-0  
access type: AUDIT OBJECT  
Permissions:  
[Write All Properties]  
----- (user) ACL (revision: 4, size: 1908, number of ACEs: 39)  
----- ACE (type: 0x00, flags: 0x00, size: 0x24, mask: 0xf01ff)  
access SID: S-1-5-21-875145104-1252384033-126309636-512  
access type: ALLOWED  
Permissions: [Full Control]  
----- ACE (type: 0x00, flags: 0x00, size: 0x18, mask: 0xf01ff)  
access SID: S-1-5-32-548  
access type: ALLOWED  
Permissions: [Full Control]  
----- ACE (type: 0x00, flags: 0x00, size: 0x14, mask: 0xf01ff)  
access SID: S-1-5-18  
access type: ALLOWED  
Permissions: [Full Control]  
----- ACE (type: 0x00, flags: 0x00, size: 0x24, mask: 0x301d4)  
access SID: S-1-5-21-875145104-1252384033-126309636-512  
access type: ALLOWED  
Permissions:  
[List Contents]  
[Read All Properties]  
[Delete Subtree]  
[List Object]  
[Change Password]  
[Reset Password]  
[Delete]  
[Read Permissions]  
----- ACE (type: 0x05, flags: 0x00, size: 0x38, mask: 0x20, object flags: 0x1)  
access SID: S-1-5-21-875145104-1252384033-126309636-512  
access type: ALLOWED OBJECT  
Permissions:  
[Write All Properties]

----- ACE (type: 0x00, flags: 0x00, size: 0x14, mask: 0x20094)  
access SID: S-1-5-11  
access type: ALLOWED  
Permissions:  
[List Contents]  
[Read All Properties]  
[List Object]  
[Read Permissions]  
----- ACE (type: 0x05, flags: 0x00, size: 0x28, mask: 0x100, object flags: 0x1)  
access SID: S-1-1-0  
access type: ALLOWED OBJECT  
Permissions:  
[Change Password]  
[Reset Password]  
----- ACE (type: 0x00, flags: 0x00, size: 0x14, mask: 0x3)  
access SID: S-1-5-10  
access type: ALLOWED  
Permissions:  
[Create All Child Objects]  
[Delete All Child Objects]  
----- ACE (type: 0x05, flags: 0x00, size: 0x2c, mask: 0x3, object flags: 0x1)  
access SID: S-1-5-32-550  
access type: ALLOWED OBJECT  
Permissions:  
[Create All Child Objects]  
[Delete All Child Objects]  
----- ACE (type: 0x05, flags: 0x00, size: 0x38, mask: 0x30, object flags: 0x1)  
access SID: S-1-5-21-875145104-1252384033-126309636-517  
access type: ALLOWED OBJECT  
Permissions:  
[Read All Properties]  
[Write All Properties]  
----- ACE (type: 0x05, flags: 0x00, size: 0x28, mask: 0x8, object flags: 0x1)  
access SID: S-1-5-10  
access type: ALLOWED OBJECT  
Permissions:  
[All validate writes]  
----- ACE (type: 0x05, flags: 0x00, size: 0x28, mask: 0x30, object flags: 0x1)  
access SID: S-1-5-10  
access type: ALLOWED OBJECT  
Permissions:  
[Read All Properties]  
[Write All Properties]  
----- ACE (type: 0x05, flags: 0x00, size: 0x28, mask: 0x8, object flags: 0x1)  
access SID: S-1-5-10  
access type: ALLOWED OBJECT  
Permissions:  
[All validate writes]  
----- ACE (type: 0x05, flags: 0x00, size: 0x38, mask: 0x8, object flags: 0x1)  
access SID: S-1-5-21-875145104-1252384033-126309636-512  
access type: ALLOWED OBJECT  
Permissions:  
[All validate writes]

----- ACE (type: 0x05, flags: 0x00, size: 0x38, mask: 0x8, object flags: 0x1)  
access SID: S-1-5-21-875145104-1252384033-126309636-512  
access type: ALLOWED OBJECT  
Permissions:  
[All validate writes]

----- ACE (type: 0x05, flags: 0x00, size: 0x48, mask: 0x20, object flags: 0x3)  
access SID: S-1-5-21-875145104-1252384033-126309636-512  
access type: ALLOWED OBJECT  
Permissions:  
[Write All Properties]

----- ACE (type: 0x05, flags: 0x00, size: 0x48, mask: 0x20, object flags: 0x3)  
access SID: S-1-5-21-875145104-1252384033-126309636-512  
access type: ALLOWED OBJECT  
Permissions:  
[Write All Properties]

----- ACE (type: 0x05, flags: 0x00, size: 0x48, mask: 0x20, object flags: 0x3)  
access SID: S-1-5-21-875145104-1252384033-126309636-512  
access type: ALLOWED OBJECT  
Permissions:  
[Write All Properties]

----- ACE (type: 0x05, flags: 0x00, size: 0x48, mask: 0x20, object flags: 0x3)  
access SID: S-1-5-21-875145104-1252384033-126309636-512  
access type: ALLOWED OBJECT  
Permissions:  
[Write All Properties]

----- ACE (type: 0x05, flags: 0x00, size: 0x2c, mask: 0x10, object flags: 0x1)  
access SID: S-1-5-32-560  
access type: ALLOWED OBJECT  
Permissions:  
[Read All Properties]

----- ACE (type: 0x00, flags: 0x12, size: 0x18, mask: 0xf01bd)  
access SID: S-1-5-32-544  
access type: ALLOWED  
Permissions:  
[Create All Child Objects]  
[List Contents]  
[All validate writes]  
[Read All Properties]  
[Write All Properties]  
[List Object]  
[Change Password]  
[Reset Password]  
[Delete]  
[Read Permissions]  
[Modify Permissions]  
[Modify Owner]

----- ACE (type: 0x00, flags: 0x12, size: 0x24, mask: 0xf01ff)  
access SID: S-1-5-21-875145104-1252384033-126309636-519  
access type: ALLOWED  
Permissions: [Full Control]

----- ACE (type: 0x00, flags: 0x12, size: 0x18, mask: 0x4)  
access SID: S-1-5-32-554  
access type: ALLOWED

Permissions:

[List Contents]

----- ACE (type: 0x05, flags: 0x1a, size: 0x3c, mask: 0x10, object flags: 0x3)

access SID: S-1-5-32-554

access type: ALLOWED OBJECT

Permissions:

[Read All Properties]

----- ACE (type: 0x05, flags: 0x1a, size: 0x3c, mask: 0x10, object flags: 0x3)

access SID: S-1-5-32-554

access type: ALLOWED OBJECT

Permissions:

[Read All Properties]

----- ACE (type: 0x05, flags: 0x1a, size: 0x3c, mask: 0x10, object flags: 0x3)

access SID: S-1-5-32-554

access type: ALLOWED OBJECT

Permissions:

[Read All Properties]

----- ACE (type: 0x05, flags: 0x1a, size: 0x3c, mask: 0x10, object flags: 0x3)

access SID: S-1-5-32-554

access type: ALLOWED OBJECT

Permissions:

[Read All Properties]

----- ACE (type: 0x05, flags: 0x1a, size: 0x3c, mask: 0x10, object flags: 0x3)

access SID: S-1-5-32-554

access type: ALLOWED OBJECT

Permissions:

[Read All Properties]

----- ACE (type: 0x05, flags: 0x1a, size: 0x2c, mask: 0x20094, object flags: 0x2)

access SID: S-1-5-32-554

access type: ALLOWED OBJECT

Permissions:

[List Contents]

[Read All Properties]

[List Object]

[Read Permissions]

----- ACE (type: 0x05, flags: 0x1a, size: 0x2c, mask: 0x20094, object flags: 0x2)

access SID: S-1-5-32-554

access type: ALLOWED OBJECT

Permissions:

[List Contents]

[Read All Properties]

[List Object]

[Read Permissions]

----- ACE (type: 0x05, flags: 0x1a, size: 0x3c, mask: 0x10, object flags: 0x3)

access SID: S-1-5-32-554

access type: ALLOWED OBJECT

Permissions:

[Read All Properties]

----- ACE (type: 0x05, flags: 0x1a, size: 0x3c, mask: 0x10, object flags: 0x3)

access SID: S-1-5-32-554

access type: ALLOWED OBJECT

Permissions:

[Read All Properties]

```
----- ACE (type: 0x05, flags: 0x1a, size: 0x3c, mask: 0x10, object flags: 0x3)
access SID: S-1-5-32-554
access type: ALLOWED OBJECT
Permissions:
  [Read All Properties]
----- ACE (type: 0x05, flags: 0x1a, size: 0x3c, mask: 0x10, object flags: 0x3)
access SID: S-1-5-32-554
access type: ALLOWED OBJECT
Permissions:
  [Read All Properties]
----- ACE (type: 0x05, flags: 0x1a, size: 0x3c, mask: 0x10, object flags: 0x3)
access SID: S-1-5-32-554
access type: ALLOWED OBJECT
Permissions:
  [Read All Properties]
----- ACE (type: 0x05, flags: 0x1a, size: 0x2c, mask: 0x20094, object flags: 0x2)
access SID: S-1-5-32-554
access type: ALLOWED OBJECT
Permissions:
  [List Contents]
  [Read All Properties]
  [List Object]
  [Read Permissions]
----- ACE (type: 0x05, flags: 0x1a, size: 0x38, mask: 0x10, object flags: 0x3)
access SID: S-1-5-9
access type: ALLOWED OBJECT
Permissions:
  [Read All Properties]
----- ACE (type: 0x05, flags: 0x1a, size: 0x38, mask: 0x10, object flags: 0x3)
access SID: S-1-5-9
access type: ALLOWED OBJECT
Permissions:
  [Read All Properties]
----- ACE (type: 0x05, flags: 0x12, size: 0x38, mask: 0x10, object flags: 0x3)
access SID: S-1-5-9
access type: ALLOWED OBJECT
Permissions:
  [Read All Properties]
----- End Of Security Descriptor
```

If all these tests are ok, your samba is joined correctly to the AD domain.

Posted by rdircio in Solaris at 15:43

Friday, January 18. 2008

## **NX+freenx server binary tarball for Solaris 10 x86**

UPDATED NOTE: for OPENSOLARIS use this link: [http://www.opensolaris.org/os/project/awards/awards\\_land/Entries](http://www.opensolaris.org/os/project/awards/awards_land/Entries)

I compiled and built this ready to use nx+freenx server tarball for solaris10 x86 (not opensolaris). This is both freenx 0.7.0 and the nomachine NX libraries version 2.1.0-2.

Tested on a fresh solaris 10 x86 u4 8/07

Here's [gnome](#) and [CDE](#) on sol10x86, for the non-believers.

Just download [nx\\_freenx\\_sol10x86.tar.gz](#)

place it in your server at /, and uncompress.

```
# cd /
```

```
# gunzip nx_freenx_sol10x86.tar.gz
```

```
# tar xvf nx_freenx_sol10x86.tar
```

Then run the nxinstall script, which will check for the required packages. all of them from blastwave.org.

**VERY IMPORTANT: if these are missing, install them!!!! nx+freenx won't work without ggetopt, expect, gmd5sum, gfind, ggetopt,netcat from blastwave at /opt/csw/bin!!!!**

```
# nxinstall
```

a log of a succesful nxinstall should look like [this](#)

After this, you'll be ready to use any nx client on another computer to connect to your Solaris 10 x86 freenx server!

ENJOY!

Posted by rdircio in Solaris at 00:04

Tuesday, January 15. 2008

## some zfs commands

To see what pools do you have:

```
bash-3.00# zpool list
NAME SIZE USED AVAIL CAP HEALTH ALTROOT
mot117zpool1 116G 24.7G 90.9G 21% ONLINE -
```

To see what disks comprise a pool

```
bash-3.00# zpool status mot117zpool1
pool: mot117zpool1
state: ONLINE
scrub: none requested
config:

NAME STATE READ WRITE CKSUM
mot117zpool1 ONLINE 0 0 0
c4t60060480000190101181533031393341d0 ONLINE 0 0 0
c4t60060480000190101181533031434543d0 ONLINE 0 0 0
```

errors: No known data errors

To check all pools health:

```
# zpool status -x
```

To mount all the zfs filesystems:

```
# zfs mount -a
```

Posted by rdircio in Solaris at 22:10

Wednesday, December 12, 2007

## **"Error with console keyboard"**

The "Error with console keyboard" appears on your Sun console, and soon you find out you can't type anything, even stop-a does not work...

This is bug #57688 as stated in <http://sunsolve.sun.com/search/document.do?assetkey=1-26-57688-1>, and all you got to do is remove the line:

```
exclude: drv/ohci
```

```
from /etc/system
```

Posted by rdircio in Solaris at 01:41

Tuesday, November 13. 2007

## Check status of your backups

A good practical command to take a look on how your TSM jobs are running is to use:

```
# egrep -i "Schedule Name|Finished" /var/tivoli/dsmsched.log
```

Posted by rdircio in Solaris at 11:30

## big concat metadvice

If you need to create a Big concat metadvice from a bunch of EMC or HITACHI disks, you can use this script:

```
#!/bin/ksh
```

```
( N=`format < /dev/null | egrep -i 'emc|hit' | wc -l | sed 's//g'`  
/usr/ucb/echo -n "metainit d1 $N "  
format < /dev/null | egrep -i 'emc|hit' | awk '{ print $2 }' | while read d;do  
/usr/ucb/echo -n "1 ${d}s2 "  
done  
echo "" ) > /tmp/makevol.ksh  
chmod 755 /tmp/makevol.ksh
```

Posted by rdircio in Solaris at 09:32

Friday, September 7. 2007

## **Unexpire an account and set to never expire**

just do

```
# passwd -x -1 <account>
```

Posted by rdircio in Solaris at 14:07

Wednesday, September 5. 2007

## Metafree

If you want to know how much space you have free for soft partitioning use Adrian Ball's script "metafree", get it at:  
<http://majords.co.uk/scripts/metafree>

or just copy/paste this:

```
#!/bin/ksh
#
# metafree - Adrian Ball (ade@majords.co.uk)- 2007/01
#
# Finds all devices containing soft partitions and shows
# size/allocation, unallocated space, devices etc
#
# NB: This is not the script that is listed on Sun's BigAdmin site
# I had written this before knowing of that one's existence, but
# subsequently found all the links to that one to be broken, so
# I have not seen it.
# Strangely I had picked the same name for the script...
# Maybe someone will find this one useful.
if [[ ! -x /usr/sbin/metastat ]] ; then
    echo "Can't run metastat binary, bogus..."
    exit 1
fi
printf "%-20s %9s %9s %10s %10s %s\n" Device GB Used Available Capacity "Soft
Partitions"
for c in $(metastat -p | nawk '#!/bin/ksh
#
# metafree - Adrian Ball (ade@majords.co.uk)- 2007/01
#
# Finds all devices containing soft partitions and shows
# size/allocation, unallocated space, devices etc
#
# NB: This is not the script that is listed on Sun's BigAdmin site
# I had written this before knowing of that one's existence, but
# subsequently found all the links to that one to be broken, so
# I have not seen it.
# Strangely I had picked the same name for the script...
# Maybe someone will find this one useful.
if [[ ! -x /usr/sbin/metastat ]] ; then
    echo "Can't run metastat binary, bogus..."
    exit 1
fi
printf "%-20s %9s %9s %10s %10s %s\n" Device GB Used Available Capacity "Soft
Partitions"
for c in $(metastat -p | nawk '$2=="-p" {print $3}' | sort -u) ; do
    if [[ ${c%[0-9]*} == d ]] ; then
```

```

### It's a metadvice, use metastat to find the size
cap=$(metastat $c 2>/dev/null| nawk '/Size:/ {print $2/2048/1024;exit}')
elif [[ ${c#/} != $c ]] ; then
### It's a real device, work out size assuming full disk/LUN is used
cap=$(prtvtoc $c | nawk '
/bytes\sector/ {bs=$2}
/sectors\cyl/ {sc=$2}
/accessible cyl/ {as=$2}
END {print bs*sc*as/1024^3}')
else
### Assume we have a cXtX form of device, use prtvtoc to work out
### the size of the slice
cap=$(prtvtoc /dev/rdisk/$c | nawk -v c=$c '
BEGIN {split(c,a,"s");s=a[2]}
/bytes\sector/ {bs=$2}
$1==s {print $5*bs/1024^3}')
fi
used=$(metastat -p | nawk -v c=$c '
$2=="-p" && $3==c {split($0,a,"-o")
for (frag in a) { split(a[frag],sz,"-b"); used+=sz[2] } }
END { print used/2048/1024 }
')
echo $cap $used | nawk '{print $1-$2, $2/$1*100}' | read avail pct
devs=$(metastat -p | sort | nawk -v c=$c '$3==c && $2=="-p" {printf "%s ", $1}')
printf "%-20s %9.2f %9.2f %10.2f %9.1f%1s %s\n" $c $cap $used $avail $pct %
"$devs"
done =="-p" {print }' | sort -u) ; do
    if [[ ${c%[0-9]*} == d ]] ; then
### It's a metadvice, use metastat to find the size
cap=$(metastat $c 2>/dev/null| nawk '/Size:/ {print #!/bin/ksh
#
# metafree - Adrian Ball (ade@majords.co.uk)- 2007/01
#
# Finds all devices containing soft partitions and shows
# size/allocation, unallocated space, devices etc
#
# NB: This is not the script that is listed on Sun's BigAdmin site
# I had written this before knowing of that one's existence, but
# subsequently found all the links to that one to be broken, so
# I have not seen it.
# Strangely I had picked the same name for the script...
# Maybe someone will find this one useful.
if [[ ! -x /usr/sbin/metastat ]] ; then
    echo "Can't run metastat binary, bogus..."
    exit 1
fi
printf "%-20s %9s %9s %10s %10s %s\n" Device GB Used Available Capacity "Soft
Partitions"
for c in $(metastat -p | nawk '$2=="-p" {print $3}' | sort -u) ; do
    if [[ ${c%[0-9]*} == d ]] ; then
### It's a metadvice, use metastat to find the size
cap=$(metastat $c 2>/dev/null| nawk '/Size:/ {print $2/2048/1024;exit}')
elif [[ ${c#/} != $c ]] ; then

```

```

### It's a real device, work out size assuming full disk/LUN is used
cap=$(prtvtoc $c | nawk '
/bytes\sector/ {bs=$2}
/sectors\cyl/ {sc=$2}
/accessible cyl/ {as=$2}
END {print bs*sc*as/1024^3}')
else
### Assume we have a cXtX form of device, use prtvtoc to work out
### the size of the slice
cap=$(prtvtoc /dev/rdisk/$c | nawk -v c=$c '
BEGIN {split(c,a,"s");s=a[2]}
/bytes\sector/ {bs=$2}
$1==s {print $5*bs/1024^3}')
fi
used=$(metastat -p | nawk -v c=$c '
$2=="-p" && $3==c {split($0,a,"-o")}
for (frag in a) { split(a[frag],sz,"-b"); used+=sz[2] } }
END { print used/2048/1024 }
')
echo $cap $used | nawk '{print $1-$2, $2/$1*100}' | read avail pct
devs=$(metastat -p | sort | nawk -v c=$c '$3==c && $2=="-p" {printf "%s ", $1}')
printf "%-20s %9.2f %9.2f %10.2f %9.1f%1s %s\n" $c $cap $used $avail $pct %
"$devs"
done /2048/1024;exit}')
elif [[ ${c#}/} != $c ]] ; then
### It's a real device, work out size assuming full disk/LUN is used
cap=$(prtvtoc $c | nawk '
/bytes\sector/ {bs=#!/bin/ksh
#
# metafree - Adrian Ball (ade@majords.co.uk)- 2007/01
#
# Finds all devices containing soft partitions and shows
# size/allocation, unallocated space, devices etc
#
# NB: This is not the script that is listed on Sun's BigAdmin site
# I had written this before knowing of that one's existence, but
# subsequently found all the links to that one to be broken, so
# I have not seen it.
# Strangely I had picked the same name for the script...
# Maybe someone will find this one useful.
if [[ ! -x /usr/sbin/metastat ]] ; then
echo "Can't run metastat binary, bogus..."
exit 1
fi
printf "%-20s %9s %9s %10s %10s %s\n" Device GB Used Available Capacity "Soft
Partitions"
for c in $(metastat -p | nawk '$2=="-p" {print $3}' | sort -u) ; do
if [[ ${c%[0-9]*} == d ]] ; then
### It's a metadvice, use metastat to find the size
cap=$(metastat $c 2>/dev/null| nawk '/Size:/ {print $2/2048/1024;exit}')
elif [[ ${c#}/} != $c ]] ; then
### It's a real device, work out size assuming full disk/LUN is used
cap=$(prtvtoc $c | nawk '

```

```

/bytes\sector/ {bs=$2}
/sectors\cyl/ {sc=$2}
/accessible cyl/ {as=$2}
END {print bs*sc*as/1024^3}' )
else
### Assume we have a cXtX form of device, use prtvtoc to work out
### the size of the slice
cap=$(prtvtoc /dev/rdsk/$c | nawk -v c=$c '
BEGIN {split(c,a,"s");s=a[2]}
/bytes\sector/ {bs=$2}
$1==s {print $5*bs/1024^3}' )
fi
used=$(metastat -p | nawk -v c=$c '
$2=="-p" && $3==c {split($0,a,"-o")}
for (frag in a) { split(a[frag],sz,"-b"); used+=sz[2] } }
END { print used/2048/1024 }
')
echo $cap $used | nawk '{print $1-$2, $2/$1*100}' | read avail pct
devs=$(metastat -p | sort | nawk -v c=$c '$3==c && $2=="-p" {printf "%s ", $1}')
printf "%-20s %9.2f %9.2f %10.2f %9.1f%1s %s\n" $c $cap $used $avail $pct %
"$devs"
done }
/sectors\cyl/ {sc=#!/bin/ksh
#
# metafree - Adrian Ball (ade@majords.co.uk)- 2007/01
#
# Finds all devices containing soft partitions and shows
# size/allocation, unallocated space, devices etc
#
# NB: This is not the script that is listed on Sun's BigAdmin site
# I had written this before knowing of that one's existence, but
# subsequently found all the links to that one to be broken, so
# I have not seen it.
# Strangely I had picked the same name for the script...
# Maybe someone will find this one useful.
if [[ ! -x /usr/sbin/metastat ]] ; then
echo "Can't run metastat binary, bogus..."
exit 1
fi
printf "%-20s %9s %9s %10s %10s %s\n" Device GB Used Available Capacity "Soft
Partitions"
for c in $(metastat -p | nawk '$2=="-p" {print $3}' | sort -u) ; do
if [[ ${c%[0-9]*} == d ]] ; then
### It's a metadvice, use metastat to find the size
cap=$(metastat $c 2>/dev/null| nawk '/Size:/ {print $2/2048/1024;exit}')
elif [[ ${c#/} != $c ]] ; then
### It's a real device, work out size assuming full disk/LUN is used
cap=$(prtvtoc $c | nawk '
/bytes\sector/ {bs=$2}
/sectors\cyl/ {sc=$2}
/accessible cyl/ {as=$2}
END {print bs*sc*as/1024^3}' )
else

```

```

### Assume we have a cXtX form of device, use prtvtoc to work out
### the size of the slice
cap=$(prtvtoc /dev/rdisk/$c | nawk -v c=$c '
BEGIN {split(c,a,"s");s=a[2]}
/bytes\sector/ {bs=$2}
$1==s {print $5*bs/1024^3}' )
fi
used=$(metastat -p | nawk -v c=$c '
$2=="-p" && $3==c {split($0,a,"-o")
for (frag in a) { split(a[frag],sz,"-b"); used+=sz[2] } }
END { print used/2048/1024 }
')
echo $cap $used | nawk '{print $1-$2, $2/$1*100}' | read avail pct
devs=$(metastat -p | sort | nawk -v c=$c '$3==c && $2=="-p" {printf "%s ", $1}')
printf "%-20s %9.2f %9.2f %10.2f %9.1f%1s %s\n" $c $cap $used $avail $pct %
"$devs"
done }
    /accessible cyl/ {as=#!/bin/ksh
#
# metafree - Adrian Ball (ade@majords.co.uk)- 2007/01
#
# Finds all devices containing soft partitions and shows
# size/allocation, unallocated space, devices etc
#
# NB: This is not the script that is listed on Sun's BigAdmin site
# I had written this before knowing of that one's existence, but
# subsequently found all the links to that one to be broken, so
# I have not seen it.
# Strangely I had picked the same name for the script...
# Maybe someone will find this one useful.
if [[ ! -x /usr/sbin/metastat ]] ; then
    echo "Can't run metastat binary, bogus..."
    exit 1
fi
printf "%-20s %9s %9s %10s %10s %s\n" Device GB Used Available Capacity "Soft
Partitions"
for c in $(metastat -p | nawk '$2=="-p" {print $3}' | sort -u) ; do
    if [[ ${c%[0-9]*} == d ]] ; then
        ### It's a metadvice, use metastat to find the size
        cap=$(metastat $c 2>/dev/null| nawk '/Size:/ {print $2/2048/1024;exit}')
    elif [[ ${c#/} != $c ]] ; then
        ### It's a real device, work out size assuming full disk/LUN is used
        cap=$(prtvtoc $c | nawk '
/bytes\sector/ {bs=$2}
/sectors\cyl/ {sc=$2}
/accessible cyl/ {as=$2}
END {print bs*sc*as/1024^3}' )
    else
        ### Assume we have a cXtX form of device, use prtvtoc to work out
        ### the size of the slice
        cap=$(prtvtoc /dev/rdisk/$c | nawk -v c=$c '
BEGIN {split(c,a,"s");s=a[2]}
/bytes\sector/ {bs=$2}

```

```

$1==s {print $5*bs/1024^3}' )
fi
used=$(metastat -p | nawk -v c=$c '
$2=="-p" && $3==c {split($0,a,"-o")
for (frag in a) { split(a[frag],sz,"-b"); used+=sz[2] } }
END { print used/2048/1024 }
')
echo $cap $used | nawk '{print $1-$2, $2/$1*100}' | read avail pct
devs=$(metastat -p | sort | nawk -v c=$c '$3==c && $2=="-p" {printf "%s ", $1}')
printf "%-20s %9.2f %9.2f %10.2f %9.1f%1s %s\n" $c $cap $used $avail $pct %
"$devs"
done }
    END      {print bs*sc*as/1024^3}' )
    else
### Assume we have a cXtX form of device, use prtvtoc to work out
### the size of the slice
cap=$(prtvtoc /dev/rdisk/$c | nawk -v c=$c '
    BEGIN {split(c,a,"s");s=a[2]}
    /bytes\/sector/ {bs=#!/bin/ksh
#
# metafree - Adrian Ball (ade@majords.co.uk)- 2007/01
#
# Finds all devices containing soft partitions and shows
# size/allocation, unallocated space, devices etc
#
# NB: This is not the script that is listed on Sun's BigAdmin site
# I had written this before knowing of that one's existence, but
# subsequently found all the links to that one to be broken, so
# I have not seen it.
# Strangely I had picked the same name for the script...
# Maybe someone will find this one useful.
if [[ ! -x /usr/sbin/metastat ]] ; then
    echo "Can't run metastat binary, bogus..."
    exit 1
fi
printf "%-20s %9s %9s %10s %10s %s\n" Device GB Used Available Capacity "Soft
Partitions"
for c in $(metastat -p | nawk '$2=="-p" {print $3}' | sort -u) ; do
    if [[ ${c%[0-9]*} == d ]] ; then
### It's a metadvice, use metastat to find the size
cap=$(metastat $c 2>/dev/null| nawk '/Size:/ {print $2/2048/1024;exit}')
    elif [[ ${c#/} != $c ]] ; then
### It's a real device, work out size assuming full disk/LUN is used
cap=$(prtvtoc $c | nawk '
/bytes\/sector/ {bs=$2}
/sectors\/cyl/ {sc=$2}
/accessible cyl/ {as=$2}
END {print bs*sc*as/1024^3}' )
    else
### Assume we have a cXtX form of device, use prtvtoc to work out
### the size of the slice
cap=$(prtvtoc /dev/rdisk/$c | nawk -v c=$c '
    BEGIN {split(c,a,"s");s=a[2]}

```

```

/bytes\sector/ {bs=$2}
$1==s {print $5*bs/1024^3}' )
fi
used=$(metastat -p | nawk -v c=$c '
$2=="-p" && $3==c {split($0,a,"-o")
for (frag in a) { split(a[frag],sz,"-b"); used+=sz[2] } }
END { print used/2048/1024 }
')
echo $cap $used | nawk '{print $1-$2, $2/$1*100}' | read avail pct
devs=$(metastat -p | sort | nawk -v c=$c '$3==c && $2=="-p" {printf "%s ", $1}')
printf "%-20s %9.2f %9.2f %10.2f %9.1f%1s %s\n" $c $cap $used $avail $pct %
"$devs"
done }
==s {print *bs/1024^3}' )
fi
used=$(metastat -p | nawk -v c=$c '
#!/bin/ksh
#
# metafree - Adrian Ball (ade@majords.co.uk)- 2007/01
#
# Finds all devices containing soft partitions and shows
# size/allocation, unallocated space, devices etc
#
# NB: This is not the script that is listed on Sun's BigAdmin site
# I had written this before knowing of that one's existence, but
# subsequently found all the links to that one to be broken, so
# I have not seen it.
# Strangely I had picked the same name for the script...
# Maybe someone will find this one useful.
if [[ ! -x /usr/sbin/metastat ]] ; then
echo "Can't run metastat binary, bogus..."
exit 1
fi
printf "%-20s %9s %9s %10s %10s %s\n" Device GB Used Available Capacity "Soft
Partitions"
for c in $(metastat -p | nawk '$2=="-p" {print $3}' | sort -u) ; do
if [[ ${c%[0-9]*} == d ]] ; then
### It's a metadvice, use metastat to find the size
cap=$(metastat $c 2>/dev/null| nawk '/Size:/ {print $2/2048/1024;exit}')
elif [[ ${c#/} != $c ]] ; then
### It's a real device, work out size assuming full disk/LUN is used
cap=$(prtvtoc $c | nawk '
/bytes\sector/ {bs=$2}
/sectors\cyl/ {sc=$2}
/accessible cyl/ {as=$2}
END {print bs*sc*as/1024^3}' )
else
### Assume we have a cXtX form of device, use prtvtoc to work out
### the size of the slice
cap=$(prtvtoc /dev/rdisk/$c | nawk -v c=$c '
BEGIN {split(c,a,"s");s=a[2]}
/bytes\sector/ {bs=$2}
$1==s {print $5*bs/1024^3}' )

```

```

fi
used=$(metastat -p | nawk -v c=$c '
$2=="-p" && $3==c {split($0,a,"-o")
for (frag in a) { split(a[frag],sz,"-b"); used+=sz[2] } }
END { print used/2048/1024 }
')
echo $cap $used | nawk '{print $1-$2, $2/$1*100}' | read avail pct
devs=$(metastat -p | sort | nawk -v c=$c '$3==c && $2=="-p" {printf "%s ", $1}')
printf "%-20s %9.2f %9.2f %10.2f %9.1f%1s %s\n" $c $cap $used $avail $pct %
"$devs"
done =="-p" && ==c {split(
#!/bin/ksh
#
# metafree - Adrian Ball (ade@majords.co.uk)- 2007/01
#
# Finds all devices containing soft partitions and shows
# size/allocation, unallocated space, devices etc
#
# NB: This is not the script that is listed on Sun's BigAdmin site
# I had written this before knowing of that one's existence, but
# subsequently found all the links to that one to be broken, so
# I have not seen it.
# Strangely I had picked the same name for the script...
# Maybe someone will find this one useful.
if [[ ! -x /usr/sbin/metastat ]] ; then
    echo "Can't run metastat binary, bogus..."
    exit 1
fi
printf "%-20s %9s %9s %10s %10s %s\n" Device GB Used Available Capacity "Soft
Partitions"
for c in $(metastat -p | nawk '$2=="-p" {print $3}' | sort -u) ; do
    if [[ ${c%[0-9]*} == d ]] ; then
        ### It's a metadvice, use metastat to find the size
        cap=$(metastat $c 2>/dev/null| nawk '/Size:/ {print $2/2048/1024;exit}')
    elif [[ ${c#/} != $c ]] ; then
        ### It's a real device, work out size assuming full disk/LUN is used
        cap=$(prtvtoc $c | nawk '
/bytes\sector/ {bs=$2}
/sectors\cyl/ {sc=$2}
/accessible cyl/ {as=$2}
END {print bs*sc*as/1024^3}')
    else
        ### Assume we have a cXtX form of device, use prtvtoc to work out
        ### the size of the slice
        cap=$(prtvtoc /dev/rdisk/$c | nawk -v c=$c '
BEGIN {split(c,a,"s");s=a[2]}
/bytes\sector/ {bs=$2}
$1==s {print $5*bs/1024^3}')
    fi
used=$(metastat -p | nawk -v c=$c '
$2=="-p" && $3==c {split($0,a,"-o")
for (frag in a) { split(a[frag],sz,"-b"); used+=sz[2] } }
END { print used/2048/1024 }

```

```
' )
echo $cap $used | nawk '{print $1-$2, $2/$1*100}' | read avail pct
devs=$(metastat -p | sort | nawk -v c=$c '$3==c && $2="-p" {printf "%s ", $1}')
printf "%-20s %9.2f %9.2f %10.2f %9.1f%1s %s\n" $c $cap $used $avail $pct %
"$devs"
done
```

```
,a,"-o")
for (frag in a) { split(a[frag],sz,"-b"); used+=sz[2] }
END { print used/2048/1024 }
```

```
)
echo $cap $used | nawk '{print -#!/bin/ksh
#
# metafree - Adrian Ball (ade@majords.co.uk)- 2007/01
#
# Finds all devices containing soft partitions and shows
# size/allocation, unallocated space, devices etc
#
# NB: This is not the script that is listed on Sun's BigAdmin site
# I had written this before knowing of that one's existence, but
# subsequently found all the links to that one to be broken, so
# I have not seen it.
# Strangely I had picked the same name for the script...
# Maybe someone will find this one useful.
```

```
if [[ ! -x /usr/sbin/metastat ]] ; then
echo "Can't run metastat binary, bogus..."
exit 1
fi
```

```
printf "%-20s %9s %9s %10s %10s %s\n" Device GB Used Available Capacity "Soft Partitions"
```

```
for c in $(metastat -p | nawk '$2=="-p" {print $3}' | sort -u) ; do
if [[ ${c%%[0-9]*} == d ]] ; then
### It's a metadvice, use metastat to find the size
cap=$(metastat $c 2>/dev/null | nawk '/Size:/ {print $2/2048/1024;exit}')
elif [[ ${c#}/ != $c ]] ; then
### It's a real device, work out size assuming full disk/LUN is used
cap=$(prtvtoc $c | nawk '
/bytes\sector/ {bs=$2}
/sectors\cyl/ {sc=$2}
/accessible cyl/ {as=$2}
END {print bs*sc*as/1024^3}')
else
### Assume we have a cXtX form of device, use prtvtoc to work out
### the size of the slice
cap=$(prtvtoc /dev/rdisk/$c | nawk -v c=$c '
BEGIN {split(c,a,"s");s=a[2]}
/bytes\sector/ {bs=$2}
$1==s {print $5*bs/1024^3}')
fi
used=$(metastat -p | nawk -v c=$c '
```

```

$2=="-p" && $3==c {split($0,a,"-o")}
for (frag in a) { split(a[frag],sz,"-b"); used+=sz[2] }
END { print used/2048/1024 }
)
echo $cap $used | nawk '{print $1-$2, $2/$1*100}' | read avail pct
devs=$(metastat -p | sort | nawk -v c=$c '$3==c && $2=="-p" {printf "%s ", $1}')
printf "%-20s %9.2f %9.2f %10.2f %9.1f%1s %s\n" $c $cap $used $avail $pct % "$devs"
done , #!/bin/ksh
#
# metafree - Adrian Ball (ade@majords.co.uk)- 2007/01
#
# Finds all devices containing soft partitions and shows
# size/allocation, unallocated space, devices etc
#
# NB: This is not the script that is listed on Sun's BigAdmin site
# I had written this before knowing of that one's existence, but
# subsequently found all the links to that one to be broken, so
# I have not seen it.
# Strangely I had picked the same name for the script...
# Maybe someone will find this one useful.

if [[ ! -x /usr/sbin/metastat ]] ; then
echo "Can't run metastat binary, bogus..."
exit 1
fi

printf "%-20s %9s %9s %10s %10s %s\n" Device GB Used Available Capacity "Soft Partitions"

for c in $(metastat -p | nawk '$2=="-p" {print $3}' | sort -u) ; do
if [[ ${c%%[0-9]*} == d ]] ; then
### It's a metadvice, use metastat to find the size
cap=$(metastat $c 2>/dev/null | nawk '/Size:/{print $2/2048/1024;exit}')
elif [[ ${c#}/ != $c ]] ; then
### It's a real device, work out size assuming full disk/LUN is used
cap=$(prtvtoc $c | nawk '
/bytes\sector/ {bs=$2}
/sectors\cyl/ {sc=$2}
/accessible cyl/ {as=$2}
END {print bs*sc*as/1024^3}')
else
### Assume we have a cXtX form of device, use prtvtoc to work out
### the size of the slice
cap=$(prtvtoc /dev/rdisk/$c | nawk -v c=$c '
BEGIN {split(c,a,"s");s=a[2]}
/bytes\sector/ {bs=$2}
$1==s {print $5*bs/1024^3}')
fi
used=$(metastat -p | nawk -v c=$c '
$2=="-p" && $3==c {split($0,a,"-o")}
for (frag in a) { split(a[frag],sz,"-b"); used+=sz[2] }
END { print used/2048/1024 }
)

```

```

echo $cap $used | nawk '{print $1-$2, $2/$1*100}' | read avail pct
devs=$(metastat -p | sort | nawk -v c=$c '$3==c && $2="-p" {printf "%s ", $1}')
printf "%-20s %9.2f %9.2f %10.2f %9.1f%1s %s\n" $c $cap $used $avail $pct % "$devs"
done /*100}' | read avail pct
devs=$(metastat -p | sort | nawk -v c=$c '==c && #!/bin/ksh
#
# metafree - Adrian Ball (ade@majords.co.uk)- 2007/01
#
# Finds all devices containing soft partitions and shows
# size/allocation, unallocated space, devices etc
#
# NB: This is not the script that is listed on Sun's BigAdmin site
# I had written this before knowing of that one's existence, but
# subsequently found all the links to that one to be broken, so
# I have not seen it.
# Strangely I had picked the same name for the script...
# Maybe someone will find this one useful.

if [[ ! -x /usr/sbin/metastat ]] ; then
echo "Can't run metastat binary, bogus..."
exit 1
fi

printf "%-20s %9s %9s %10s %10s %s\n" Device GB Used Available Capacity "Soft Partitions"

for c in $(metastat -p | nawk '$2=="-p" {print $3}' | sort -u) ; do
if [[ ${c%%[0-9]*} == d ]] ; then
### It's a metadvice, use metastat to find the size
cap=$(metastat $c 2>/dev/null | nawk '/Size:/ {print $2/2048/1024;exit}')
elif [[ ${c#}/ != $c ]] ; then
### It's a real device, work out size assuming full disk/LUN is used
cap=$(prtvtoc $c | nawk '
/bytes\sector/ {bs=$2}
/sectors\cyl/ {sc=$2}
/accessible cyl/ {as=$2}
END {print bs*sc*as/1024^3}')
else
### Assume we have a cXtX form of device, use prtvtoc to work out
### the size of the slice
cap=$(prtvtoc /dev/rdisk/$c | nawk -v c=$c '
BEGIN {split(c,a,"s");s=a[2]}
/bytes\sector/ {bs=$2}
$1==s {print $5*bs/1024^3}')
fi

used=$(metastat -p | nawk -v c=$c '
$2=="-p" && $3==c {split($0,a,"-o")
for (frag in a) { split(a[frag],sz,"-b"); used+=sz[2] } }
END { print used/2048/1024 }
')
echo $cap $used | nawk '{print $1-$2, $2/$1*100}' | read avail pct
devs=$(metastat -p | sort | nawk -v c=$c '$3==c && $2="-p" {printf "%s ", $1}')
printf "%-20s %9.2f %9.2f %10.2f %9.1f%1s %s\n" $c $cap $used $avail $pct % "$devs"

```

```
done="-p" {printf "%s ", }'  
printf "%-20s %9.2f %9.2f %10.2f %9.1f%1s %s\n" $c $cap $used $avail $pct % "$devs"  
done
```

Posted by rdircio in Solaris at 17:03

Sunday, September 2. 2007

## NX and Freenx tarball for Solaris 10 SPARC

After doing the slackware build, i felt i could do the solaris 10 build... it required some tweaking on the freenx scripts to make it all work. This is both freenx 0.7.0 and the nomachine NX libraries version 2.1.0-2.

Download [nx-sol10-sparc.tar.gz](http://www.nomachine.com/download-package.php?Prod_Id=65)

place it in your server at /, and uncompress.

```
# cd /
```

```
# gunzip nx-sol10-sparc.tar.gz
```

```
# tar xvf nx-sol10-sparc.tar
```

Then run the nxinstall script, which will check for the required packages. all of them from blastwave.org.

**VERY**

**IMPORTANT: if these are missing, install them!!!! nx+freenx won't work without ggetopt, expect, gmd5sum, gfind, ggetopt,netcat from blastwave at /opt/csw/bin!!!**

```
# nxinstall
```

a log of a succesful nxinstall should look like [this](#)

After this, you'll be ready to use any nx client on another computer to connect to your Solaris 10 SPARC freenx server!

Then install the client on your PC from:

[http://www.nomachine.com/download-package.php?Prod\\_Id=65](http://www.nomachine.com/download-package.php?Prod_Id=65)

And connect to your solaris server :)

ENJOY

Posted by rdircio in Solaris at 20:01

Saturday, May 19. 2007

## RDAC

if you see some "rdnexus" strings in your format output, you have RDAC (and yes, it is a disease)

```
32. c9t0d14 <STK-OPENstorageD220-0533 cyl 34965 alt 2 hd 64 sec 64>
/pseudo/rdnexus@4/rdriver@0,e
```

RDAC is a multipathing software, (like vxdump or powerpath). Other symptom includes finding the module loaded.

```
bash-2.03# modinfo | grep rd
42 103231bf 8f2 235 1 rdnexus (RDAC Pseudo Controller Driver)
43 1032391b 19fb4 236 1 rdriver (Solaris RDAC Driver)
```

well, you can find out how the paths are shown by issuing:

```
bash-2.03# SMdevices
SANtricity Storage Manager Devices, Version 08.30.02.00
Built Wed Sep 11 15:25:19 CDT 2002
Copyright (C) 2001 LSI Logic Corporation. All rights reserved.
Copyright (C) 2001 Storage Technology Corporation. All rights reserved.
/dev/rdisk/c5t0d11s2 [Disk Subsystem ImageDisk1, Volume Volume 12, LUN 11, Volume WWN
<600a0b8000f9d5300000014464c4753>, Preferred Path (Controller-B): In Use ]
/dev/rdisk/c5t0d13s2 [Disk Subsystem ImageDisk1, Volume Volume 14, LUN 13, Volume WWN
<600a0b8000f9d5300000015464c47a1>, Preferred Path (Controller-B): In Use ]
/dev/rdisk/c6t0d3s2 [Disk Subsystem ImageDisk2, Volume Volume 4, LUN 3, Volume WWN
<600a0b80001318bd000010c540605ca1>, Preferred Path (Controller-B): In Use]
/dev/rdisk/c6t0d1s2 [Disk Subsystem ImageDisk2, Volume Volume 2, LUN 1, Volume WWN
<600a0b80001318bd000010bd400ceaa7>, Preferred Path (Controller-B): In Use]
/dev/rdisk/c3t0d0s2 [Disk Subsystem ImageDisk2, Volume Volume 1, LUN 0, Volume WWN
<600a0b800013186c000010d9400cea9e>, Preferred Path (Controller-A): In Use]
/dev/rdisk/c3t0d2s2 [Disk Subsystem ImageDisk2, Volume Volume 3, LUN 2, Volume WWN
<600a0b800013186c000010e440605c8e>, Preferred Path (Controller-A): In Use]
/dev/rdisk/c3t0d10s2 [Disk Subsystem ImageDisk2, Volume Volume 11, LUN 10, Volume WWN
<600a0b800013186c000010ef464c4d3a>, Preferred Path (Controller-A): In Use ]
/dev/rdisk/c9t0d14s2 [Disk Subsystem ImageDisk1, Volume Volume 15, LUN 14, Volume WWN
<600a0b80001256a900000038464c4ca2>, Preferred Path (Controller-A): In Use ]
```

Posted by rdircio in Solaris at 11:22

Friday, May 4, 2007

## **TSM logs**

Again another basic tivoli thing... in solaris TSM logs are at:

`/var/tivoli`

Posted by rdircio in Solaris at 11:05

## Make a copy of your root disk on solaris 10

if you need to make a copy of your boot disk, and you don't have disksuite, and you only have one slice as "/", and this is solaris 10, you can do:

```
#prtvtoc /dev/rdisk/c0t0d0s0 | fmthard -s - /dev/rdisk/c0t1d0s0

#newfs /dev/dsk/c0t1d0s0

#mount /dev/dsk/c0t1d0s0 /a

#cd /

#find . -mount | cpio -pmdv /a

#installboot /usr/platform/`uname -i`/lib/fs/ufs/bootblk /dev/rdisk/c0t1d0s0

#vi /a/etc/vfstab(change references from c0t0d0s0 to c1)

#mkdir /a/tmp

#mkdir /a/dev

#mkdir /a/proc

#touch /a/etc/mnttab

#mkdir /a/etc/svc/volatile

#mkdir /a/system/object

#mkdir /a/system/contract

#umount /a

#init 0

#boot otherdisk -r
```

And the system will boot fine from the other disk.

Remember it should be better to do this in single user mode, and with apps down.

Posted by rdircio in Solaris at 15:36

Wednesday, February 14, 2007

## **zlogin escape character**

You can bring a server to the LOM if you don't specify a different escape sequence for zlogin.

To login to the console, for example, you can

```
# zlogin -e @ -C thezone
```

That will make "@" the escape sequence to get out of the zone's console

Posted by rdircio in Solaris at 06:56

Tuesday, February 6, 2007

## Build openssh with chroot jail patch on Solaris 10

The compiler is gcc version 3.4.5, from blastwave.

Got the sources for:

zlib-1.2.3  
openssl-0.9.8d  
openssh-4.5.1

Untar them all and apply the chroot patch "openssh-4.5p1-sshjail.patch" to openssh-4.5.1

Build them in that order.

To build zlib do:

```
# ./configure --prefix=/opt/acsssh --shared  
# make && make install
```

To build openssl:

```
# ./config shared --prefix=/opt/acsssh  
# make && make install
```

To build openssh

```
#!/usr/bin/ksh  
# ./configure --with-pam --prefix=/opt/acsssh --with-ssl-dir=/opt/acsssh --with-zlib=/opt/acsssh  
# make && make install
```

To build the chrooted jail i modded a script, jail.ksh, with contents:

```
#!/usr/bin/ksh  
JAILUSER=juser  
JAILGROUP=jgroup  
echo "deleting previous structures"  
userdel -r $JAILUSER  
groupdel $JAILGROUP  
rm -fR /export/home/jail  
echo "deleted"  
SDIR=opt/acsssh  
/usr/sbin/groupadd $JAILGROUP  
mkdir -p /export/home/jail  
chown root:$JAILGROUP /export/home/jail  
chmod 750 /export/home/jail  
/usr/sbin/useradd -g $JAILGROUP -c "Jail user $JAILUSER" \  
-d /export/home/jail/$JAILUSER/export/home/$JAILUSER -s /bin/sh $JAILUSER  
mkdir -p /export/home/jail/$JAILUSER
```

```

chown $JAILUSER:$JAILGROUP /export/home/jail/$JAILUSER
cd /export/home/jail/$JAILUSER
mkdir etc
mkdir bin
mkdir usr
mkdir usr/bin
mkdir -p $SDIR/bin
mkdir -p $SDIR/lib
mkdir -p $SDIR/sbin
mkdir -p $SDIR/etc
mkdir -p $SDIR/ssl/lib
mkdir -p $SDIR/libexec
mkdir usr/lib
mkdir usr/platform
mkdir usr/platform/`uname -i`
mkdir usr/platform/`uname -i`/lib
mkdir -p usr/platform/sun4u/lib/sparcv9
mkdir -p platform/sun4u/lib/sparcv9
mkdir lib
mkdir dev
mkdir devices
mkdir devices/pseudo
cd /export/home/jail/$JAILUSER
APPS='bin/cp bin/ls bin/mkdir bin/mv bin/pwd bin/rm bin/rmdir bin/sh'
for i in $APPS; do
cp /$i ./$i
LIBS=`ldd ./$i | awk '{print $3}'`
for l in $LIBS; do
if [[ ! -d `dirname $l` ]]; then
mkdir `dirname $l` > /dev/null
fi
cp $l ./$l
done
done
cd /export/home/jail/$JAILUSER/devices/pseudo
mknod mm@0:zero c 13 12
mknod mm@0:null c 13 2
cd /export/home/jail/$JAILUSER/dev
ln -s ../devices/pseudo/mm@0:zero zero
ln -s ../devices/pseudo/mm@0:null null
cd /export/home/jail/$JAILUSER
BINS="lib/ld.so.1 $SDIR/bin/ssh $SDIR/libexec/sftp-server $SDIR/sbin/sshd $SDIR/lib/libz.so
$SDIR/lib/libcrypto.so.0.9.8 usr/lib/ld.so.1 usr/lib/nss_files.so.1 platform/sun4u/lib/libc_psr.so.1
platform/sun4u/lib/sparcv9/libc_psr.so.1"
for i in $BINS; do
cp /$i ./$i
done
mkdir -p /export/home/jail/$JAILUSER/export/home/$JAILUSER
chown $JAILUSER:$JAILGROUP /export/home/jail/$JAILUSER/export/home/$JAILUSER
touch /export/home/jail/$JAILUSER/etc/passwd
touch /export/home/jail/$JAILUSER/etc/group
echo "$JAILUSER:x:`usr/xpg4/bin/id -u $JAILUSER`:`usr/xpg4/bin/id -g
$JAILUSER`::/export/home/$JAILUSER:/bin/sh" > /export/home/jail/$JAILUSER/etc/passwd

```

```
echo "$JAILGROUP::`/usr/xpg4/bin/id -g $JAILUSER`:$JAILUSER" > \
/export/home/jail/$JAILUSER/etc/group
echo "done!"
```

get it [here!](#)

Posted by rdircio in Solaris at 21:07

## Powerpath, Mpxio or Vxdmp

You want to know what DMP technology you are using at your server, and if having it activated, how do you see the disks....

Wrote this script:

```
#!/bin/ksh
echo "running check"
(
echo "MPXIO ?"
echo "-----"
modinfo | grep -i vhci
M=`modinfo |grep -i vhci | wc -l`
if [ $M -gt 0 ];then
echo "enabled"
else
echo "disabled"
fi
echo ""
echo "POWERPATH ?"
echo "-----"
modinfo | grep -i emcp
P=`modinfo |grep -i emcp | wc -l`
if [ $P -gt 0 ];then
echo "enabled"
else
echo "disabled"
fi
echo ""
echo "VXDMP ?"
echo "-----"
modinfo | grep -i vxdmp
V=`modinfo |grep -i vxdmp | wc -l`
if [ $V -gt 0 ];then
echo "enabled"
else
echo "disabled"
fi
echo ""
if [ $P -gt 0 ];then
echo "POWERPATH"
echo "-----"
/etc/powermt display dev=all
fi
if [ $V -gt 0 ];then
echo "VXDMP"
echo "-----"
```

```

vxdmpadm listctlr all
fi
if [ $M -gt 0 ];then
echo "LUXADM"
echo "-----"
for d in `format < /dev/null| egrep 'HIT|EMC' | awk '{ print $2 }';do
echo $d
luxadm display /dev/rdisk/${d}s2
done
fi
if [ -f /usr/sbin/vxdisk ];then
echo "VXDISK LIST"
echo "-----"
for d in `usr/sbin/vxdisk -o alldgs list| awk '{ print $1 }';do
echo $d
/usr/sbin/vxdisk list $d
done
fi
) > /tmp/results_`uname -n`.txt
echo "finished running check"

```

Run it, relax, and enjoy

Thx to Hiram Ruiz and Corey Joiner

Posted by rdircio in Solaris at 07:55

Monday, January 29. 2007

## Restore file from TSM

Issue

```
# dsmc rest -pick=yes /some/file /destination/file
```

This will show you a menu of different backup dates for /some/file, and after you pick one, you will be able to restore it

To restore from a file list, include inactive copies and do not replace files:

```
# dsmc restore -replace=no -ina -filelist=f
```

Hey!, and if you have access to X on that server you can also use the gui:

```
# dsm
```

Thanks to Fergus O'Hea

Posted by rdircio in Solaris at 12:21

Friday, December 8, 2006

### **create big concat using svm (disksuite)**

Please use the right syntax to concatenate disks and  
create a big metadvice

```
# metainit d40 4 1 c0t1d0s0 1 c0t2d0s0 1 c0t2d0s0 1 c0t2d1s0
```

the "4" is for the number of disks we are using.

Posted by rdircio in Solaris at 06:28

Sunday, October 22, 2006

## **To re-attach a failed disk on VXVM**

vxmend - recover plexes that are in a "DISABLED/RECOVER" state.

vxdiskadm option 1

(/usr/sbin/vxmend)

vxmend -o force off <plex>

vxmend on <plex>

vxmend fix clean <plex>

vxvol start <volume>

fsck the filesystems and mount them

This is Hiram Ruiz's

Posted by rdircio in Solaris at 18:48

Wednesday, October 4, 2006

## routes

If when you do netstat -rn you see routes that look the same...

```
bash-3.00$ netstat -rn
Routing Table: IPv4
Destination Gateway Flags Ref Use Interface
-----
198.148.129.64 198.148.129.110 U 1 385521 ce0:1
198.148.129.64 198.148.129.98 U 1 0 ce0:2
198.148.129.64 198.148.129.110 U 1 0 ce0
198.148.129.64 198.148.129.110 U 1 346097 qfe0
```

To differentiate them do a netstat -nr

```
bash-3.00$ netstat -nr
IRE Table: IPv4
Destination Mask Gateway Device Mxfrg Rtt Ref Flg Out In/Fwd
-----
198.148.129.64 255.255.255.192 198.148.129.110 ce0:1 1500* 0 1 U 385521 0
198.148.129.64 255.255.255.192 198.148.129.98 ce0:2 1500* 0 1 U 0 0
198.148.129.64 255.255.255.192 198.148.129.110 ce0 1500* 0 1 U 0 0
198.148.129.64 255.255.255.192 198.148.129.110 qfe0 1500* 0 1 U 346097 0
```

then you can add/delete, differentiating them by the netmask, like:

```
route delete 198.148.129.64 198.148.129.110 255.255.255.192
```

Thanks to Alejandro Siller

Posted by rdircio in Solaris at 13:22

Friday, September 29. 2006

### **enable ftp access, disable telnet for only some users**

To make "only ftp" users, meaning they can only use ftp, no telnet or ssh, do this:

add /bin/false to the list of shells in /etc/shells

change your user's shell to /bin/false

To make users be able to sftp & ftp, but not telnet or ssh change their shell to /usr/sbin/sftp-server

done!

Posted by rdircio in Solaris at 16:41

### **testing ipmp**

To make one of your interfaces fail:

```
# ifconfig adapter_name modinsert ldterm@2
```

To get it back

```
# ifconfig adapter_name modremove ldterm@2
```

You can also do the supported command:

```
#if_mpadm -d qfe1
```

to get it back

```
# if_mpadm -r qfe1
```

Posted by rdircio in Solaris at 16:38

Saturday, September 2. 2006

**nfs client entry for vfstab on solaris**

```
somehost:/shared - /mydir nfs - yes rw,soft
```

Posted by rdircio in Solaris at 17:16

Friday, September 1. 2006

## Netbackup 5.1 client on Solaris 10

If you push the netbackup client from the master server to a solaris 10 box it thinks it is solaris 8, so it writes entries to `/etc/inetd.conf`

Solaris 10 does not have `inetd` but `SMF`, so the `bpcd` port is not activated, you can test by doing a `telnet localhost bpcd...` it does not answer.

So you have to convert your `/etc/inetd.conf` file to `SMF` descriptors, you do so by running

```
# inetconv
```

You make sure it is imported by doing

```
# inetadm | grep -i bpcd  
# telnet localhost bpcd
```

if you get a connection to the `bpcd` port you're all set and you're ready for backup/restore

:)

Posted by `rdircio` in `Solaris` at 23:30

Saturday, August 26. 2006

## Find a file in package contents

Sometimes you need to know in which package a particular file is in

```
grep <full_path_to_filename> /var/sadm/install/contents| awk '{ print $1 " " $10 }'
```

Posted by rdircio in Solaris at 08:46

## **solaris patchadd return codes**

0. No error
1. Usage error
2. Attempt to apply a patch that's already been applied
3. Effective UID is not root
4. Attempt to save original files failed
5. pkgadd failed
6. Patch is obsoleted
7. Invalid package directory
8. Attempting to patch a package that is not installed
9. Cannot access /usr/sbin/pkgadd (client problem)
10. Package validation errors
11. Error adding patch to root template
12. Patch script terminated due to signal
13. Symbolic link included in patch
14. NOT USED
15. The prepatch script had a return code other than 0.
16. The postpatch script had a return code other than 0.
17. Mismatch of the -d option between a previous patch install and the current one.
18. Not enough space in the file systems that are targets of the patch.
19. \$SOFTINFO/INST\_RELEASE file not found
20. A direct instance patch was required but not found
21. The required patches have not been installed on the manager
22. A progressive instance patch was required but not found
23. A restricted patch is already applied to the package
24. An incompatible patch is applied
25. A required patch is not applied
26. The user specified backout data can't be found
27. The relative directory supplied can't be found
28. A pkginfo file is corrupt or missing
29. Bad patch ID format
30. Dryrun failure(s)
31. Path given for -C option is invalid
32. Must be running Solaris 2.6 or greater
33. Bad formatted patch file or patch file not found
34. Incorrect patch spool directory
35. Later revision already installed

Posted by rdircio in Solaris at 15:10

Wednesday, August 16. 2006

## **Change timezone in solaris**

Edit `/etc/TIMEZONE` and use one of the typical:

US/Eastern  
US/Central  
US/Mountain  
US/Pacific

or look for an appropriate timezone in `/usr/share/lib/zoneinfo`  
After this reboot your server with `shutdown` or `init`.

Posted by `rdircio` in Solaris at 10:13

Tuesday, August 15. 2006

## **sendmail queue management**

To know what is in your queue you use

```
# sendmail -bp
```

or

```
#mailq
```

To clear a job you just might delete it from `/var/spool/mqueue`

Posted by rdircio in Solaris at 14:18

Monday, August 14. 2006

## copy files using netbackup

You want to copy some files from and to a server and you don't have root access to it?

But you do have access to its netbackup master server?

yeah!

use bpgp

*The command should be executed on the Master Server.*

*The formats are:*

*A. To get a file from a client to the MasterServer:*

*bpgp from <client-name> <full-pathname-of-file-to-get> <full-pathname-of file-on-MS>*

*B. To put a file from the Master Server to a client:*

*bpgp to <client-name> <full-pathname-of file-on-MS> <full-path-name-of-file-to create-on-client>*

thx to Alberto Vazquez :)

Posted by rdircio in Solaris at 13:25

Tuesday, July 25. 2006

## **rstatd performance daemon**

Well, i didn't know that rstatd can be used to obtain a remote's server performance counters....

It is usually disabled for security purposes.

Just make sure it is configured at /etc/inetd.conf, by uncommenting the proper line

```
bash-2.05# grep rstat /etc/services  
bash-2.05# grep rstat /etc/inetd.conf  
rstatd/2-4 tli rpc/datagram_v wait root /usr/lib/netsvc/rstat/rpc.rstatd rpc.rstatd
```

Then restart inetd

```
bash-2.05# ps -ef | grep inetd  
root 355 1 0 Apr 26 ? 0:00 /usr/sbin/inetd -s -t  
bash-2.05# kill -HUP 355
```

It will be running

```
bash-2.05# ps -ef | grep rstat  
root 23620 23616 0 18:01:48 pts/1 0:00 grep rstat  
root 21694 355 0 17:39:37 ? 0:01 rpc.rstatd  
You can use a fine rstatd client called java perfmeter: http://jperfmeter.sourceforge.net/
```

If you have CDE you can use sdtperfmeter

```
bash-3.00$ /usr/dt/bin/sdtperfmeter othermachine
```

If you have gnome, you might as well use gnome-perfmeter

```
# gnome-perfmeter othermachine
```

Posted by rdircio in Solaris at 18:02

Monday, July 24. 2006

## Unresponsive keyboard on vnc with gnome sessions on Solaris

You might get a problem when using a gnome session on Solaris with vnc where keyboard is unresponsive....

According to <http://sunsolve.sun.com/search/document.do?assetkey=1-9-81175-1> :

To work around this issue, touch the following file:

*/etc/gconf/gconf.xml.defaults/apps/gnome\_settings\_daemon/keybindings/%gconf.xml*

and edit the file with:

```
<?xml version="1.0"?>
<gconf>
<entry name="volume_up" mtime="1110896708" type="string">
<stringvalue></stringvalue>
</entry>
<entry name="volume_mute" mtime="1110896705" type="string">
<stringvalue></stringvalue>
</entry>
<entry name="volume_down" mtime="1110896702" type="string">
<stringvalue></stringvalue>
</entry>
<entry name="help" mtime="1110896698" type="string">
<stringvalue></stringvalue>
</entry>
</gconf>
```

Then just stop and start your vnc server

Posted by rdircio in Solaris at 13:56

Sunday, July 23, 2006

## Suntrunking 1.3 install

This is only how the install looks on solaris 9, not much options. It looks as it can be used only on qfe, ce and ge interfaces. We'll later see how to configure it

```
bash-2.05# pwd
```

```
/tmp/SUNTRUNKING/suntrunking
```

```
bash-2.05# ls
```

```
Copyright  FR_Copyright  README  Solaris_8  install
```

```
Docs      License  Solaris_7  Solaris_9  remove
```

```
bash-2.05# ./install
```

*Sun Trunking 1.3 Utility Installation.*

*Copyright 2003 Sun Microsystems, Inc. All rights reserved.  
Use is subject to license terms.*

**\*\* Checking for existing Sun Trunking 1.3 Utility packages...**

**\*\*Checking for supporting Ethernet drivers and patches...**

*This script is about to take the following actions:*

- Install Sun Trunking 1.3 Utility packages.*
- Install Vlan Ethernet Utility: 114600-02 required patch(es).*
- Install Gigabit Ethernet: 113361-06 required patch(es).*

*Press return to continue, or 'Q' followed by a return to quit:*

**\*\* Installing patch 114600-02 for Solaris 9...**

**Checking installed patches...**

**Verifying sufficient filesystem capacity (dry run method)...**

**Installing patch packages...**

**Patch number 114600-02 has been successfully installed.**

**See /var/sadm/patch/114600-02/log for details**

**Patch packages installed:**

**SUNWvld**

**SUNWvldu**

**SUNWvldx**

**\*\*Installing patch 113361-06 for Solaris 9...**

*Checking installed patches...*

*This patch is obsoleted by patch 112233-12 which has already been applied to this system.*

*Patchadd is terminating.*

**\*\* Installing Sun Trunking 1.3 Utility packages...**

**Copyright 2003 Sun Microsystems, Inc. All rights reserved.**

**Use is subject to license terms.**

---

---

**Please modify /etc/opt/SUNWconn/bin/nettr.sh to set up trunking configuration.**

**Please create /etc/hostname.qfeN entries for QFE trunk heads.**

**Please create /etc/hostname.geN entries for GEM trunk heads.**

**Please create /etc/hostname.ceN entries for CE trunk heads.**

**Reboot your system.**

---

---

**Installation of <SUNWtrku> was successful.**

**Copyright 2003 Sun Microsystems, Inc. All rights reserved.**

**Use is subject to license terms.**

**Installation of <SUNWtrkm> was successful.**

**\*\*Done.**

*A log of this Install can be found at:*

*/var/tmp/Trunking.install.2006.07.25*

Posted by rdircio in Solaris at 21:51

Wednesday, July 19. 2006

## Calculate your netmasks in hex & decimal

if you're bad at transforming hex to dec or viceversa

<http://www.softwareeffect.com/tools/netmask.html>

- rdircio

Posted by rdircio in Solaris at 16:18

Tuesday, July 18. 2006

### **Mount cdrom without vold**

If your vold is fsck'ed up:

```
# mount -F hsfs -o nomap,case,ro /dev/sr0 /CDROM
```

Thanks to Hiram Ruiz

- rdircio

Posted by rdircio in Solaris at 17:46

Monday, July 17, 2006

## Get epoch in decimals

### 1. Solaris 2.6 / 7

```
1. /usr/bin/truss /usr/bin/date 2>&1 | /usr/bin/awk '/^time/ {print $NF}'
```

### 2. Solaris 8 / 9 / 10

```
1. /usr/bin/perl -e 'printf "%d\n", time;'
```

Posted by rdircio in Solaris at 15:35

Monday, July 10, 2006

## **tar and gzip**

this is first grade, but in case you forget how to create a tar and gzip on the fly if you do not have gnu tar...

to unpack on the fly:

```
# gunzip < filename.tar.gz | tar xvf -
```

to pack on the fly:

```
# tar cvf - files_to_tar | gzip -c > filename.tar.gz
```

Posted by rdircio in Solaris at 13:11

Tuesday, June 13. 2006

### **is your solaris running on 64bit or/and 32bit?**

Just type

```
# isainfo -v
```

Posted by rdircio in Solaris at 11:39

### **Get WWN tree info for an hba**

To get all the wwns attached to an hba you can use this

```
# cfgadm -al -o show_FCP_dev | egrep 'disk|tape'
```

Posted by rdircio in Solaris at 09:01

Monday, June 12, 2006

## Add new filesystem to existing zone

You have to create it in the global zone, add it with zoneconfig, then mount it manually from the global zone

```
bash-3.00# zonecfg -z zones004z3
zonecfg:zones004z3> add fs
zonecfg:zones004z3:fs> set dir=/myzone/cpbatch
zonecfg:zones004z3:fs> set special=/dev/md/dsk/d132
zonecfg:zones004z3:fs> set raw=/dev/md/rdisk/d132
zonecfg:zones004z3:fs> set type=ufs
zonecfg:zones004z3:fs> end
zonecfg:zones004z3> verify
zonecfg:zones004z3> commit
zonecfg:zones004z3> exit
bash-3.00# mount /dev/md/dsk/d132 /zones/CHB/myzone/cpbatch
```

Posted by rdircio in Solaris at 17:45

Sunday, June 4, 2006

### **Create/grow volume with maxsize (Veritas)**

Create volume of maxsize (max size of all the DG's disks):

```
# vxassist -g ora2dg maxsize
```

```
Maximum volume size: 39522304 (19298Mb)
```

```
# vxassist -g ora2dg make oratmp 39522304
```

```
# mkfs -F vxfs /dev/vx/rdisk/ora2dg/oratmp
```

```
# mkdir /db/prd3/archive_test2
```

```
# mount -F vxfs /dev/vx/dsk/ora2dg/oratmp /db/prd3/archive_test2
```

*Grow a volume to maxgrow:*

```
# vxassist -g ora1dg maxgrow archive_temp
```

```
Volume archive_temp can be extended by 10772480 to 94658560 (46220Mb)
```

```
# vxresize -g ora1dg archive_temp +4g
```

Posted by rdircio in Solaris at 13:47

Friday, June 2. 2006

### **Get disk drive information using luxadm**

```
format < /dev/null | grep -iv search | grep -i "\." | awk '{ print $2 }' | while read line;do luxadm display /dev/rdisk/"$line"s2;
done
```

Posted by rdircio in Solaris at 12:51

Thursday, May 25. 2006

## **Disk Array DMP support in veritas volume manager**

You might want to know if you have support for dmp in your array for veritas, you do this by checking the right ASL package is installed.

```
# vxddladm listsupport
```

Then check if there's an entry for your type of disk array

Posted by rdircio in Solaris at 12:41

Wednesday, May 24. 2006

## NetBackup 101 (for me)

ok, just hang with me while I learn NB.

some jobs were stalled in a machine, so we went to the media server and checked if the jobs were pending and the tape drives were not in "TSD":

```
bash-3.00# /usr/opensv/volmgr/bin/vmoprcmd
PENDING REQUESTS
<NONE>
DRIVE STATUS
Drv Type Control User Label RVSN EVSN Ready Wr.Enbl. ReqId
0 dlt TSD - No - -
1 dlt TSD root Yes EBI219 EBI219 Yes Yes 1
ADDITIONAL DRIVE STATUS
Drv DriveName Multihost Assigned Comment
0 QUANTUMDLT70000 No -
1 QUANTUMDLT70001 No ebitdb1
```

In the "pending" section there were jobs pending on drv 1 and 0, and the tape drives were unavailable, so we reset them:

```
# vmoprcmd -deny 0
# vmoprcmd -deny 1
# vmoprcmd -up 0
# vmoprcmd -up 1
```

We assumed the drives were recently fed with tapes, so after this jobs were running ok.

To see the active jobs...

```
#bpdjobs -report |grep -i act
```

To see the queued jobs...

```
#bpdjobs -report |grep -i que
```

To kill a job

```
#bpdjobs -kill <job#>
```

To see the job summary...

```
##/usr/opensv/netbackup/bin/bpdjobs -summary
#bpps -a
```

To run an inventory on a tape library...

```
#bpinvent 0
```

Check last 24 hour's errors

```
#bperror -U -backstat -hoursago 24
```

See available tapes

```
# available_media
```

Or go to bpadm - media mgmt - special actions - Inventory a Robot and Compare with Volume Configuration

To see if a tape is in the correct pool

```
#vmquery -m <mediaid>
```

To change a tape of pool

```
# vmchange -h ntxldb2 -p 1 -m NTXD53
```

(Where 1 is NetBackup and 2 is systems)

To see the pools

```
#vmpool -listall
```

A monitoring loop:

```
# while true;do clear; vmopr cmd; echo "active"; bpdjobs -report |grep -i act;echo "queued"; bpdjobs -report |grep -i que;bpdjobs -summary;bperror -U -backstat -hoursago 24| tail -15; sleep 5; clear; done
```

Run bpadm for interactive actions

```
# bpadm
```

Please check the netbackup quick reference card [document](#) too

Posted by rdircio in Solaris at 15:20

Tuesday, May 23. 2006

## **soft partition create, grow, delete on SVM (Sun Volume Manager/Solstice DiskSuite)**

You have a partition d50 comprising slice 0 of a disk:

```
d50 1 1 c7t60060E80042AEE0000002AEE00000019d0s0
```

now you want to create soft partitions in it, you can create a 5gb soft partition just by:

```
# metainit d55 -p d50 5gb  
# newfs /dev/md/dsk/d55
```

You can destroy it by unmounting it, and clearing it

```
# metaclear d55
```

You can grow this partition

```
# metattach d55 1gb  
# growfs -M /mount/point /dev/md/rdisk/d55
```

Posted by rdircio in Solaris at 14:30

Saturday, May 13. 2006

## Configure the VCS hagui service

for those who feel comfortable with the hagui in Veritas Cluster, and cannot use it because you don't have the daemon running, you must add this section to your `/etc/VRTSvcs/conf/config/main.cf`:

```
group ClusterService (  
  
  SystemList = { hst101 = 0, hst102 = 1 }  
  AutoStartList = { hst101, hst102 }  
  OnlineRetryLimit = 3  
  Tag = CSG  
)
```

```
IP webip (  
  Device = eth0  
  Address = "172.20.162.15"  
  NetMask = "255.255.254.0"  
)
```

Where you need to change "hst101" and "hst102" to the hostnames of your nodes, and configure your network interface accordingly.

Posted by rdircio in Solaris at 22:44

## Edit VCS main.cf using vi

Sometimes you need to edit your veritas cluster configuration file "main.cf" and you feel the urge to do it using vi, here's how to do it:

```
hastop -all  
cp main.cf main.cf.copy  
vi main.cf  
hacf -verify /etc/VRTSvcs/conf/config  
hacf -generate /etc/VRTSvcs/conf/config  
hastart
```

Posted by rdircio in Solaris at 22:37

Friday, May 12. 2006

## no lsof?

```
cd /proc
pfiles * | egrep "^[0-9]sockname" | more
```

Or use PCP, found at <http://www.unix.ms/pcp/>

```
#!/usr/bin/ksh
#
# PCP (PID con Port)
# v1.07 20/05/2008 sam@unix.ms
#
# If you have a Solaris 8, 9 or 10 box and you can't
# install lsof, try this. It maps PIDS to ports and vice versa.
# It also shows you which peers are connected on which port.
# Wildcards are accepted for -p and -P options.
#
# The script borrows from Eric Steed's excellent "getport.sh" script.
#
#
if [ $# -lt 1 ]
then
echo >&2 "usage: $0 [-p PORT] [-P PID] [-a ALL ] (Wildcards OK)"
exit 1
fi
while getopts :p:P:a opt
do
case "${opt}" in
p ) port=${OPTARG};;
P ) pid=${OPTARG};;
a ) all=all;;
[?]) # unknown flag
echo >&2 "usage: $0 [-p PORT] [-P PID] [-a ALL ] (Wildcards OK) "
exit 1;;
esac
done
shift `expr $OPTIND - 1`
if [ $port ]
then
# Enter the port number, get the PID
#
port=${OPTARG}
echo "PID\tProcess Name and Port"
echo "_____ "
for proc in `ptree -a | grep -v ptree | awk '{print $1};`
do
```

```

result=`pfiles $proc 2> /dev/null| grep "port: $port"`
if [ ! -z "$result" ]
then
program=`ps -fo comm -p $proc | tail -1`
echo "$proc\t$program\t$port\n$result"
echo "_____ "
fi
done
elif [ $pid ]
then
# Enter the PID, get the port
#
pid=$OPTARG
# Print out the information
echo "PID\tProcess Name and Port"
echo "_____ "
for proc in `ptree -a | grep -v ptree | grep $pid| awk '{print $1};`
do
result=`pfiles $proc 2> /dev/null| grep port:`
if [ ! -z "$result" ]
then
program=`ps -fo comm -p $pid | tail -1`
echo "$proc\t$program\n$result"
echo "_____ "
fi
done
elif [ $all ]
then
# Show all PIDs, Ports and Peers
#
echo "PID\tProcess Name and Port"
echo "_____ "
for pid in `ptree -a | grep -v ptree |sort -n | awk '{print $1};`
do
out=`pfiles $pid 2>/dev/null| grep "port:"`
if [ ! -z "$out" ]
then
name=`ps -fo comm -p $pid | tail -1`
echo "$pid\t$name\n$out"
echo "_____ "
fi
done
fi
exit 0

```

Posted by rdircio in Solaris at 11:21

Tuesday, April 25. 2006

## A (very) poor man's nmap

I cannot install packages in a machine, so i wrote this very basic substitute to nmap called nmap.ksh

```
#!/bin/ksh
H=$1
X=0
L=600
while [ $X -lt $L ];do
    O=`echo "~." | telnet -r $H $X 2>&1 | grep -i connected | wc -l | sed 's/ //g'`
    if [ $O -eq 1 ];then
        echo "port $X open "
    fi
    X=`echo "$X + 1" | bc`
done
```

Posted by rdircio in Solaris at 16:45

Thursday, April 20. 2006

## Automate ssh login with expect

Sometimes you are not allowed to use keys to autologin to your machines, so you must find a way to automate your login using password authentication.

This expect script can help you do it:

```
#!/usr/local/bin/expect --
set timeout 30
spawn ssh user@ipaddress
while 1 {
expect {
"*(yes/no)?" { send "yes\n"}
"*password:" { send "password\n"}
"*user]#" { send "shutdown -h now\nexit\n"}
}
}
```

Posted by rdircio in Solaris at 09:25

Wednesday, April 5, 2006

## **configuring mpzio**

I found a nice document on how to enable mpzio in solaris:

<http://www.petertribble.co.uk/Solaris/mpzio.html>

Posted by rdircio in Solaris at 15:07

Tuesday, April 4, 2006

## **physical slot location**

Do you want to know where a I/O card is physically located in the server?

If you have the device path you can know.

Read this doc:

<http://sunsolve.sun.com/search/document.do?assetkey=1-9-21216>

Get the important part [here](#)

Posted by rdircio in Solaris at 15:06

## Raid Manager commands

Got these from elsewhere:

Term/Cmd Description Command / File

=====

SUNWosafw Open Storage Array Firmware

SUNWosamn Open Storage Array Man Pages

SUNWosar Open Storage Array (Root)

SUNWosau Open Storage Array (Usr)

arraymon Array monitor daemon < /usr/lib/osa/bin/rdaemon 28 140 > in background

rdaemon Redundant Controller Resolution Daemon < /usr/lib/osa/bin/arraymon > in background

rm6 Raid Manager GUI

# /usr/lib/osa/bin/rm6 &

raidutil RAID Management Configuration Utility

# raidutil -c c6t5d0s2 -A 65 1 -T 99 1 -f 1 1-C off 1 -m off 1 -p disable -S 8

raidutil Create new lun

# raidutil -c (raid controller) -g (disks) -l (raid level) -n (number of LUN) -s (capacity)

raidutil Create RAID5 lun 0

# raidutil -c c3t4d1s0 -s 1000 -n 0 -l 5 -g 11,21,31,41,51

raidutil Create RAID1 lun 2 (18g)

# raidutil -c c2t5d0s0 -l 1 -n 2 -r fast -s 17234 -g 50,11

raidutil Create Hot Spare

# raidutil -c (raid controller) -h (disk)

raidutil -c c2t5d0s0 -h 213

raidutil Delete a LUN

# raidutil -c (raid controller) -D (LUN number)

raidutil Delete lun 0

# raidutil -c c3t4d1s0 -D 0

raidutil Delete Hot Spare

# raidutil -c (raid controller) -H (disk)

raidutil -c c3t4d1s0 -H 113 (row,column 113=1,13)

lad List Array Devices

# lad

healthck perform a health check on RAID Modules # healthck -a

healthck Get raid module specifier

# healthck -a (ex: da103\_001: Optimal)

drivutil Display Drives associated with groups

# drivutil -d

drivutil Fail a Drive

```
# drivutil -f 01
drivutil Report Status or reconstruction status
# drivutil -p
drivutil Unfail a drive and reconstruct data
# drivutil -U
```

```
log Log file /usr/lib/osa/rmlog.log
logutil Log format utility - View log # logutil
```

```
rmparams Parameters file /usr/lib/osa/rmparams
```

```
rdacctl Correct problem of LUNS on different
# rm /dev/rdisk/c3*,/dev/rdisk/c5*,/dev/dsk/c3*
controllers on shared A3500 in cluster
# rm /dev/dsk/c5*,/dev/osa/dev/dsk,rdisk c5* c3*

# rdacctl config;drvconfig -i sd, disks
```

```
* BEGIN RAID Manager addition
```

```
* DO NOT EDIT from BEGIN above to END below...
```

```
forceload: drv/sd
```

```
forceload: drv/rdnexus
```

```
forceload: drv/rdriver
```

```
* END RAID Manager addition
```

```
Posted by rdircio in Solaris at 10:15
```

Wednesday, February 15, 2006

## **SAN migration paper**

I've written a detailed procedure on how to migrate a server using SAN disks as intermediate storage.

[Read it!](#)

it is fun.

Posted by rdircio in Solaris at 13:47

Tuesday, February 7. 2006

### change ip on created zone

```
zonecfg:z> select net address=<old-ip>
zonecfg:z:net> set address=<new-ip>
zonecfg:z:net> end
zonecfg:z> commit
```

Posted by rdircio in Solaris at 19:28

### copy dvd on solaris10 to an image

I've copied and used a DVD image from the eiscd DVD (enterprise installation standards), to add it to my remote JET(jumpstart enterprise toolkit) servers using:

```
1- insert the dvd in the drive
2- observe the mountpoint (cxtxdx)
# df -k | grep cdrom
/vol/dev/dsk/c0t1d0/eis-dvd 2290588 2290588    0 100% /cdrom/eis-dvd
3- Stop volume management, because it prevents from reading from the raw device
# /etc/init.d/volmgt stop
4- Use dd on slice 2
# dd if=/dev/dsk/c0t1d0s2 of=/eis dvd.iso
5- transfer eis dvd.iso to other machine, and test:
# mkdir /eis
# mount -F hsfs `lofiadm -a /eis dvd.iso` /eis
```

Success!

Posted by rdircio in Solaris at 12:50

Sunday, January 29. 2006

### **Configure resolution and depth in a Sun xvr-600 graphics card**

I have this graphics card in a sunblade150, and thought it could be configured using m64config. I tried to, but it did not work.

For this card, you must use "fbconfig"

```
# fbconfig -dev /dev/fbs/jfb0 -defdepth 24 -res 1280x1024x76 now
```

did the trick

Posted by rdircio in Solaris at 13:59

Wednesday, October 19. 2005

## non-root sessions with vnc for solaris sparc from sunfreeware

When you want to run a vncserver session on a non-root user from a solaris sparc machine, it just ends with this message:

```
_XSERVTransSocketCreateListener: failed to bind listener
```

After going around some places i found something to fix this at <http://www.webservertalk.com/archive100-2004-2-131777.html>

That involves creating a script in /usr/bin/fixXdir.ksh and putting it at the crontab, like this:

```
#!/usr/bin/ksh
    PATH="
        IFS=' '
            ice_dir=/tmp/.ICE-unix
                x11_dir=/tmp/.X11-unix
                    for dirname in $ice_dir $x11_dir
; do
    {
        if ([ -d $dirname ]
            then
                /usr/bin/mkdir $dirname
                    if [ $? != 0 ];
                        then
                            /usr/bin/echo "Warning: $0: cannot
mkdir $dirname.";
                                fi
                                    fi
                                        if ([ -d $dirname ]
                                            then
                                                /usr/bin/chown root:root $dirname
                                                    if [ $? != 0 ];
                                                        then
                                                            /usr/bin/echo "Warning: $0: cannot chown $dirname.";
                                                                fi
                                                                    /usr/bin/chmod 1777 $dirname
                                                                        if [ $? != 0 ];
                                                                            then
                                                                                /usr/bin/echo "Warning: $0: cannot chmod $dirname to 1777.";
                                                                                    fi
                                                                                        fi
                                                                                          }
done
```

The crontab line is this:

```
0,5,10,15,20,25,30,35,40,45,50,55 /usr/bin/fixXdir.ksh
```

Posted by rdircio in Solaris at 18:53

Tuesday, October 11. 2005

## Setting a local NTP server and client

You want to synchronize your solaris servers? don't have access to an atomic clock?

You must pick one of your solaris servers to be the time server, in this case i picked a sun server called "server". You will trust in this server's clock to be correct most of the time.

At "server" do:

```
root@server# cd /etc/inet
root@server# vi ntp.conf
```

And paste this inside:

```
server 127.127.1.0 prefer
fudge 127.127.1.0 stratum 0
broadcast 224.0.1.1 ttl 4
enable auth monitor
driftfile /var/ntp/ntp.drift
statsdir /var/ntp/ntpstats/
filegen peerstats file peerstats type day enable
filegen loopstats file loopstats type day enable
filegen clockstats file clockstats type day enable
keys /etc/inet/ntp.keys
trustedkey 0
requestkey 0
controlkey 0
```

Then, start ntp:

```
root@server# /etc/rc2.d/S74xntpd start
```

Verify its status by checking the process and querying the server:

```
root@server# ps -ef | grep ntp
root 8339 6236 0 11:39:16 pts/6 0:00 grep ntp
root 6725 1 0 11:15:48 ? 0:00 /usr/lib/inet/xntpd
root@server# ntpq -pn
remote refid st t when poll reach delay offset disp
```

---

```
*127.127.1.0 .LCL. 0 1 1 64 377 0.00 0.000 10.01
224.0.1.1 0.0.0.0 16 - - 64 0 0.00 0.000 16000.0
```

now ON THE CLIENT MACHINE:

```
root@client# cd /etc/inet
root@client# vi ntp.conf
```

Paste the contents:

```
driftfile /var/lib/ntp/ntp.drift
server 150.250.250.18
```

Notice that 150.250.250.18 is the ip of the timeserver called "server".

Now initialize the ntp daemon:

```
root@client# /etc/rc2.d/S74xntpd start
```

Check if it's running:

```
root@client # ps -ef | grep ntp
```

```
root 23579 1 0 11:20:12 ? 0:00 /usr/lib/inet/xntpd
```

```
root 1054 18332 0 11:41:34 pts/8 0:00 grep ntp
```

```
root@client # ntpq -pn
```

```
remote refid st t when poll reach delay offset disp
```

---

---

```
*150.250.250.18 .LCL. 1 u 4 64 377 18.83 0.473 1.80
```

That's it!

now type the "date" command, and see the results!

Posted by rdircio in Solaris at 11:41

Monday, October 3. 2005

## Solaris Recommended Patches Mirror

Playing around with wget i found the site where Sun keeps their patches, so we are now mirroring it, in our slackware machine, with a script in /etc/cron.daily/patches.sh, which contains the following:

```
#!/bin/sh
set -x
(
rm /software/SOLARIS/patches.sun.com/*index.html*
cd /software/SOLARIS/patches.sun.com
wget --tries=10 -r -nd -N http://patches.sun.com/clusters
) > /root/patches.log 2>&1
```

No more suffering for ugly patch downloads, we will always have the latest and greatest!

Posted by rdircio in Solaris at 16:33

Friday, September 30. 2005

## Free inodes

Sometimes your filesystems are full, but you don't know why!, check if you have enough free inodes first!

```
# df -k -o i
Filesystem      iused ifree %iused Mounted on
/dev/md/dsk/d0  11192 355144  3% /
/dev/md/dsk/d5  50042 286214 15% /usr
/dev/md/dsk/d4  21752 476872  4% /var
/dev/md/dsk/d2  17876 348460  5% /opt
```

if %iused=100% or near, you must tune your FS to have more inodes!

Posted by rdircio in Solaris at 22:11

Sunday, September 4, 2005

## Orca, the orcallator and the SE Toolkit

Before reading, look at my [orca monitored systems](#)!

After some weeks of hearing too much about performance tuning on solaris, I decided to walk on my own road, and found orca and the orcallator, which are tools related to the SE toolkit that use RRDtool, the same as used in MRTG, to put up some nice graphs on your system's performance.

We have several systems, so the approach is to have a central node in which we concentrate all performance data.

I decided to run the rrdtool, the image creation, and web page creation part in linux, and only the statistics collection in each solaris node.

The first step is to get orca at <http://www.orcaware.com/orca/>, the version I got was 0.27, at <http://www.orcaware.com/orca/pub/orca-0.27.tar.gz>, and put it somewhere in my Slackware 10.1 machine.

```
# tar zxvf orca-0.27.tar.gz
# cd orca-0.27
# ./configure -with-html-dir=/usr/local/apache2/htdocs/orca
```

this configures everything!  
then the usual

```
# make && make install
```

and you're done

```
# echo "/usr/local/bin/orca /usr/local/lib/orcallator.cfg &" > /etc/rc.d/rc.orca
# chmod 755 /etc/rc.d/rc.orca
# echo "/etc/rc.d/rc.orca" >> /etc/rc.d/rc.M
# /etc/rc.d/rc.orca
```

then created a tarball for the solaris clients, that includes de SE-Toolkit, from

<http://www.setoolkit.com/setoolkit.html>

you can get this tarball I made clicking [here](#)

Copy it to your solaris machine into /tmp , gunzip it and untar it, then:

```
# ./installorca.ksh
```

It will bail out if you don't have gawk, which you can get from [www.sunfreeware.com](http://www.sunfreeware.com).

If you don't have the SE-Toolkit installed, it will install it, and it will also install the appropriate rc script in /etc/rc3.d/S99orcallator

to start it:

```
# /etc/rc3.d/S99orcallator start
```

it will create files in /usr/local/var/orca/orcallator

So for each Solaris machine, I'm getting all created files in this directory with scp (with RSA authentication) in a crontab line like this inside the linux slackware machine:

```
0,5,10,15,20,25,30,35,40,45,50,55 /usr/bin/run-parts /etc/cron.fivemin 1> /dev/null
```

which calls a script inside the /etc/cron.fivemin, each five minutes, called getorca.ksh:

```
#!/bin/ksh
for host in `cat /etc/cron.fivemin/hosts.txt`;do
echo "----- $host"
mkdir -p /usr/local/var/orca/orcallator/$host > /dev/null 2>&1
cd /usr/local/var/orca/orcallator/$host
scp $host:/usr/local/var/orca/orcallator/$host/* .
done
```

This copies all data from the solaris hosts to the linux machine, which in turn processes all data into the rrdtool engine, and generates pages in /usr/local/apache2/htdocs/orca

done!

Posted by rdircio in Solaris at 20:05

Thursday, September 1. 2005

## My first zone with solaris10

I decided to create a zone in one of my sf25k domains, using 50% of the cpu/mem on the machine, these are the commands I used:

```
bash-3.00# zonecfg -z zone1
zone1: No such zone configured
Use 'create' to begin configuring a new zone.
zonecfg:zone1> create
zonecfg:zone1> set zonepath=/zone1
zonecfg:zone1> set address=150.250.250.21
Cannot set a resource-specific property from the global scope.
zonecfg:zone1> add net
zonecfg:zone1:net> set address=150.250.250.21
zonecfg:zone1:net> set physical=ce1
zonecfg:zone1:net> end
zonecfg:zone1> set autoboot=true
zonecfg:zone1> commit
zonecfg:zone1> exit
bash-3.00# chmod 700 /zone1
bash-3.00# zoneadm -z zone1 install
bash-3.00# zoneadm -z zone1 boot
bash-3.00# zoneadm list -cv
ID NAME STATUS PATH
0 global running /
1 zone1 running /zone1
bash-3.00# zonecfg -z zone1
zonecfg:zone1> add rctl
zonecfg:zone1:rctl> set name=zone.cpu-shares
zonecfg:zone1:rctl> set value=(priv=privileged,limit=50,action=none)
zonecfg:zone1:rctl> end
zonecfg:zone1> verify
zonecfg:zone1> commit
zonecfg:zone1> exit
# zlogin -C zone1
```

The first time you connect to it, the zone will ask you for its initial settings, after this is done, it reboots and is ready to work.

Posted by rdircio in Solaris at 11:19

Tuesday, August 30, 2005

## Script to delete old files in directories

I needed to delete 30 days old files from some log directories automatically, in a crontab entry.

This is a ksh shell to do it, it takes a list of directories from "dirs.txt", and for each directory deletes 30 day old files.

```
#!/bin/ksh
test=0 #---- Is this a test?
age="+30" #---- How old do we want the files to be deleted?
dir_list=`cat /scripts/depura/dirs.txt`
for dir in $dir_list;
do
echo "Checking: $dir"
cd $dir
if [ $test -eq 1 ];then
/usr/bin/find . -type f -mtime $age -exec rm {} \;
else
/usr/bin/find . -type f -mtime $age -exec ls -ld {} \;
fi
done
```

This deletes even files in subdirectories inside the chosen dirs, so if you want to delete only the files in the directories, not in subdirs, substitute the find commands with a -prune option:

```
#!/bin/ksh
test=0 #---- Is this a test?
age="+30" #---- How old do we want the files to be deleted?
dir_list=`cat /scripts/depura/dirs.txt`
for dir in $dir_list;
do
echo "Checking: $dir"
cd $dir
if [ $test -eq 1 ];then
/usr/bin/find * -prune -type f -mtime $age -exec rm {} \;
else
/usr/bin/find * -prune -type f -mtime $age -exec ls -ld {} \;
fi
done
```

Posted by rdircio in Solaris at 15:46

Sunday, August 28, 2005

## BMC controlm semi-automated install on Solaris9

We've decided to install both controlm and em on the same machine. We'll have to install it on several machines, because we'd like having a remote failover on two cities +700 miles from each. All of this complicates our existence.

We've also been told to separate the database repository on a different disk/partition.

Given the long wishlist, a custom setup was necessary, and several options were to be chosen in the installation process.

We wish in this unix world that everything should be files, but the folks at BMC made a nice little install shell that is INTERACTIVE, so it gives less room to automation....

So we proceeded to build our 4 filesystems, 2 for controlm and 2 for em.

/em and /ora02 for em

/ctrlm and /ora01

for controlm

after this we built a script to appropriately build these users and place their appropriate permissions:

makeusers1.ksh

```
#!/bin/ksh
```

```
set -x
```

```
rm
```

```
-fR /ora02/oradata
```

```
rm -fR /ora01/oradata
```

```
mkdir -p
```

```
/export/home
```

```
groupadd dba
```

```
userdel -r ctm613
```

```
userdel -r
```

```
em613
```

```
userdel -r agt613
```

```
useradd -s /bin/csh -g dba -d
```

```
/ctrlm/ctm613 -m ctm613
```

```
useradd -s /bin/csh -g dba -d /em/em613 -m
```

```
em613
```

```
useradd -s /bin/csh -g dba -d /export/home/agt613 -m
```

```
agt613
```

```
mkdir -p /ora02/oradata/EM613
```

```
mkdir -p
```

```
/ora01/oradata/ctrlm
```

```
chown -fR em613:dba /em/em613
```

```
chown -fR
```

```
ctm613:dba /ctrlm/ctm613
chown -fR em613:dba /ora02
chown -fR
ctm613:dba /ora01
No way we can have problems repeating
this task...
We've tuned the kernel properly in /etc/system:
```

```
set msgsys:msginfo_msgmap=500
set
msgsys:msginfo_msgmnb=65536
set msgsys:msginfo_msgmni=1024
set
msgsys:msginfo_msgseg=8192
set msgsys:msginfo_msgssz=16
set
msgsys:msginfo_msgtql=1024
set semsys:seminfo_semmap=384
set
semsys:seminfo_semmni=2048
set semsys:seminfo_semmns=2048
set
semsys:seminfo_semmnu=1024
set semsys:seminfo_semmsl=2048
set
semsys:seminfo_semopm=128
set semsys:seminfo_semume=1024
set
semsys:seminfo_semvmx=65534
set
shmsys:shminfo_shmmax=4294967295
set shmsys:shminfo_shmmin=1
set
shmsys:shminfo_shmmni=600
set shmsys:shminfo_shmseg=1024
```

Remember to do a "init 6" after this kernel tuning. After this we proceeded to make use of standard redirection to answer all our setup questions(it was a bit overkill to use "expect")!

Suppose you've copied all the Oracle, controlm and em CDs to /export/home/control-m  
As root,

```
# cd /em/em613
#
```

```
/export/home/control-m/fd5380_ho7/orapreinst.sh <
/scripts/controlm/1/orapreinst-em.txt
```

Here's where the fun starts!

orapreinst-em

contains:

2

s

which basically are the answers to the questions the  
"orapreinst.sh" asks, with the line breaks it needs...

hehe

After this has finished. log in as em613, with a graphical  
connection, so u're able to open a terminal in em613's session. Run  
the em oracle installation:

```
$/export/home/control-m/fd5380_ho7/inst_oracle <
/scripts/controlm/1/instoracle-em.txt
```

instoracle-em.txt

contents:

2

L

3

adminpass04

adminpass04

4

/export/home/control-m/TEMP/ora\_inst\_dir/Disk1

n

9

/ora02/oradata/EM613/EM613indx01.dbf

11

/ora02/oradata/EM613/log01.dbf

12

/ora02/oradata/EM613/log02.dbf

n

15

/ora02/oradata/EM613/system01.dbf

17

/ora02/oradata/EM613/rbs01.dbf

n

19

*/ora02/oradata/EM613/temp01.dbf*

*b*

*Well,*

*you get it, all of the entries in the files have the answers for the questions on the installer script...*

*While running this, find a*

*root.sh shell in /em/em613 and run it as root before the first installation gets to 10%*

*When all of the instoracle procedure*

*ends, go to /em/em613 as root and run:*

```
# /export/home/control-m/fd5380_ho7/orapostinst.sh
```

login to em613 and run the components installer,  
again with an input file:

```
$ /export/home/control-m/TEMP/em613_sparc5309_3/INSTALL/LOAD <  
/scripts/controlm/1/emcomponents.txt
```

After this, you're done with em613!!!!

Now we'll

do the same with controlm....

Login as root, go to /ctrlm/ctm613,

and run:

```
# /export/home/control-m/fd5380_ho7/orapreinst.sh <  
/scripts/controlm/1/orapreinst-ctrlm.txt
```

Then login as ctm613 in a graphical X session, and  
issue the command:

```
$ /export/home/control-m/fd5380_ho7/inst_oracle <  
/scripts/controlm/1/instoracle-ctrlm.txt
```

before the first stage gets to 10%, find root.sh  
inside /ctrlm/ctm613 and run it. As root run:

```
# /export/home/control-m/fd5380_ho7/orapostinst.sh
```

As ctm613 run:

```
$ /export/home/control-m/controlm-server/CD_INST <  
/scripts/controlm/1/ctrlmcomponents.txt  
$ install/custom <  
/scripts/controlm/1/installcustom.txt
```

You can download these "answer files" as [scripts-controlm.tar.gz](http://scripts-controlm.tar.gz)

That's it!

Posted by rdircio in Solaris at 19:14

Wednesday, August 17. 2005

## Veritas Foundation suite 4.1 Installation in Solaris 10

To remember which options to choose, this is the log of all the interaction I had with the veritas foundation suite installation script:

```
bash-3.00# ./installsf
```

VERITAS STORAGE FOUNDATION 4.1 INSTALLATION PROGRAM

Copyright (c) 2005 VERITAS Software Corporation. All rights reserved.

VERITAS, the VERITAS Logo and all other VERITAS product names and slogans are trademarks or registered trademarks of VERITAS Software Corporation. VERITAS and the VERITAS Logo Reg. U.S. Pat. & Tm. Off. Other product names and/or slogans mentioned herein may be trademarks or registered trademarks of their respective companies.

Enter the system names separated by spaces on which to install SF: distmty2-g

```
Checking OS version on distmty2-g ..... SunOS 5.10
Verifying global zone on distmty2-g ..... global
```

Initial system check completed successfully.

Press [Return] to continue:

VERITAS STORAGE FOUNDATION 4.1 INSTALLATION PROGRAM

VERITAS Infrastructure package installation:

Installing VERITAS Infrastructure packages on distmty2-g:

```
Checking VRTScpi package ..... not installed
Checking VRTSvlic package ..... not installed
Checking file system space ..... required space is available
Installing VRTScpi 4.1.0.92 on distmty2-g ..... Done
Installing VRTSvlic 3.02.005k on distmty2-g ..... Done
```

VERITAS Infrastructure packages installed successfully.

VERITAS STORAGE FOUNDATION 4.1 INSTALLATION PROGRAM

SF Licensing Verification:

```
Checking SF license key on distmty2-g ..... not licensed
Enter a SF license key for distmty2-g: [?] 38PU-VIRB-RNBU-T5TF-KN3B-IGTP-NC3
Registering VERITAS Storage Foundation Standard PERMANENT_NODE_LOCK key on distmty2-g
```

Do you want to enter another license key for distmty2-g? [y,n,q,?] (n)  
SF licensing completed successfully.

Press [Return] to continue:

## VERITAS STORAGE FOUNDATION 4.1 INSTALLATION PROGRAM

installsf can install the following optional SF packages:

VRTSobgui VERITAS Enterprise Administrator  
VRTSvmman VERITAS Volume Manager Manual Pages  
VRTSvmdoc VERITAS Volume Manager Documentation  
VRTStep VERITAS Task Provider  
VRTSap VERITAS Action Provider  
VRTSfsman VERITAS File System Manual Pages  
VRTSfsdoc VERITAS File System Documentation  
VRTSfsmnd VERITAS File System Software Developer Kit Manual Pages  
VRTSfasdc VERITAS FlashSnap Agent for Symmetrix Documentation  
VRTSvrw VERITAS Volume Replicator Web Console  
VRTSvrdoc VERITAS Volume Replicator Documentation

- 1) Install all of the optional packages
- 2) Install none of the optional packages
- 3) View package descriptions and select optional packages

Select the optional packages to be installed on all systems? [1-3,q,?] (1)

## VERITAS STORAGE FOUNDATION 4.1 INSTALLATION PROGRAM

installsf will install the following SF packages:

VRTSperl VERITAS Perl 5.8.0 Redistribution  
VRTSob VERITAS Enterprise Administrator Service  
VRTSobgui VERITAS Enterprise Administrator  
VRTSvxvm VERITAS Volume Manager Binaries  
VRTSvmman VERITAS Volume Manager Manual Pages  
VRTSvmdoc VERITAS Volume Manager Documentation  
VRTSvmpro VERITAS Volume Manager Management Services Provider  
VRTSfspro VERITAS File System Management Services Provider  
VRTSalloc VERITAS Volume Manager Intelligent Storage Provisioning  
VRTSddlpr VERITAS Device Discovery Layer Services Provider  
VRTSvxfs VERITAS File System  
VRTStep VERITAS Task Provider  
VRTSap VERITAS Action Provider  
VRTSfsman VERITAS File System Manual Pages  
VRTSfsdoc VERITAS File System Documentation  
VRTSfssdk VERITAS File System Software Developer Kit  
VRTSfsmnd VERITAS File System Software Developer Kit Manual Pages  
VRTSvail VERITAS Array Provider  
VRTSfas VERITAS FlashSnap Agent for Symmetrix

Press [Return] to continue:

..continued:

VRTSfasdc VERITAS FlashSnap Agent for Symmetrix Documentation  
VRTSvrpro VERITAS Volume Replicator Client Extension and Provider for  
VERITAS Enterprise Administrator  
VRTSvcsvr VERITAS Cluster Server Agents for VVR  
VRTSjre VERITAS Java Runtime Environment Redistribution  
VRTSweb VERITAS Java Web Server

VRTSvrw VERITAS Volume Replicator Web Console  
VRTSvrdoc VERITAS Volume Replicator Documentation  
Press [Return] to continue:  
VERITAS STORAGE FOUNDATION 4.1 INSTALLATION PROGRAM

Checking system installation requirements:

Checking SF installation requirements on distmty2-g:

Checking VRTSperl package ..... not installed  
Checking VRTSob package ..... not installed  
Checking VRTSobgui package ..... not installed  
Checking VRTSvxvm package ..... not installed  
Checking VRTSvmman package ..... not installed  
Checking VRTSvmdoc package ..... not installed  
Checking VRTSvmpro package ..... not installed  
Checking VRTSfspro package ..... not installed  
Checking VRTSalloc package ..... not installed  
Checking VRTSddlpr package ..... not installed  
Checking VRTSvxfs package ..... not installed  
Checking VRTStep package ..... not installed  
Checking VRTSap package ..... not installed  
Checking VRTSfsman package ..... not installed  
Checking VRTSfsdoc package ..... not installed  
Checking VRTSfssdk package ..... not installed  
Checking VRTSfsmnd package ..... not installed  
Checking VRTSvail package ..... not installed  
Checking VRTSfas package ..... not installed  
Checking SYMCLI required for VRTSfas ..... not installed  
Checking VRTSfasdc package ..... not installed  
Checking VRTSvrpro package ..... not installed  
Checking VRTSvcsvr package ..... not installed  
Checking VRTSjre package ..... not installed  
Checking VRTSweb package ..... not installed  
Checking VRTSvrw package ..... not installed  
Checking VRTSvrdoc package ..... not installed  
Checking VERITAS patch 115209 ..... not installed  
Checking VERITAS patch 115210 ..... not installed  
Checking for any Solaris patch issues ..... None  
Checking file system space ..... required space is available  
Checking for any AP driver issues ..... None  
Checking vxsvc process ..... not running  
Checking VRTSweb process ..... not running  
Checking vxportal driver ..... not running  
Checking fdd driver ..... not running  
Checking qllog driver ..... not running  
Checking vxfs driver ..... not running

Installation requirement checks completed successfully.

Press [Return] to continue:  
VERITAS STORAGE FOUNDATION 4.1 INSTALLATION PROGRAM

It is possible to install SF packages without performing configuration.

It is optional to configure SF now. If you choose to configure SF later, you can either do so manually or run the `installsf -configure` command.

Are you ready to configure SF? [y,n,q] (y) y

## VERITAS STORAGE FOUNDATION 4.1 INSTALLATION PROGRAM

Installing Storage Foundation Standard 4.1 on distmty2-g:

Installing VRTSperl 4.0.12 on distmty2-g ..... Done 1 of 26 steps  
Installing VRTSob 3.2.514.0 on distmty2-g ..... Done 2 of 26 steps  
Adding patch 115209-16 on distmty2-g ..... Done 3 of 26 steps  
Installing VRTSobgui 3.2.514.0 on distmty2-g ..... Done 4 of 26 steps  
Adding patch 115210-16 on distmty2-g ..... Done 5 of 26 steps  
Installing VRTSvxvm 4.1 on distmty2-g ..... Done 6 of 26 steps  
Installing VRTSvmman 4.1 on distmty2-g ..... Done 7 of 26 steps  
Installing VRTSvmdoc 4.1 on distmty2-g ..... Done 8 of 26 steps  
Installing VRTSvmpro 4.1 on distmty2-g ..... Done 9 of 26 steps  
Installing VRTSfspro 4.1 on distmty2-g ..... Done 10 of 26 steps  
Installing VRTSalloc 4.1 on distmty2-g ..... Done 11 of 26 steps  
Installing VRTSddlpr 4.1 on distmty2-g ..... Done 12 of 26 steps  
Installing VRTSvxfs 4.1 on distmty2-g ..... Done 13 of 26 steps  
Installing VRTStep 1.20.028 on distmty2-g ..... Done 14 of 26 steps  
Installing VRTSap 2.00.023 on distmty2-g ..... Done 15 of 26 steps  
Installing VRTSfsman 4.1 on distmty2-g ..... Done 16 of 26 steps  
Installing VRTSfsdoc 4.1 on distmty2-g ..... Done 17 of 26 steps  
Installing VRTSfssdk 4.1 on distmty2-g ..... Done 18 of 26 steps  
Installing VRTSfsmnd 4.1 on distmty2-g ..... Done 19 of 26 steps  
Installing VRTSvail 4.2.45 on distmty2-g ..... Done 20 of 26 steps  
Installing VRTSvrpro 4.1 on distmty2-g ..... Done 21 of 26 steps  
Installing VRTSvcsvr 4.1 on distmty2-g ..... Done 22 of 26 steps  
Installing VRTSjre 1.4 on distmty2-g ..... Done 23 of 26 steps  
Installing VRTSweb 4.2 on distmty2-g ..... Done 24 of 26 steps  
Installing VRTSvrw 4.1 on distmty2-g ..... Done 25 of 26 steps  
Installing VRTSvrdoc 4.1 on distmty2-g ..... Done 26 of 26 steps

Storage Foundation Standard installation completed successfully.

Press [Return] to continue:

The enclosure-based naming scheme is a feature of Volume Manager. It allows one to reference disks using a symbolic name that is more meaningful than the operating system's normal device access name. This symbolic name is typically derived from the array name.

Evaluating which systems can be configured for the enclosure-based naming scheme...

System distmty2-g is eligible -- can configure.

Do you want to set up the enclosure-based naming scheme? [y,n,q,?] (n) n

## VERITAS STORAGE FOUNDATION 4.1 INSTALLATION PROGRAM

Configuring Storage Foundation Standard:

Configuring Storage Foundation Standard on distmty2-g ..... Done

Storage Foundation Standard configured successfully.

Press [Return] to continue:

VERITAS STORAGE FOUNDATION 4.1 INSTALLATION PROGRAM

Do you want to start Storage Foundation Standard processes now? [y,n,q] (y) y

Note: The vxconfigd daemon will be started, which can take a while depending upon the hardware configuration.

Disabling enclosure-based naming on distmty2-g ..... Done

Starting vxconfigd for VxVM on distmty2-g ..... Started

Volume Manager default disk group configuration:

Many Volume Manager commands affect the contents or configuration of a disk group. Such commands require that the user specify a disk group. This is accomplished by using the -g option of a command or setting the VXVM\_DEFAULTDG environment variable. An alternative to these two methods is to configure the default disk group of a system.

Do you want to set up the default disk group for each system? [y,n,q,?] (y) n

Volume Manager default disk group setup and daemon startup

You declined to set up the default disk group for distmty2-g.

Starting vxcached on distmty2-g ..... Started

Storage Foundation Standard was started successfully.

Press [Return] to continue:

VERITAS STORAGE FOUNDATION 4.1 INSTALLATION PROGRAM

Installation of Storage Foundation Standard 4.1 has completed successfully.

The installation summary is saved at:

/opt/VRTS/install/logs/installsf817194957.summary

The installsf log is saved at:

/opt/VRTS/install/logs/installsf817194957.log

The installation response file is saved at:

/opt/VRTS/install/logs/installsf817194957.response

Reboot all systems on which VxFS was installed or upgraded.

shutdown -y -i6 -g0

See the VERITAS File System Administrators Guide for information on using VxFS.

Posted by rdircio in Solaris at 20:10

Friday, August 12. 2005

### **How to create a file in solaris**

You can create a big file to test, for example if your filesystem supports large files, or to transfer it to other host and diagnose network speed, or for anything you want...

```
# mkfile 6g some.file
```

creates a 6 gigabyte file

Posted by rdircio in Solaris at 13:29

Tuesday, August 9, 2005

## getting the wwpn from a qllogic card in a SF25k

just write

```
# prtctl -v -c scsi-fcp
```

or

```
# prtconf -vp | grep -i wwn
```

Posted by rdircio in Solaris at 13:23

Wednesday, July 27. 2005

## sshd connections dropping in Solaris 9, after applying 9\_recommended patches

Yes, these guys at sun sometimes cause problems rather than fixing them.

If you use their stock sshd, with X11Forwarding enabled, sshd will kick your login after applying their recommended patches

To get around this edit /etc/rc3.d/S89sshd and change

```
[ -x /usr/lib/ssh/sshd ] && /usr/lib/ssh/sshd &
```

to

```
[ -x /usr/lib/ssh/sshd ] && /usr/lib/ssh/sshd -4 &
```

After doing this edit the /etc/ssh/sshd\_config and comment the line

```
#ListenAddress ::
```

so it doesn't use IPv6

Just restart it

```
#!/etc/rc3.d/S89sshd stop
```

```
#!/etc/rc3.d/S89sshd start
```

After this you'll have a working ssh with X11Forwarding

Good Luck.

Posted by rdircio in Solaris at 22:21

Thursday, July 14. 2005

## Configuring IPMP in Solaris 9

We want to balance an ip address within two network cards, using two cables. So we need 3 IPS, one for each cable, and the other to be flying around on both.

To exemplify, we have these 3 addresses:

```
100.100.100.1
100.100.100.2
100.100.100.3
```

We have to assign two of them to physical interfaces (100.100.100.1=ce0, 100.100.100.2=ce1), and the other(100.100.100.3) should be either ce0:1 or ce1:1, depending on the availability of the physical interface.

We have to change the local-mac-address param in the eeprom to "true", it is a requirement to have different MAC addresses in all cards for IPMP to work. First we ask its value:

```
# eeprom local-mac-address?
```

```
local-mac-address?=false
```

Then, if false, we change it:

```
# eeprom local-mac-address?=true
```

We need a reboot to see the effects, but we then continue to configure other IPMP necessary files and then reboot.

So, we edit /etc/hostname.ce0 and put this, creating group "migroup":

```
100.100.100.1 deprecated -failover netmask + broadcast + group migroup up addif 100.100.100.3 netmask + broadcast + up
```

In /etc/hostname.ce1 we put this:

```
100.100.100.2 deprecated -failover netmask + broadcast + group migroup up
```

After this we can reboot and then we have this:

```
ce0 100.100.100.1
ce1 100.100.100.2
ce0:1 100.100.100.3
```

In case ce0 fails we will have:

```
ce1 100.100.100.2
ce1:1 100.100.100.3
```

Optionally, and for your control, add each ip to /etc/hosts, associated with a name.

Ok?

Posted by rdircio in Solaris at 11:02

Wednesday, July 13. 2005

## **Mount ISO images on Solaris**

You can mount ISO images in solaris without burning them to a CD.

You can use a one liner to do this:

```
# mount -F hsfs -o ro `lofiadm -a /files/disc.iso` /mountpoint
```

To unmount it do this:

```
# umount /mountpoint; lofiadm -d /files/disc.iso
```

Posted by rdircio in Solaris at 12:59

Tuesday, July 12. 2005

### **openssh sshd log levels on solaris 8**

Sometimes you can get your syslog facility flooded with ssh connection messages, for instance, I have an automatic performance meter that uses ssh to connect. This keeps appending messages to /var/adm/messages like this:

```
sshd: public key accepted for 150.100.106.32 for root
```

And I hate it!

So you can edit your /usr/local/etc/sshd\_config, and change the "LogLevel" directive to "Quiet", in my case, or to some value in this table, according to the level of verbosity you want:

```
QUIET<-> 0
```

```
FATAL <-> 1
```

```
ERROR <-> 2
```

```
INFO <-> 3
```

```
VERBOSE <-> 4
```

```
DEBUG1 <-> 5
```

```
DEBUG2 <-> 6
```

```
DEBUG3 <-> 7
```

Posted by rdircio in Solaris at 22:46

Wednesday, April 27. 2005

## MT driver for tivoli tapes in Solaris 9, using a FC64-1063 sbus card in an E10K

While installing tivoli storage manager's (TSM) scsi tape driver for solaris, we had this message

```
# add_drv -m '* 0666 bin bin' /usr/kernel/drv/sparcv9/mt
driver (mt) successfully added to system but failed to attach
```

It loaded the driver, but devices in /dev/rmt never appeared. The mt driver is supposed to create tape entries in /dev/rmt to use them as tsm tape devices.

Googling into the error message, we found this document for another jni card model and driver: [http://www.jni.com/drivers/Files/Solaris/Solaris\\_5311/readme\\_Sol\\_5311\\_v2.txt](http://www.jni.com/drivers/Files/Solaris/Solaris_5311/readme_Sol_5311_v2.txt)

This document had this hint

5. Configure ACSLS with the "install\_scsi\_sol.sh" script. When running this you will expect the error: "devfsadm: driver failed to attach: mchanger" "Warning: Driver (mchanger) successfully added

to system but failed to attach" This process failed, but it is the process that creates the

/usr/kernel/drv/mchanger.conf file...

6. Edit /usr/kernel/drv/mchanger.conf to add: name="mchanger" parent="jnic146x" target=0;

Translating it to our machine and driver problem... mchanger=mt, and jnic146x=fcaw, because we are trying to use the mt driver, and attach it to the "fcaw"

driver of our specific fc64-1063 card. So we proceed to edit /usr/kernel/drv/mt.conf, and add a parent="fcaw" section to each line identifying a device.

```
# vi /usr/kernel/drv/mt.conf
name="mt" parent="fcaw" class="fp" target=0 lun=0;
name="mt" parent="fcaw" class="fp" target=1 lun=0;
name="mt" parent="fcaw" class="fp" target=2 lun=0;
```

Now we unload the unattached driver.

```
# rem_drv mt
```

We load it again, and it reads its /usr/kernel/drv/mt.conf file, attaching the driver to its parent, fcaw.

```
# add_drv -m '* 0666 bin bin' /usr/kernel/drv/sparcv9/mt
```

Now we see all those mt devices in /dev/rmt

```
# ls /dev/rmt/*mt*
/dev/rmt/3mt /dev/rmt/3mtt /dev/rmt/4mtn /dev/rmt/5mt /dev/rmt/5mtt /dev/rmt/3mtn
/dev/rmt/4mt /dev/rmt/4mtt /dev/rmt/5mtn
```

nice!

Posted by rdircio in Solaris at 16:21

Tuesday, April 12. 2005

## How to extract files from a flar archive

1. Split the Flash Archive into the cookie, identification, and archive files.

```
# flar split flash_archive.flar
```

2. Extract the file(s) from the cpio archive.

```
# cpio -idv var/adm/messages < archive
```

Posted by rdircio in Solaris at 16:41

Saturday, April 9, 2005

## How to check for/enable largefile support on vxfs

To check if largefiles are enabled:

```
# /usr/lib/fs/vxfs/fsadm /mount/point  
nolargefiles
```

Unlike ufs, you can enable vxfs largefile support on the fly with fsadm:

```
# /usr/lib/fs/vxfs/fsadm -o largefiles /mount/point
```

Posted by rdircio in Solaris at 01:00

Friday, April 1. 2005

## Verify daylight savings time zone changes in Solaris 8,9

*If you doubt on how your time will change, or if it will change on your servers according to your timezone, take a look at this command:*

```
# /usr/sbin/zdump -v $TZ
```

```
Mexico/General Sun Apr 3 07:59:59 2005 UTC = Sun Apr 3 01:59:59 2005 CST isdst=0
```

```
Mexico/General Sun Apr 3 08:00:00 2005 UTC = Sun Apr 3 03:00:00 2005 CDT isdst=1
```

```
Mexico/General Sun Oct 30 06:59:59 2005 UTC = Sun Oct 30 01:59:59 2005 CDT isdst=1
```

```
Mexico/General Sun Oct 30 07:00:00 2005 UTC = Sun Oct 30 01:00:00 2005 CST isdst=0
```

*The first time indicates GMT time, the second is local time, you can see on April 3, the machine will change one hour forward at 2:00 am, it will set the clock to 3:00 am.*

Posted by rdircio in Solaris at 11:07

Thursday, March 31. 2005

## Uninstalling/Reinstalling forte developer 7 on Solaris Sparc

*Got a demo license problem on forte developer 7, we're using the trial version!. To get a new trial license you must uninstall and install everything, first, to uninstall:*

```
# cd /var/sadm/prod/com.sun.forte_developer_7
```

```
# ls
```

```
uninstall_Forte_Developer_7_Compilers_SPARC.class uninstall_Forte_Developer_7_Tools_SPARC.class
```

```
uninstall_Forte_Developer_7_Documentation_SPARC.class uninstall_Solaris_patches_for_Forte_Developer_7.class
```

To uninstall do:

```
# /usr/bin/java uninstall_Forte_Developer_7_Compilers_SPARC -nodisplay -noconsole
```

```
# /usr/bin/java uninstall_Forte_Developer_7_Documentation_SPARC -nodisplay -noconsole
```

```
# /usr/bin/java uninstall_Forte_Developer_7_Tools_SPARC -nodisplay -noconsole
```

```
# /usr/bin/java uninstall_Solaris_patches_for_Forte_Developer_7 -nodisplay -noconsole
```

Then reinstall, in the installation directory, like:

```
# ./batch_installer -t
```

To check your license:

```
# /opt/SUNWspro/bin/cc -xlicinfo
```

Number of serial numbers found: 1

Total RTUs available: 1

License file used:

```
/opt/SUNWspro/bin/./prod/bin/./lib/serial.dat
```

License tokens accepted by this product: FDT070

Serial Number (SN): FDT070-M99BDISNU-031042412

Is SN valid?: yes

RTUs: 1


Is trial SN?: yes

Days until expiration: 60

Expiration date: 05/31/05

Posted by rdircio in Solaris at 14:47

## Create new instance of jfserver in Adobe Central Pro Output Server 5.5



I needed to create a new instance of jfserver, in order to make a difference between jobs of two applications, so that forms, incoming and outgoing jobs fall into different directories.

This is how i did it.

My ACPOS5.5 installation is in /usr/local/adobe/central, inside of it i have an instance of jfserver, inside the /usr/local/adobe/central/server directory.

I proceeded to build a directory schema with only the necessary files in /usr/local/adobe/central/mx:

```
bash-2.03# find mx
```

```
mx
mx/control
mx/data
mx/backup
mx/error
mx/etc
mx/etc/exprint
mx/etc/exprint/pdf
mx/etc/exprint/forms
mx/jfserver.log
mx/jfserver.jmd
mx/jfserver.ini
mx/jfserver.sts
```

Jfserver.jmd is a copy of /usr/local/adobe/central/server/jfserver.jmd

Jfserver.ini is a copy of /usr/local/adobe/central/server/jfserver.ini, but edited to match the "mx" directory.

I'll make a special mention to these lines:

### [Paths]

```
ExePath=/usr/local/adobe/central/bin
MSTPath=/usr/local/adobe/central/mx/mst
ControlPath=/usr/local/adobe/central/mx/control
DataPath=/usr/local/adobe/central/mx/data
BackupPath=/usr/local/adobe/central/mx/backup
ErrorPath=/usr/local/adobe/central/mx/error
EtcPath=/usr/local/adobe/central/mx/etc
```

### [ReturnParms]

```
NumberPages=0
SpoolId=
OutputFileName=
User1=/usr/local/adobe/central/mx/etc/exprint/pdf/
User2=
```

After this, just edited /usr/local/adobe/central/install.ini, and added an entry for the new server:

*[Configuration]*

```
server@afpmxc=/usr/local/adobe/central/server/jfserver.ini  
server2@afpmxc=/usr/local/adobe/central/mx/jfserver.ini
```

*[Paths]*

```
server=/usr/local/adobe/central/server  
server2=/usr/local/adobe/central/mx
```

*[JetForm]*

```
Node=afpmxc
```

This adjusts ACPOS to launch our server issuing the command:

```
# /usr/local/adobe/central/bin/jfdaemon
```


You'll get something like this:

```
bash-2.03# !ps  
ps -ef | grep jf  
root 10361 1 0 13:47:01 ? 0:00 /usr/local/adobe/central/bin/jfdaemon  
root 10363 10361 0 13:47:01 ? 0:00 jfserver -aai/usr/local/adobe/central/mx/jfserver.ini  
root 10362 10361 0 13:47:01 ? 0:00 jfserver -aai/usr/local/adobe/central/server/jfserver.ini
```

Posted by rdircio in Solaris at 13:49

Wednesday, March 9. 2005

## flarcreate and multiple excludes in solaris 8



Jumpstart flash is very useful to create an "image" of your OS, it's like a perfect copy of your installation, and it installs really easily.

In solaris 8, however, we have a limitation that ends up the party: you can't exclude multiple directories!!!! This means that if you want to backup only the OS, but not the mounted filesystems of your T3, you're out at the cold.

The command:

```
# flarcreate -n "afpmxc" -a "rdircio@gmail.com" -R / -x /var/tmp -x /nacar -x /software -x /Netscape -x /WebSphere -x /qa -x /waslogs -x /DB2 -x /software2 -x /patrol -x /opt/mqm -x /opt/admini3 -x /export/home -x /var/mqm -x /var/mqm/log -x /opt/IBM/db2/V8.1 -x /fo -x /was40 -x /chgmanQA -x /chgmands /var/tmp/afpmxc.flar
```

is not accepted, and only excludes the first entry, in this case /var/tmp.

After googling a bit, and finding very [creative solutions](#), including a use of find & grep & awk, I found out that this limitation is fixed in a solaris 8 patch: 109318, and a Solaris 9 patch: 113434 (yes, solaris 9 has also limitations in this field)

This patch surely comes in the "recommended" patches bundle.

Posted by rdircio in Solaris at 11:28

## Speed of qfe,eri,hme and others, in solaris 8,9

Just use this script, uncommenting the appropriate lines.....  
Except for ce, for which you should use the ce.conf file, explained in my previous blog entry. Read the comment inside it for installation instructions!

```
#!/sbin/sh
# /etc/rc2.d/S68net-tune /etc/init.d/net-tune
# Copyright (c) 2000 by Sun Microsystems, Inc.
# All rights reserved.
#
#ident "@(#)net-tune 1.7.3 031106 SMI ES-PTS-NET/Lor"
#
#####
# This script is not needed for most installations ! #
# #
# It MAY be needed for rare special customizations. In this case, #
# copy the script and edit the following example code to suit your needs. #
#####
#
# Depending on your changes, this script overwrites the Sun[TM] recommended
# default values (To use the default values, you do not need this script).
# In particular, Sun's[TM] recommendation is to leave Ethernet Auto-negotiation
# ON at both link partners (this is the default). See Product Documentation and
# SunSolve Infodocs for proper use first (e.g. Infodocs 17416, 23041, 41665,
# 70282). Most parameters must be set on both sides to have effect
# (e.g check also switch and IP partner settings).
# To install:
# 1) cp S68net-tune to /etc/rc2.d
# 2) perform edits on the script as required
# 3) chmod 744 /etc/rc2.d/S68net-tune
# 4) chown root:sys /etc/rc2.d/S68net-tune
# 5) ln /etc/rc2.d/S68net-tune /etc/init.d/net-tune

PATH=/usr/bin:/usr/sbin

case "$1" in
'start')
echo "Implementing Solaris Network Tuning."

# hme-Interfaces
# hme0
#ndd -set /dev/hme instance 0
#ndd -set /dev/hme adv_100T4_cap 0
#ndd -set /dev/hme adv_100fdx_cap 1
#ndd -set /dev/hme adv_100hdx_cap 0
#ndd -set /dev/hme adv_10fdx_cap 0
```

```
#nnd -set /dev/hme adv_10hdx_cap 0
#nnd -set /dev/hme adv_autoneg_cap 0
# hme1
#nnd -set /dev/hme instance 1
#nnd -set /dev/hme adv_100T4_cap 0
#nnd -set /dev/hme adv_100fdx_cap 1
#nnd -set /dev/hme adv_100hdx_cap 0
#nnd -set /dev/hme adv_10fdx_cap 0
#nnd -set /dev/hme adv_10hdx_cap 0
#nnd -set /dev/hme adv_autoneg_cap 0
```

#### # eri-Interfaces

```
# eri0
#nnd -set /dev/eri instance 0
#nnd -set /dev/eri adv_100T4_cap 0
#nnd -set /dev/eri adv_100fdx_cap 1
#nnd -set /dev/eri adv_100hdx_cap 0
#nnd -set /dev/eri adv_10fdx_cap 0
#nnd -set /dev/eri adv_10hdx_cap 0
#nnd -set /dev/eri adv_autoneg_cap 0
```

#### # qfe-Interfaces

```
# qfe0
nnd -set /dev/qfe instance 0
nnd -set /dev/qfe adv_100T4_cap 0
nnd -set /dev/qfe adv_100fdx_cap 1
nnd -set /dev/qfe adv_100hdx_cap 0
nnd -set /dev/qfe adv_10fdx_cap 0
nnd -set /dev/qfe adv_10hdx_cap 0
nnd -set /dev/qfe adv_autoneg_cap 0
# qfe3
nnd -set /dev/qfe instance 3
nnd -set /dev/qfe adv_100T4_cap 0
nnd -set /dev/qfe adv_100fdx_cap 1
nnd -set /dev/qfe adv_100hdx_cap 0
nnd -set /dev/qfe adv_10fdx_cap 0
nnd -set /dev/qfe adv_10hdx_cap 0
nnd -set /dev/qfe adv_autoneg_cap 0
```

#### # dmfe-Interfaces (Netra X1)

```
# dmfe0
#nnd -set /dev/dmfe0 adv_100fdx_cap 1
#nnd -set /dev/dmfe0 adv_100hdx_cap 0
#nnd -set /dev/dmfe0 adv_10fdx_cap 0
#nnd -set /dev/dmfe0 adv_10hdx_cap 0
#nnd -set /dev/dmfe0 adv_autoneg_cap 0
# dmfe1
#nnd -set /dev/dmfe1 adv_100fdx_cap 1
#nnd -set /dev/dmfe1 adv_100hdx_cap 0
#nnd -set /dev/dmfe1 adv_10fdx_cap 0
#nnd -set /dev/dmfe1 adv_10hdx_cap 0
#nnd -set /dev/dmfe1 adv_autoneg_cap 0
```

```

# vge-Interfaces (Gigabit 1.x)
# vge0
#nnd -set /dev/vge instance 0
#nnd -set /dev/vge fdr_filter 1
#nnd -set /dev/vge link_negotiation 0
# vge1
#nnd -set /dev/vge instance 1
#nnd -set /dev/vge fdr_filter 1
#nnd -set /dev/vge link_negotiation 0

# ge-Interfaces (Gigabit 2.x and 3.x)
# example: forced 1000 Mbit/s, 802.3x Flow Control send and receive
# check if your NIC is 802.3x capable (nnd parameter pause_cap, asm_dir_cap)
# ge0
#nnd -set /dev/ge instance 0
#nnd -set /dev/ge adv_1000fdx_cap 1
#nnd -set /dev/ge adv_1000hdx_cap 0
#nnd -set /dev/ge adv_pauseTX 1
#nnd -set /dev/ge adv_pauseRX 1
#nnd -set /dev/ge adv_1000autoneg_cap 0
# ge1
#nnd -set /dev/ge instance 1
#nnd -set /dev/ge adv_1000fdx_cap 1
#nnd -set /dev/ge adv_1000hdx_cap 0
#nnd -set /dev/ge adv_pauseTX 1
#nnd -set /dev/ge adv_pauseRX 1
#nnd -set /dev/ge adv_1000autoneg_cap 0

# ce-Interfaces (GigaSwift 1.0 UTP)
# example: forced 1000 Mbit/s, 802.3x Flow Control send and receive
# ce0
#nnd -set /dev/ce instance 0
# Master values, change in forced 1000 Mbit/s back-to-back config only:
#nnd -set /dev/ce link_master 0 # one link partner 1, the other 0 (ce <=1.118)
#nnd -set /dev/ce master_cfg_enable 1 # enable usage of master_cfg_value (ce >1.118)
#nnd -set /dev/ce master_cfg_value 0 # one link partner 1, the other 0 (ce >1.118)
# Speed/Mode values, set the desired to 1, all other to 0:
#nnd -set /dev/ce adv_1000fdx_cap 1
#nnd -set /dev/ce adv_1000hdx_cap 0
#nnd -set /dev/ce adv_100fdx_cap 0
#nnd -set /dev/ce adv_100hdx_cap 0
#nnd -set /dev/ce adv_10fdx_cap 0
#nnd -set /dev/ce adv_10hdx_cap 0
#nnd -set /dev/ce adv_asmpause_cap 0
#nnd -set /dev/ce adv_pause_cap 1
#nnd -set /dev/ce adv_autoneg_cap 0
# ce1
#nnd -set /dev/ce instance 1
# Master values, change in forced 1000 Mbit/s back-to-back config only:
#nnd -set /dev/ce link_master 0 # one link partner 1, the other 0 (ce <=1.118)
#nnd -set /dev/ce master_cfg_enable 1 # enable usage of master_cfg_value (ce >1.118)
#nnd -set /dev/ce master_cfg_value 0 # one link partner 1, the other 0 (ce >1.118)
# Speed/Mode values, set the desired to 1, all other to 0:

```

```

#nnd -set /dev/ce adv_1000fdx_cap 1
#nnd -set /dev/ce adv_1000hdx_cap 0
#nnd -set /dev/ce adv_100fdx_cap 0
#nnd -set /dev/ce adv_100hdx_cap 0
#nnd -set /dev/ce adv_10fdx_cap 0
#nnd -set /dev/ce adv_10hdx_cap 0
#nnd -set /dev/ce adv_asmpause_cap 0
#nnd -set /dev/ce adv_pause_cap 1
#nnd -set /dev/ce adv_autoneg_cap 0

# bge-Interfaces (SF-V210, SF-V240, SB-1500, ... on-board only)
# example: forced 1000 Mbit/s, 802.3x Flow Control send and receive
# bge0
# Speed/Mode values, set the desired to 1, all other to 0:
#nnd -set /dev/bge0 adv_1000fdx_cap 1
#nnd -set /dev/bge0 adv_1000hdx_cap 0
#nnd -set /dev/bge0 adv_100fdx_cap 0
#nnd -set /dev/bge0 adv_100hdx_cap 0
#nnd -set /dev/bge0 adv_10fdx_cap 0
#nnd -set /dev/bge0 adv_10hdx_cap 0
#nnd -set /dev/bge0 adv_asym_pause_cap 0
#nnd -set /dev/bge0 adv_pause_cap 1
#nnd -set /dev/bge0 adv_autoneg_cap 0
# bge1
# Speed/Mode values, set the desired to 1, all other to 0:
#nnd -set /dev/bge1 adv_1000fdx_cap 1
#nnd -set /dev/bge1 adv_1000hdx_cap 0
#nnd -set /dev/bge1 adv_100fdx_cap 0
#nnd -set /dev/bge1 adv_100hdx_cap 0
#nnd -set /dev/bge1 adv_10fdx_cap 0
#nnd -set /dev/bge1 adv_10hdx_cap 0
#nnd -set /dev/bge1 adv_asym_pause_cap 0
#nnd -set /dev/bge1 adv_pause_cap 1
#nnd -set /dev/bge1 adv_autoneg_cap 0

# switch MTU path discovery off
#nnd -set /dev/ip ip_path_mtu_discovery 0
# TCP/UDP buffer tunings
#nnd -set /dev/tcp tcp_xmit_hiwat 65536
#nnd -set /dev/tcp tcp_rcv_hiwat 65536
#nnd -set /dev/udp udp_xmit_hiwat 65536
#nnd -set /dev/udp udp_rcv_hiwat 65536
# TCP congestion window size, for Solaris <= 2.5, only
#nnd -set /dev/tcp tcp_cwnd_max 65535
# Change TCP maximum retransmission interval, for Solaris <= 7, only
#nnd -set /dev/tcp tcp_rexmit_interval_max 60000
# Change number of TCP slow start initial packets to work around
# slow start bug in BSD stack and Windows. For Solaris <= 7, only
#nnd -set /dev/tcp tcp_slow_start_initial 2

;;

'stop')

```

```

echo "No kernel parameters changed."
;;


*)
echo "Usage: $0 {start|stop}"
;;

esac
exit 0

```

Posted by rdircio in Solaris at 10:28

## Speed of ce network interfaces in Solaris9



We have crappy switches where I work at, and we have to mind the speed and configuration of our cards, there's not much of a problem if we use hme,eri or qfe, but the ce... it is a different story, which we will explain now.

To configure the characteristics for each ce port (ce0, ce1, ce2...), we need to find out its device string, like this:

```

root@my_machine # grep ce /etc/path_to_inst
"/pci@11d,700000/pci@1/network@0" 0 "ce"
"/pci@11d,700000/pci@1/network@1" 1 "ce"
"/pci@13d,700000/pci@1/network@0" 2 "ce"
"/pci@13d,700000/pci@1/network@1" 3 "ce"

```

And write a file at /platform/sun4u/kernel/drv/ce.conf, like this:

```

#----- this is ce0
name="ce" parent="/pci@11d,700000/pci@1" unit-address="0"
adv_autoneg_cap=1 adv_1000fdx_cap=1 adv_1000hdx_cap=0 adv_100fdx_cap=0 adv_100hdx_cap=0
adv_10fdx_cap=0 adv_10hdx_cap=0 adv_100T4_cap=0;
#----- this is ce1
name="ce" parent="/pci@11d,700000/pci@1" unit-address="1"
adv_autoneg_cap=0 adv_1000fdx_cap=0 adv_1000hdx_cap=0 adv_100fdx_cap=1 adv_100hdx_cap=0
adv_10fdx_cap=0 adv_10hdx_cap=0 adv_100T4_cap=0;
#----- this is ce2
name="ce" parent="/pci@13d,700000/pci@1" unit-address="0"
adv_autoneg_cap=0 adv_1000fdx_cap=0 adv_1000hdx_cap=0 adv_100fdx_cap=1 adv_100hdx_cap=1
adv_10fdx_cap=0 adv_10hdx_cap=0 adv_100T4_cap=0;
#----- this is ce3
name="ce" parent="/pci@13d,700000/pci@1" unit-address="1"
adv_autoneg_cap=1 adv_1000fdx_cap=1 adv_1000hdx_cap=1 adv_100fdx_cap=1 adv_100hdx_cap=1
adv_10fdx_cap=0 adv_10hdx_cap=0 adv_100T4_cap=0;

```

You need to do an init 6, or reboot, in order for this to work.

You can see there are all these "adv" stuff, which doesn't mean "advanced", but "advertised", which means that the interface announces the switch what it can do...

1.- If `adv_autoneg_cap=0`, the ce interface chooses the highest value it can, for example, in ce2, it chooses 100fdx instead of 100hdx, although both have a "1" value. If the switch is 100hdx, it will never bring the link up.

2.- If `adv_autoneg_cap=1`, like on ce3, the ce interface tries to negotiate the highest value it has on, but can also go to the other value, if the switch can't do the highest speed the ce port can.

Posted by rdircio in Solaris at 07:08

Monday, February 14. 2005

## Miniexplorer new features

New features!

The miniexplorer for solaris now shows you:

- For each network interface, the speed(s) the link partner is capable of.
- The latest error messages in the system
- A little improvement on the interface

Check it out at <http://rdircio.no-ip.org/page>, at downloads!

Posted by rdircio in Solaris at 08:34

Friday, January 14. 2005

## To use a remote serial port, as if it was local!!!

Check out serial port sharing daemons:

Sercd:

<http://www.lysator.liu.se/~astrand/projects/sercd/>

Cyclades-serial-client:

<http://www.lysator.liu.se/~astrand/projects/cyclades-serial-client/>

Some of them are made for BSD or linux, so you might have to patch the code for Solaris...  
substitute:

change SOL\_IP to IPPROTO\_IP  
change ioctl FIONBIO to fcntl O\_NDELAY

Solaris does not have cfmakeraw!!!!

```
+ #if 0
cfmakeraw (&termios);
+ #else
+ termios.c_iflag &= ~(IGNBRK|BRKINT|PARMRK|ISTRIP|INLCR|IGNCR|ICRNL|IXON);
+ termios.c_oflag &= ~OPOST;
+ termios.c_lflag &= ~(ECHO|ECHONL|ICANON|ISIG|IEXTEN);
+ termios.c_cflag &= ~(CSIZE|PARENB);
+ termios.c_cflag |= CS8;
+ #endif
```

And you might not link setsockopt!!!!

Yes ... they are in libnsl and libsocket ... this is from "man  
setsockopt":

SYNOPSIS

```
cc [ flag ... ] file ... -lsocket -lnsl [ library ... ]
```

```
#include
```

```
#include
```

This goes into the Makefile of sercd

```
sercd_LDFLAGS = -lsocket -lnsl
```

Posted by rdircio in Solaris at 23:59

## Gemplus Card Readers and Solaris 9 on an Enterprise 220

Some strange project here requires to connect a gemplus cardreader using OpenCardFramework on Solaris...  
that is a Java implementation, for this OS.

You might not be in this situation anyhow, but at least, if you have a serial card reader like this, you must have these  
files, at these paths...

/usr/j2se/jre/lib/sparc/libSolarisSerialParallel.so  
and  
/usr/j2se/jre/lib/ext/comm.jar

they're easy to find elsewhere, just put them at those paths!

- rdircio

Posted by rdircio in Solaris at 14:54

Tuesday, January 11. 2005

## Adobe Central Output Server 5.5 installation, on Solaris 8

Copy the appropriate software to a temporary directory

```
# mkdir /mytmp
# cp /cdrom/cdrom0/solaris/software/central_output_server_5_5.tar /mytmp
# cd /mytmp ; tar xvf central_output_server_5_5.tar
```

Call setup

```
# cd /mytmp
# ./setup
```

This will start the installation procedure, asking these questions:

Name? (your company name)

PAC CODE (your PAC code)

Accept license? Yes

Products to install....

```
# Install Adobe Central Output Server products
#
# 1 + Central Utilities
# 2 + Central
# 3 + Central Agent(s)
# 4 + Mail Reader
# 5 Access remote Central
```

# Enter another item number or enter to finish your selection:

-

# Full pathname of Central home directory: (default

# '/usr/local/adobe/central')

- /usr/local/adobe/central

# Directory: '/usr/local/adobe/central' doesn't exist. Create it? (Yes/No):# (default 'Yes')

- Yes

# Directory in which to install executables? (default # '/usr/local/adobe/central/bin')

- /usr/local/adobe/central/bin

# Directory: '/usr/local/adobe/central/bin' doesn't exist. Create it? # (Yes/No): (default 'Yes')

- Yes

# Name to use for this computer? (default 'afpcal')

- afpcal

# Short name for this instance of Central: (default 'server')

- server

# Directory to install Central: (default '/usr/local/adobe/central/server')

```
- /usr/local/adobe/central/server
# Directory: '/usr/local/adobe/central/server' doesn't exist. Create it?# (Yes/No): (default Yes')
- Yes
# Collector directory for this instance of Central: (default # /usr/local/adobe/central/server/data')
- /usr/local/adobe/central/server/data
# Directory: '/usr/local/adobe/central/server/data' doesn't exist. Create it?
# (Yes/No): (default 'Yes')
- Yes
# Is this Central to be managed by the Central Daemon? (Yes/No): (default
# 'Yes')
- Yes
# Is this Central to start automatically whenever the Central Daemon starts?
# (Yes/No): (default 'Yes')
- Yes
```

Nota: Inicia el proceso de creaci?n de ambiente

```
# Creating subdirectories of /usr/local/adobe/central/server
# Copying jfserver to /usr/local/adobe/central/bin
# Copying jfserver.ini to /usr/local/adobe/central/server
# Copying jfserver.jmd to /usr/local/adobe/central/server
# Copying jfdaemon to /usr/local/adobe/central/bin
# Copying jfdaemon.ini to /usr/local/adobe/central/etc
# Copying jfnojob to /usr/local/adobe/central/bin
# Copying pacset to /usr/local/adobe/central/bin
# Copying eula_full.txt to /usr/local/adobe/central/bin
# Copying eula_eval.txt to /usr/local/adobe/central/bin
# Copying jfnojob.ini to /usr/local/adobe/central/server/etc
# copying default.def to /usr/local/adobe/central/server/etc
# copying default.def to /usr/local/adobe/central/server/etc
# Updating /usr/local/adobe/central/server/jfserver.ini
# Updating /usr/local/adobe/central/install.ini
# Updating /usr/local/adobe/central/server/etc/jfreg.ini
# Copying jfc to /usr/local/adobe/central/bin
# Copying jfc.ini to /usr/local/adobe/central/etc
# Copying jfjmd to /usr/local/adobe/central/bin
# Copying jfjmd.ini to /usr/local/adobe/central/etc
# Copying jfkick to /usr/local/adobe/central/bin
# Copying jfkick.ini to /usr/local/adobe/central/etc
# Copying jflisten to /usr/local/adobe/central/bin
# Copying jflisten.ini to /usr/local/adobe/central/etc
# Copying jfp to /usr/local/adobe/central/bin
# Copying jfp.ini to /usr/local/adobe/central/etc
# Copying jfq to /usr/local/adobe/central/bin
# Copying jfq.ini to /usr/local/adobe/central/etc
# Copying jfrm to /usr/local/adobe/central/bin
# Copying jfrm.ini to /usr/local/adobe/central/etc
# Copying jfstat to /usr/local/adobe/central/bin
# Copying jfstat.ini to /usr/local/adobe/central/etc
# Copying jfreg.ini to /usr/local/adobe/central
# Updating /usr/local/adobe/central/jfreg.ini
# Creating /usr/local/adobe/central/install.ini
# Creating /usr/local/adobe/central/install.ini
#
# Install agents
```

```
#
# 1+ Messaging Agent
# 2 + Print Agent
# 3 + Preview Agent
# 4 + Transformation Agent
# 5 + XML Import Agent
# Enter the number for the item you want to install:
(copy more files)
```

(Accept to see README)  
Adobe Central Output Server includes technology from Symbol  
Technologies, Inc. Portions (c) 1994 Symbol Technologies, Inc.

\*\*\*\*\*END OF FILE\*\*\*\*\*

Readme File  
Adobe Central Output Server, Version 5.5  
November 2002

```
#
# (when finished presents the next message)
```

```
# Adobe Central Output Server installation complete.
```

-----  
Add two ports to /etc/services:  
Jetform 1706/udp #event notification  
Jetform 1706/tcp #data transfer

Add a path to PATH:

```
PATH=${PATH}:/usr/local/adobe/central/bin
```

```
kill -HUP inet process ID
```

Posted by rdircio in Solaris at 16:50

Thursday, January 6. 2005

## dhcp client on solaris 8 or 9

The first thing to do is to create an empty

/etc/hostname.le0

- or-

/etc/hostname.hme0

- or-

/etc/hostname.qfe0

- or-

/etc/hostname.ce0

Depending on the interface you're going to use:

Next, you create

/etc/dhcp.le0

- or-

/etc/dhcp.hme0

- or-

/etc/dhcp.qfe0

- or-

/etc/dhcp.ce0

To get the hostname right, put this little script in to /etc/init.d/set\_hostname, with a symbolic link to it from /etc/rc2.d/S70set\_hostname.

```
-----  
#!/sbin/sh
```

```
HOSTNAME=`cat /etc/nodename`
```

```
echo "Setting hostname to $HOSTNAME... \c"
```

```
uname -S $HOSTNAME
```

```
echo "Done."  
-----
```

Posted by rdircio in Solaris at 15:03

Tuesday, December 14. 2004

## **pseudo-jni cards in SunFire 25K, FCX2-6562-L**

we've got some FCX2-6562-L cards, using a JNI chipset, not sold by JNI, but Sun. These have a JNI sibling, the FCX2-6562.

While one could easily get confused on these two models, they are not the same thing!!!! YOU SHOULD NOT use the driver for FCX2-6562, from [www.amcc.com](http://www.amcc.com), but the one that is integrated to the 25k.

The driver comes in one of these packages:

```
system SUNWjfca JNI Fibre Channel Adapter (FCA) Driver
system SUNWjfcau JNI Fibre Channel Adapter (FCA) (usr)
system SUNWjfcaux JNI Fibre Channel Adapter (FCA) (usr) (64-bit)
system SUNWjfcax JNI Fibre Channel Adapter (FCA) Driver (64-bit)
```

These packages are part of SFS 4.2, which in most cases is installed by Sun personnel with the OS. There is a list of patches that are necessary to be installed in order for this card to work, so don't go on a vanilla SFS 4.2.

These cards don't have a GUI like ezfiber, and are administered by the "cfgadm" command.

The manual we used to configure these cards is at:

<http://docs-pdf.sun.com/817-1245-10/817-1245-10.pdf>

See particularly page 18 "To Detect Fabric Devices Visible on a Host", to discover your SAN disks.

The document explains several methods to configure your cards, so it's worth to read it from start to end.

You need to do a

```
# cfgadm -al
```

to show all the controllers you can see, in our case, we see a controller like this:

```
c5 fc-fabric connected unconfigured unknown
```

so we proceed to map all the drives from that controller

```
# cfgadm -c configure c5
```

after that you'll do a `cfgadm -al` and see disks in that controller, you're ready then to run

```
# format
```

and see all your disks!

Posted by rdircio in Solaris at 11:21

## **FC64-1063 JNI Card**

I've got several of these plugged in to an Sun E10K on Solaris 8:

FC64-1063 JNI Card, single port, fiber channel, 1 Gb speed.

The driver I used for these is the one found at:

[amcc's page](#)

But I've put up a tarball with both the driver and its corresponding EZFiber here:

<http://rdircio.no-ip.org/jni>

You just download the file and do:

```
# gunzip JNI-SINGLE-E10K-DRIVER-USETHIS.tar.gz
# tar xvf JNI-SINGLE-E10K-DRIVER-USETHIS.tar
# cd JNI-SINGLE-E10K-DRIVER-USETHIS
# pkgadd -d ./fcaw.pkg
# tar xvf EZF_222.tar
# cd EZF_222
# ./install -silent
# init 6
```

After this, configured the port to be in "fabric mode", as opposed to "private loop", and I got away connecting to a CNT "director"

Posted by rdircio in Solaris at 10:40

Tuesday, December 7, 2004

## **Solaris truss error codes, explained**

Just check this file in your Solaris machine...


`/usr/include/sys/errno.h`

Credits go for this fine page: <http://www.princeton.edu/~psg/unix/Solaris/troubleshoot>

Posted by rdircio in Solaris at 08:14

Wednesday, December 1. 2004

## Veritas Administration Console (gui), a.k.a. vea



Just in case you wondered where...

Where the %\$#!#\$ is "vea" in Veritas Foundation Suite 4.0????

If you only need to install the gui in a workstation, install package VRTSobgui, usually in the "pkgs" directory inside your Veritas installation cd.

Posted by rdircio in Solaris at 09:19

## Copy a filesystem using ufsdump!


only for UFS!!!

This method is useful when you're using native Solaris UFS, or Solstice DiskSuite, it copies an entire FS to a directory, using a pipe.

```
ufsdump 0f - /dev/md/dsk/d32 | (cd /newopt;ufsrestore xf -)
```

Posted by rdircio in Solaris at 07:55

## Migrate your Solstice DiskSuite metadb from one disk to another



1. Annotate the # of cylinders in your original metadb...

```
partition> p
```

Current partition table (original):

Total disk cylinders available: 7506 + 2 (reserved cylinders)

Part	Tag	Flag	Cylinders	Size	Blocks
------	-----	------	-----------	------	--------

0	root	wm	446 - 7059	14.86GB (6614/0/0)	31165168
---	------	----	------------	--------------------	----------

1	swap	wu	0 - 412	950.22MB (413/0/0)	1946056
---	------	----	---------	--------------------	---------

2	backup	wm	0 - 7505	16.86GB (7506/0/0)	35368272
---	--------	----	----------	--------------------	----------

3	unassigned	wm	7060 - 7505	1.00GB (446/0/0)	2101552
---	------------	----	-------------	------------------	---------

```
4 unassigned wm 413 - 432 46.02MB (20/0/0) 94240
5 unassigned wu 0 0 (0/0/0) 0
6 unassigned wu 0 0 (0/0/0) 0
7 unassigned wu 0 0 (0/0/0) 0
```

In this case it is in slice 4 and it has 20 cylinders. Totals 46.02 Mb.

2. Next, create another partition using format, same size!

```
partition> p
```

```
Current partition table (original):
```

```
Total disk cylinders available: 7506 + 2 (reserved cylinders)
```

```
Part Tag Flag Cylinders Size Blocks
0 root wm 446 - 7059 14.86GB (6614/0/0) 31165168
1 swap wu 0 - 412 950.22MB (413/0/0) 1946056
2 backup wm 0 - 7505 16.86GB (7506/0/0) 35368272
3 unassigned wm 7060 - 7505 1.00GB (446/0/0) 2101552
4 unassigned wm 0 0 (0/0/0) 0
5 unassigned wm 413 - 432 46.02MB (20/0/0) 94240
6 unassigned wm 0 0 (0/0/0) 0
7 unassigned wm 0 0 (0/0/0) 0
```

Now, we created it on slice 5 on another disk!....

3. Add a metadb replica on c0t0d0s5

```
metadb -a -c 2 c0t0d0s5
```

```
Check it out!
```

```
bash-2.05# metadb -i
```

```
flags first blk block count
```

```
a u 16 8192 /dev/dsk/c0t0d0s5
```

```
a u 8208 8192 /dev/dsk/c0t0d0s5
```

```
a m p luo 16 8192 /dev/dsk/c0t1d0s4
```

```
a p luo 8208 8192 /dev/dsk/c0t1d0s4
```

4. The dangerous step!!!!, delete the original metadb!

```
metadb -d -f c0t1d0s4
```

```
bash-2.05# metadb -i
```

```
flags first blk block count
```

```
a u 16 8192 /dev/dsk/c0t0d0s5
```

```
a u 8208 8192 /dev/dsk/c0t0d0s5
```

5. Test it? REBOOT

```
# reboot
```

Posted by rdircio in Solaris at 07:53

## Settings in Solaris OBP to connect to a serial port

Like this!

```
#eeprom setenv auto-boot? false (optional, to make it stop before booting solaris)
#eeprom setenv ttya-rts-dtr-off true
#eeprom setenv ttya-ignore-cd true
#eeprom setenv input-device ttya (if using ttya)
#eeprom setenv output-device ttya (if using ttya)
#eeprom setenv diag_switch? true
#eeprom setenv diag-level max
#reset-all
```

when finishing connect, to reset default values use

```
#eeprom set-defaults
#eeprom reset
```

on machine from which you're connecting


```
set up /etc/remote to point "hardwire" to your serial port
use admintool to configure serial ports in expert mode and disable "initialize only" and enable "bidirectional"
rm /var/spool/locks/*
chmod 777 /dev/term/a
tip hardwire
```

when you want to exit use ~.

you could also use pmadm or/and sacadm to configure your port, instead of admintool

Posted by rdircio in Solaris at 07:51

## How to mirror a disk using Sun Volume Manager (Solstice Disk Suite)



```
#---
```

```
#--- Instructions i followed to create a mirror on 2 gb disks
```

```
#--- On solaris 9 with 3 partitions: /, /opt, and swap
```

```
#---
```

```
#-- Detach swap, and get a 50 Mb chunk from it to create the metadb
```

```
swap -d /dev/dsk/c0t0d0s1
```

```
df -k
```

```
#--- Use format to resize swap and leave a chunk of 50 mb, create s4 of 50mb
```

```
format
```

```
#--- Remount swap
```

```
swap -a /dev/dsk/c0t0d0s1
```

```
#--- copy the partition table to the mirror disk  
prtvtoc /dev/rdisk/c0t0d0s2 | fmthard -s - /dev/rdisk/c0t1d0s2
```

```
#--- Initialize the metadb  
metadb -a -f -c2 /dev/dsk/c0t0d0s4 /dev/dsk/c0t1d0s4
```

```
#--- Mirror on the / (root) partition  
metainit -f d10 1 1 c0t0d0s0  
metainit d20 1 1 c0t1d0s0  
metainit d30 -m d10  
metaroot d30  
cat /etc/vfstab
```

```
#--- Mirror on the swap partition  
metainit -f d11 1 1 c0t0d0s1  
metainit d21 1 1 c0t1d0s1  
metainit d31 -m d11
```

```
#--- Now for the /opt partition  
metainit -f d12 1 1 c0t0d0s3  
metainit d22 1 1 c0t1d0s3  
metainit d32 -m d12
```

```
#--- edit vfstab and replace raw disks for d31 and d32  
# #device device mount FS fsck mount mount  
# #to mount to fsck point type pass at boot options  
# #  
# fd - /dev/fd fd - no -  
# /proc - /proc proc - no -  
# /dev/md/dsk/d31 /dev/md/rdisk/d31 - - swap - no -  
# /dev/md/dsk/d30 /dev/md/rdisk/d30 / ufs 1 no logging  
# /dev/md/dsk/d32 /dev/md/rdisk/d32 /opt ufs 2 yes logging  
# swap - /tmp tmpfs - yes -
```

```
#--- Create a dummy empty hot spare pool  
metainit hsp001
```

```
#--- Reboot the system  
lockfs -fa  
init 6
```

```
#--- After reboot attach mirrors  
metattach d30 d20  
metattach d31 d21  
metattach d32 d22
```

```
#--- Check percentage of mirroring process  
metastat
```

```
#--- Make the mirror disk bootable  
installboot /usr/platform/`uname -i`/lib/fs/ufs/bootblk /dev/rdisk/c0t1d0s0
```

```
#--- Setup NVRAM to boot from disk0 or mirror
eeprom "boot-device=disk0 disk1"
eeprom "use-nvramrc?=false"
```

```
#--- in case of disk failure use "boot mirror"
  Posted by rdircio in Solaris at 07:49
```

Friday, November 26. 2004

## **copy filesystem using tar**

use it, it's good...

This is a useful way of copying different types of filesystems.

I used to copy from ufs to vxfs, got symlinks, permissions and creation dates right!

```
tar cpf - . | ( cd /target; tar xfp - )
```

Posted by rdircio in Solaris at 19:46

Wednesday, October 20, 2004

## samba for solaris from blastwave.org

Ok, previous article was, howto jump a proxy to get to blastwave...

Now, blastwave knows about package dependencies, but the machine i'm installing samba is not the same as the one i'm using to download packages...

You must then use "pkg-get -d " instead of "pkg-get -i ", the "-d" option does not get the packages we depend upon, so you have to download all manually.

You can get the package names from <http://www.blastwave.org/packages.php>, for example, to get berkeleydb4-4.1.25p1-SunOS5.8-sparc-CSW.pkg.gz, you type:

```
# pkg-get berkeleydb4
```

So... to install samba from blastwave you need these packages:

```
berkeleydb4-4.1.25p1-SunOS5.8-sparc-CSW.pkg.gz
common-1.4.1-SunOS5.8-sparc-CSW.pkg.gz
expat-1.95.7-SunOS5.8-sparc-CSW.pkg.gz
gettext-0.14.1-SunOS5.8-sparc-CSW.pkg.gz
gsed-4.1.1,REV=2004.07.14-SunOS5.8-sparc-CSW.pkg.gz
krb5_lib-1.3.1,REV=2004.03.25-SunOS5.8-sparc-CSW.pkg.gz
libconv-1.8,REV=2003.10.23-SunOS5.8-sparc-CSW.pkg.gz
libpopt-1.7,REV=2004.05.15-SunOS5.8-sparc-CSW.pkg.gz
libtool-1.5.2-SunOS5.8-sparc-CSW.pkg.gz
openldap-2.1.22-SunOS5.8-sparc-CSW.pkg.gz
openssl-0.9.7,REV=2004.04.26_rev=d-SunOS5.8-sparc-CSW.pkg.gz
samba-3.0.7,REV=2004.09.14-SunOS5.8-sparc-CSW.pkg.gz
samba_common-3.0.7,REV=2004.09.14-SunOS5.8-sparc-CSW.pkg.gz
samba_swat-3.0.7,REV=2004.09.14-SunOS5.8-sparc-CSW.pkg.gz
sasl-2.1.13,REV=2003.06.11-SunOS5.8-sparc-CSW.pkg.gz
```

you just need to do some pkgadd, and find who depends on who, install them in the right order, and you're done!

to start smbd you just do

```
# /opt/csw/sbin/smbd
```

BUT... before you do that, create a basic smb.conf, mine is in /opt/csw/etc/samba/smb.conf, and has this:

```
[global]
security = SHARE
```

```
[qa]
comment = qa disk
path = /qa
read only = No
guest ok = Yes
```

This allows everyone everywhere to see your share, and read and write to it, as long as it has 777 mode.

You might want to do other more secure stuff, but this config here is just to make sure it works.

Posted by rdircio in Solaris at 17:11

## **Blastwave.org and wget and a proxy**

wget can get heavy loads

Again on jumping over the proxy....

blastwave.org has lots of useful packages for solaris, which install in a non-standard path, this is useful sometimes to help unmix your environment, and isolate your errors.

Yea, but I've got a proxy... again, so I've got the useful pkg-get tool from <http://www.blastwave.org/> and downloaded wget (required) from the same page.

Out of the box, you should be able to use it as:

```
# pkg-get -i samba
```

But it relies on wget, and wget must do a connection on http or ftp to get the packages....

So, doing a truss on wget, got the preferences file!!!!

```
/opt/csw/etc/wgetrc
```

It was not present, so I just created it with a one-liner:

```
http_proxy = http://proxyuser:proxypass@150.100.204.19:8080
```


Where proxyuser is my proxy's user, and proxy pass is my proxy's pass(obvious).

And it just works.

Posted by rdircio in Solaris at 11:34

Friday, October 15. 2004

## Obsoleted Kernel Parameters in Solaris 9



As obscure as usual, today I found out that BMC's controlm does recognise my Solaris 9 as Solaris 9, but.... When it checks for prerequired kernel parameters, it checks for some that are obsoleted in Solaris 9!

These are:

shmsys:shminfo\_shmmin

shmsys:shminfo\_shmseg

You can find a complete list at:

<http://docs.sun.com/db/doc/806-7009/6jftnqskp?q=shmmin+solaris+9&a=view>

if the url does not take you directly, try navigate:

Solaris 9 System Administrator Collection >> Solaris Tunable Parameters Reference Manual >> A. Tunable Parameter Change History

Posted by rdircio in Solaris at 10:20

Monday, October 11. 2004

## Sun One Directory Server 5.2 installation on solaris 9

Sun One Directory Server 5.2 installation on solaris 9

This is the package-based installation

Install order:

```
#pkgadd -d . SUNWtls
#pkgadd -d . SUNWsas1
#pkgadd -d . SUNWpr
#pkgadd -d . SUNWldk
#pkgadd -d . SUNWjss
#pkgadd -d . SUNWicu
#cd SunOS5.9/
#patchadd 114049-04
#patchadd 115342-01
#cd ..
#pkgadd -d . SUNWicu
#pkgadd -d . SUNWdsvu
#pkgadd -d . SUNWdsvpl
#pkgadd -d . SUNWdsvr
#pkgadd -d . SUNWdsvpl
#pkgadd -d . SUNWdsvh
#pkgadd -d . SUNWasvc
#pkgadd -d . SUNWdsvcp
#pkgadd -d . SUNWasvu
#pkgadd -d . SUNWasvr
#pkgadd -d . SUNWasvcp
```

Then run:

```
# /usr/sbin/directoryserver idsktune -q > idsktune.out
```

and check idsktune.out, fix everythng marked as ERROR, in my case, none.

Ok, then run

```
# /usr/sbin/directoryserver configure
```

had these errors:

```
---[slapd-netra]: starting up server ...
---[slapd-netra]: [14/Oct/2004:00:08:23 -0600] - ERROR<4167> - Startup - ---conn=-1
---op=-1 msgId=-1 - System error Load library ---/var/mps/serverroot/lib/acl-plugin.
---so: error ld.so.1: ./ns-slapd: fatal: relocation error: file
---/var/mps/serverroot/lib/acl-plugin.so: symbol PR_GetSysfdTableMax: referenced---symbol not found

---[slapd-netra]: [14/Oct/2004:00:08:23 -0600] - ERROR<4140> - plugin_setup -
---conn=-1 op=-1 msgId=-1 - Configuration error Could not load symbol ---"acl_init"
---from library "/var/mps/serverroot/lib/acl-plugin.so" for plugin ACL Plugin

---[slapd-netra]: [14/Oct/2004:00:08:23 -0600] - ERROR<4112> - Bootstrap config -
---conn=-1 op=-1 msgId=-1 - Configuration error Unable to load plugin "cn=ACL
```

---Plugin,cn=plugins,cn=config".

---[slapd-netra]: Server not running!! Failed to start ns-slapd process.

APPLIED PATCH 114049-12, then ran  
# /usr/sbin/directoryserver unconfigure  
# /usr/sbin/directoryserver configure  
OK, the server STARTED!

root# /usr/sbin/mpsadmserver configure

Had problems with X11 to linux, so  
NOTE:Fixed X11 problems by adding the linux IP address and hostname to /etc/hosts in the solaris box

# /usr/sbin/mpsadmserver configure -nodisplay  
WARNING: Package SUNWj3rt has multiple instances:  
SUNWj3rt  
1.4.0,REV=2002.06.20.00.00 /usr  
SUNWj3rt.2  
1.4.2,REV=2003.09.16.20.27 /usr  
WARNING: only SUNWj3rt is used by Administration Server 5.2

Enter the fully qualified name of the computer

Fully Qualified Computer Name [netra.domain.com] {"<" goes back, "!" exits):

Server Root Information

Server Root: : /var/mps/serverroot

Choose the type of installation you prefer from the following choices:

Express - Installation option choices are made automatically. The easiest installation and recommended for evaluating the product.

Typical - Software will be installed with the most common options. Recommended for most deployments.

1. Express
2. Typical

What would you like to do [2] {"<" goes back, "!" exits)?

Configuration Directory Server Administrator

Administrator ID [admin] {"<" goes back, "!" exits):  
Password:

Administration Domain

Administration Domain [casa.com] {"<" goes back, "!" exits}:

The Administration Server runs on a different network port from other servers.  
Specify the number of the port.

Administration Port [390] {"<" goes back, "!" exits}:

Installing Administration Server

| -1%-----25%-----50%-----75%-----100%| Checking connection to the Configuration Directory Server... done.

Updating Administration Server layout... done.

Updating links between Server Root and Administration Server layout... done.

Registering Administration Server with Configuration Directory Server... done.

Loading Administration Server tasks... done.

Loading global Administration Server configuration... done.

Generating configuration files ... done.

Configuration of the Administration Server succeeded.

SunONE-WebServer-Enterprise/6.0SP3 B05/14/2003 17:58


warning: daemon is running as super-user

[LS ls1] http://netra.domain.com, port 390 ready to accept requests

startup: server started successfully

Posted by rdircio in Solaris at 23:52

## Just remembering how to connect a v480 to an annex serial terminal concentrator



Some time ago, i remember that i used to connect v480s to an annex terminal concentrator, using the serial port, and Netras 20 at the RSC card.

Both the Sun v480 and the annex have rj45 ports.

The cable pinout was found at:

<http://www.stokely.com/unix.serial.port.resources/A-B-Ycablepinout.html>

IMPORTANT NOTICE: The cable only works one way!!!, it means, there's a specific end that goes to the v480 and another to the annex.

Try these two, most likely the first one works

annex SUN V480

Pin Signal Pin Signal

-----

1 RTS ----- 8 CTS

2 DSR ----- 2 DTR

3 CD --Not Used-- 4 GRND  
4 RD ----- 3 TD  
5 TD ----- 6 RD  
6 GRND ----- 5 GRND  
7 DTR ----- 7 DSR  
8 CTS ----- 1 RTS

V480 LOM RJ45 ANNEX RJ45  
Pin Signal Pin Signal

-----  
8 CTS --- 1 RTS  
7 DSR --- 2 DTR  
6 RXD --- 3 TXD  
4 GND --- 6 GND  
3 TXD --- 5 RXD  
2 DTR --- 7 DSR  
1 RTS --- 8 CTS

Posted by rdircio in Solaris at 11:22

## Useful commands of Veritas Foundation Suite 4.0

veritas!

These commands may help while adding and deleting disks, creating disk groups or volumes.

Initialize disk to be used by Veritas

```
# /etc/vx/bin/vxdisksetup -i c3t0d15
```

Create disk group

```
# vxdg init prueba_dg disco1=c3t0d15s2
```

Add disks to disk group

```
# vxdg -g prueba_dg adddisk disco2=c3t0d16s2
```

Create a volume, using vxassist. See vxedit

```
# vxassist -g prueba_dg make vol01 200m
```

(Create a concatenation)

```
# vxprint
```

```
bash-2.03# mkfs -F vxfs /dev/vx/rdisk/prueba_dg/vol01
```

Mount the volume

```
# mount -F vxfs /dev/vx/dsk/prueba_dg/vol01 /vol01
```

To verify largefiles support on a fs

```
# fsadm /vol01
```

To start a volume

```
# vxvol -g prueba_dg start vol01
```

Remove a volume

```
# vxassist -g prueba_dg remove volume vol02
```

Create a mirrored volume

```
# vxassist -g prueba_dg make vol02 2g layout=mirror
```

Resize a filesystem

```
# /etc/vx/bin/vxresize -F vxfs -g prueba_dg vol01 210m
```

Resize a filesystem, adding space

```
# /etc/vx/bin/vxresize -F vxfs -g prueba_dg vol01 +20m
```

Posted by rdircio in Solaris at 14:37

## Veritas Foundation Suite 4.0 on Solaris 8 (just the installation steps)

Veritas Foundation Suite on Sun

1. First, install SUNWsan, from SFS\_base\_packages.S8
2. Install patches  
108982-09 111412-15  
111095-17 111413-14
3. init 6

bash-2.03# ./installsf

VERITAS STORAGE FOUNDATION 4.0 INSTALLATION PROGRAM

Copyright (c) 2004 VERITAS Software Corporation. All rights reserved.

VERITAS, the VERITAS Logo and all other VERITAS product names and slogans are trademarks or registered trademarks of VERITAS Software Corporation. VERITAS and the VERITAS Logo Reg. U.S. Pat. & Tm. Off. Other product names and/or slogans mentioned herein may be trademarks or registered trademarks of their respective companies.

Enter the system names separated by spaces on which to install SF: bastanteo

Checking OS version on bastanteo ..... SunOS 5.8

Initial system check completed successfully.

Press [Return] to continue:

VERITAS STORAGE FOUNDATION 4.0 INSTALLATION PROGRAM

VERITAS Infrastructure package installation:

Installing VERITAS Infrastructure packages on bastanteo:

Checking VRTScpi package ..... not installed  
Checking VRTSvlic package ..... not installed  
Checking file system space ..... required space is available  
Installing VRTScpi 4.0.5 on bastanteo ..... Done  
Installing VRTSvlic 3.02.005d on bastanteo ..... Done

VERITAS Infrastructure packages installed successfully.

Press [Return] to continue:

VERITAS STORAGE FOUNDATION 4.0 INSTALLATION PROGRAM

SF Licensing Verification:

Checking SF license key on bastanteo ..... not licensed  
Enter a SF license key for bastanteo: [?] 3JPU-TDNC-NW42-WSHS-LPPZ-RCC  
Registering VERITAS Storage Foundation Standard DEMO key on bastanteo

Do you want to enter another license key for bastanteo? [y,n,q,?] (n)

SF licensing completed successfully.

Press [Return] to continue:

## VERITAS STORAGE FOUNDATION 4.0 INSTALLATION PROGRAM

installsf can install the following optional SF packages:

VRTSobgui VERITAS Enterprise Administrator  
VRTSvmman VERITAS Volume Manager Manual Pages  
VRTSvmdoc VERITAS Volume Manager Documentation  
VRTStep VERITAS Task Provider  
VRTSap VERITAS Action Provider  
VRTSfsman VERITAS File System Manual Pages  
VRTSfsdoc VERITAS File System Documentation

- 1) Install all of the optional packages
- 2) Install none of the optional packages
- 3) View package descriptions and select optional packages

Select the optional packages to be installed on all systems? [1-3,q,?] (1)

## VERITAS STORAGE FOUNDATION 4.0 INSTALLATION PROGRAM

installsf will install the following SF packages:  
VRTSperl VERITAS Perl 5.8.0 Redistribution  
VRTSob VERITAS Enterprise Administrator Service  
VRTSmuob VERITAS Enterprise Administrator Service Localized Package  
VRTSobgui VERITAS Enterprise Administrator  
VRTSvxvm VERITAS Volume Manager Binaries  
VRTSvmman VERITAS Volume Manager Manual Pages  
VRTSvmdoc VERITAS Volume Manager Documentation  
VRTSvmpro VERITAS Volume Manager Management Services Provider  
VRTSfspro VERITAS File System Management Services Provider  
VRTSalloc VERITAS Volume Manager Intelligent Storage Provisioning

VRTSddlpr VERITAS Device Discovery Layer Services Provider  
VRTSvxfs VERITAS File System  
VRTStep VERITAS Task Provider  
VRTSap VERITAS Action Provider  
VRTSfppm VERITAS File Placement Policy Manager  
VRTSfsman VERITAS File System Manual Pages  
VRTSfsdoc VERITAS File System Documentation

Press [Return] to continue:

## VERITAS STORAGE FOUNDATION 4.0 INSTALLATION PROGRAM

Checking system installation requirements:

Checking SF installation requirements on bastanteo:

Checking VRTSperl package ..... not installed  
Checking VRTSob package ..... not installed  
Checking VRTSmuob package ..... not installed  
Checking VRTSobgui package ..... not installed  
Checking VRTSvxvm package ..... not installed  
Checking VRTSvmman package ..... not installed  
Checking VRTSvmdoc package ..... not installed  
Checking VRTSvmpro package ..... not installed  
Checking VRTSfspro package ..... not installed  
Checking VRTSalloc package ..... not installed  
Checking VRTSddlpr package ..... not installed  
Checking VRTSvxfs package ..... not installed  
Checking VRTStep package ..... not installed  
Checking VRTSap package ..... not installed  
Checking VRTSfppm package ..... not installed  
Checking VRTSfsman package ..... not installed  
Checking VRTSfsdoc package ..... not installed  
Checking required SunOS patch 111413-06 on bastanteo ..... 111413-14 installed  
Checking required SunOS patch 108993-18 on bastanteo ..... 108993-36 installed  
Checking required SunOS patch 108528-18 on bastanteo ..... 108528-29 installed  
Checking required SunOS patch 108528-02 on bastanteo ..... 108528-29 installed  
Checking required SunOS patch 108901-03 on bastanteo ..... not installed

The following required OS patches were not found on bastanteo:

108901-03

It is strongly recommended these patches be installed on systems before installing Storage Foundation

Installation of Storage Foundation without required patches should only be performed

if you are certain that a superceded version of the patch is already installed on this system

Do you want to continue installation of Storage Foundation on these systems? [y,n,q] (n) y

Checking VERITAS patch 115209 ..... not installed  
Checking VERITAS patch 115212 ..... not installed  
Checking VERITAS patch 115210 ..... not installed  
Checking VERITAS patch 115217 ..... not installed  
Checking for any Solaris patch issues ..... None  
Checking file system space ..... required space is available  
Checking for any AP driver issues ..... None  
Checking vxsvc process ..... not running  
Checking vxportal driver ..... not running  
Checking fdd driver ..... not running  
Checking qlong driver ..... not running  
Checking vxfs driver ..... not running

Installation requirement checks completed successfully.

Press [Return] to continue:

## VERITAS STORAGE FOUNDATION 4.0 INSTALLATION PROGRAM

### VERITAS STORAGE FOUNDATION 4.0 INSTALLATION PROGRAM

It is possible to install SF packages without performing configuration.

It is optional to configure SF now. If you choose to configure SF later, you can either do so manually or run the `installsf -configure` command.

Are you ready to configure SF? [y,n,q] (y)

### VERITAS STORAGE FOUNDATION 4.0 INSTALLATION PROGRAM

`installsf` will now ask sets of SF configuration-related questions.

When a [b] is presented after a question, 'b' may be entered to go back to the first question of the configuration set.

When a [?] is presented after a question, '?' may be entered for help or additional information about the question.

Following each set of questions, the information you have entered will be presented for confirmation. To repeat the set of questions and correct any previous errors, enter 'n' at the confirmation prompt.

No configuration changes are made to the systems until all configuration questions are completed and SF is installed successfully.

Press [Return] to continue:

## VERITAS STORAGE FOUNDATION 4.0 INSTALLATION PROGRAM

Installing Storage Foundation Standard 4.0 on bastanteo:

Installing VRTSperl 4.0.2 on bastanteo ..... Done 1 of 21 steps  
Installing VRTSob 3.2.514.0 on bastanteo ..... Done 2 of 21 steps  
Installing VRTSmuob 3.2.514.0 on bastanteo ..... Done 3 of 21 steps  
Installing VRTSobgui 3.2.514.0 on bastanteo ..... Done 4 of 21 steps  
Installing VRTSvxvm 4.0 on bastanteo ..... Done 5 of 21 steps  
Installing VRTSvmman 4.0 on bastanteo ..... Done 6 of 21 steps  
Installing VRTSvmdoc 4.0 on bastanteo ..... Done 7 of 21 steps  
Installing VRTSvmpro 4.0 on bastanteo ..... Done 8 of 21 steps  
Installing VRTSfspro 4.0 on bastanteo ..... Done 9 of 21 steps  
Installing VRTSalloc 4.0 on bastanteo ..... Done 10 of 21 steps  
Installing VRTSddlpr 4.0 on bastanteo ..... Done 11 of 21 steps  
Installing VRTSvxfs 4.0 on bastanteo ..... Done 12 of 21 steps  
Installing VRTStep 1.20.025 on bastanteo ..... Done 13 of 21 steps  
Installing VRTSap 2.00.015 on bastanteo ..... Done 14 of 21 steps  
Installing VRTSfppm 4.0 on bastanteo ..... Done 15 of 21 steps  
Installing VRTSfsman 4.0 on bastanteo ..... Done 16 of 21 steps  
Installing VRTSfsdoc 4.0 on bastanteo ..... Done 17 of 21 steps  
Adding patch 115209-05 on bastanteo ..... Done 18 of 21 steps  
Adding patch 115212-05 on bastanteo ..... Done 19 of 21 steps  
Adding patch 115210-05 on bastanteo ..... Done 20 of 21 steps  
Adding patch 115217-01 on bastanteo ..... Done 21 of 21 steps

Storage Foundation Standard installation completed successfully.

Press [Return] to continue:

The enclosure-based naming scheme is a feature of Volume Manager. It allows one to reference disks using a symbolic name that is more meaningful than the operating

system's normal device access name. This symbolic name is typically derived from the array name.

Evaluating which systems can be configured for the enclosure-based naming scheme...

System bastanteo is eligible -- can configure.

Do you want to set up the enclosure-based naming scheme? [y,n,q,?] (n)

## VERITAS STORAGE FOUNDATION 4.0 INSTALLATION PROGRAM

Configuring Storage Foundation Standard:

Storage Foundation Standard configured successfully.

Press [Return] to continue:

## VERITAS STORAGE FOUNDATION 4.0 INSTALLATION PROGRAM

Do you want to start Storage Foundation Standard processes now? [y,n,q] (y)

Evaluating which systems can now be started...

System bastanteo is eligible -- can be started.

Preparing to start VxVM on target systems...

Begin starting VxVM on system bastanteo

Disabling enclosure-based naming ..... Done

Starting vxconfigd for VxVM ..... Succeeded

Done starting VxVM on system bastanteo

Done with starting VxVM on target systems...

Press [Return] to continue:

Many Volume Manager commands affect the contents or configuration of a disk group.

Such commands require that the user specify a disk group. This is accomplished by

using the -g option of a command or setting the VXVM\_DEFAULTDG environment variable.

An alternative to these two methods is to configure the default disk group of a system.

Evaluating which systems can now have their default disk group configured...

System bastanteo is eligible -- can set the default disk group.

Do you want to set up the default disk group for each system? [y,n,q,?] (y) n

Preparing to start daemons on target systems...

Starting vxcached on bastanteo ..... Success

Storage Foundation Standard was started successfully.

Press [Return] to continue:

## VERITAS STORAGE FOUNDATION 4.0 INSTALLATION PROGRAM

Installation of Storage Foundation Standard 4.0 has completed successfully.

The installation summary is saved at:

/opt/VRTS/install/logs/installsf1007135010.summary

The installsf log is saved at:

/opt/VRTS/install/logs/installsf1007135010.log

The installation response file is saved at:

/opt/VRTS/install/logs/installsf1007135010.response

Reboot all systems on which VxFS was installed or upgraded.

shutdown -y -i6 -g0

See the VERITAS File System Administrators Guide for information on using VxFS.

```
-----  
bash-2.03# vxdisk list  
DEVICE TYPE DISK GROUP STATUS  
c0t0d0s2 auto:none - - online invalid  
c0t1d0s2 auto:none - - online invalid  
c0t2d0s2 auto:none - - online invalid  
c0t8d0s2 auto:none - - online invalid
```

```
c0t9d0s2 auto:none - - online invalid
c0t10d0s2 auto:none - - online invalid
c3t0d15s2 auto - - error
c3t0d16s2 auto - - error
c3t0d17s2 auto - - error
c3t0d18s2 auto - - error
c3t0d19s2 auto - - error
c3t0d20s2 auto - - error
c3t0d21s2 auto - - error
c3t0d22s2 auto - - error
```

---

Use format to label all your disks.

Check for new disk configuration after labeling  
bash-2.03# vxdctl enable

Done!

Posted by rdircio in Solaris at 11:50