

Image Backup of Small Files

David Stes

email: stes@pandora.be

Februari 12, 2010

Abstract

We show in this paper some experiences and performance numbers of making image backups and guest OS level backups of Oracle SUN VirtualBox virtual machines using EMC NetWorker. Specifically, we do some *postmark* benchmarks to check how the performance numbers are affected by the number of small files to backup. Also, the tests were done on different two SUN Solaris machines, one AMD Opteron and the other one with Intel XEON processors, with different RAID cards and network cards, and we discuss BMR (bare metal recovery) of a virtual machine on different physical hardware.

1 Test hardware

For our test we use EMC NetWorker version 7.5.1 and Oracle SUN VirtualBox version 3.1.2 on two different machines :

```
# uname -a
SunOS opensolaris 5.11 snv_125 i86pc i386 i86pc
# uname -a
SunOS vbox 5.11 snv_125 i86pc i386 i86pc
```

They both run a fairly recent version of OpenSolaris :

```
# cat /etc/release
                                OpenSolaris Development snv_125 X86
Copyright 2009 Sun Microsystems, Inc. All Rights Reserved.
Use is subject to license terms.
                                Assembled 12 October 2009
```

The difference is that one machine is a rather old W2100z SUN desktop server, while the other machine is a modern and powerful X4275 rackmounted server.

The machine *opensolaris* is an AMD machine with only 1 GB RAM and with a Broadcom *bge0* network card, using an internal 80 GB pci-ide disk :

```
# prtdiag | head
System Configuration: Sun Microsystems W1100z/2100z
BIOS Configuration: Sun Microsystems R01-B5 S1 09/29/2006
```

```
==== Processor Sockets =====
```

Version	Location Tag
AMD	CPU1

```
# prtconf | head
System Configuration: Sun Microsystems i86pc
Memory size: 1023 Megabytes
```

```
# dladm show-link
LINK      CLASS      MTU      STATE      BRIDGE      OVER
bge0      phys       1500     up         --          --
vboxnet0  phys       1500     unknown   --          --
```

The other server is a powerful X4275 with two XEON CPU's and much more RAM, built-in RAID controller using six SAS disks and SSD (solid state disk) and with a quad-ethernet Intel card :

```
# prtdiag
System Configuration: SUN MICROSYSTEMS SUN FIRE X4275 SERVER
BIOS Configuration: American Megatrends Inc. 07060215 06/19/2009
BMC Configuration: IPMI 1.5 (KCS: Keyboard Controller Style)
```

```
==== Processor Sockets =====
```

Version	Location Tag
Intel(R) Xeon(R) CPU	X5570 @ 2.93GHz CPU 1
Intel(R) Xeon(R) CPU	X5570 @ 2.93GHz CPU 2

```
# prtconf | more
System Configuration: Sun Microsystems i86pc
Memory size: 24568 Megabytes
```

```
# dladm show-link
LINK      CLASS      MTU      STATE      BRIDGE      OVER
igb0      phys       1500     up         --          --
igb1      phys       1500     down       --          --
igb2      phys       1500     down       --          --
igb3      phys       1500     down       --          --
vboxnet0  phys       1500     unknown   --          --
```

So the hardware is rather different, with the new X4275 of course dramatically outperforming the old W2100z box, as we will see.

On the two OpenSolaris machines, we installed the EMC NetWorker software following the usual procedure :

```
# pkginfo -l | grep LGT
  PKGINST: LGTOclnt
  PKGINST: LGToman
  PKGINST: LGTOnode
  PKGINST: LGTOserv
```

Also, the Oracle SUN VirtualBox is installed :

```
# pkginfo -l SUNWvbox
  PKGINST: SUNWvbox
  NAME: Sun VirtualBox
  CATEGORY: application
  ARCH: i386
  VERSION: 3.1.2,REV=2009.12.17.15.01.56127
  BASEDIR: /
  VENDOR: Sun Microsystems, Inc.
  DESC: A powerful PC virtualization solution
```

2 Configuring a Windows 2003 test server

All commands will be performed as user *root*.

This is important since EMC NetWorker runs as root, so if we will use scripts (savenpc commands) to make backups, these scripts will also run as *root*.

On the (slow) AMD machine, we create a virtual machine. As user *root*, we execute the following command :

```
# VBoxManage list ostypes | grep 2003
ID: Windows2003
Description: Windows 2003
ID: Windows2003_64
Description: Windows 2003 (64 bit)
```

Next, we create a virtual machine that we will locate on a special filesystem */vbox* :

```
# VBoxManage createvm --name win2003 --basefolder /vbox/Machines --ostype Windows2003
VirtualBox Command Line Management Interface Version 3.1.2
(C) 2005-2009 Sun Microsystems, Inc.
All rights reserved.

Virtual machine 'win2003' is created.
UUID: 96be1e8c-854a-4dc8-a66a-c9acd57f93d8
```

Next, we register the virtual machine, and check the VirtualBox directory in the */root* home directory :

```
# VBoxManage registervm /vbox/Machines/win2003/win2003.xml
VirtualBox Command Line Management Interface Version 3.1.2
(C) 2005-2009 Sun Microsystems, Inc.
All rights reserved.
```

```
# ls /root/.VirtualBox
compreg.dat      VirtualBox.xml  xpti.dat
```

Finally, we launch the X11 front-end to a remote machine, to configure the Oracle SUN VirtualBox as user *root*. We do this remotely because we do not logon locally as user *root* on the machines :

```
# DISPLAY=10.0.0.1:0.0
# export DISPLAY
# VirtualBox &
```

In the GUI, we add some storage (IDE hard drive and CD-rom) and we install the Windows 2003 R2 operating system from an ISO image of Microsoft Windows 2003 R2.

```
# VBoxManage showvminfo win2003
```

```
Storage Controller Name (0):          IDE Controller
Storage Controller Type (0):          PIIX4
Storage Controller Instance Number (0): 0
Storage Controller Max Port Count (0): 2
Storage Controller Port Count (0):    2
IDE Controller (0, 0): /vbox/Harddisks/win2003.vdi (UUID: 83eaa4ba-cc69-474d-9bbc-2e2a87
```

The virtual machine thinks it uses an IDE 20GB harddisk (which is mapped to the `/vbox/Harddisks/win2003.vdi` file).

Next we start the machine and grant VRDP access to the virtual machine, so to control the console via RDP (remote desktop). We do this via the `VBoxHeadless` command.

```
# VBoxHeadless -s win2003 &
VirtualBox Headless Interface 3.1.2
(C) 2008-2009 Sun Microsystems, Inc.
All rights reserved.
```

```
Listening on port 3389.
```

On a different machine we use `rdesktop` to connect to the console, and then we can login into the win2003 server and type commands such as `ipconfig` :

```
ipconfig /all
Ethernet adapter Local Area Connection:

    Connection-specific DNS Suffix  . : source.be
    Description . . . . . : AMD PCNET Family PCI Ethernet Adapter
```

Note that this virtual machine believes that it runs on some generic PC hardware (as configurable in Oracle SUN VirtualBox).

The network card is bridged to the *bge0* adapter, so that it can obtain an ip address via DHCP, for example, or do EMC NetWorker backups over the LAN.

Inside the virtual machine, we install the Oracle SUN VirtualBox guest additions (via the CD-ROM). This ISO image provides some guest additions that make the mouse pointer and keyboard behave friendly.

Also, we install EMC NetWorker (in the guest OS), and run the usual Windows setup to install this package.

Finally, we also install the *postmark* utility in the Windows 2003 machine. This package can be downloaded from:

```
http://www.linuxbox.com/tiki/downloads/postmark151w
```

We use the *postmark* tool to create lots of small files in the Windows operating system.

3 Image Backup using savenpc on the old W2100z

The backup command *savenpc* is part of EMC NetWorker and it allows us to temporarily save the virtual machine to disk, stopping it, and then to make a backup of the image files, and using a post command, to restart the virtual machine.

We use the following configuration file for a group called VBox :

```
# cat /nsr/res/VBox.res
type: savenpc;
precmd: "/usr/bin/VBoxPrcmd >> /nsr/logs/savestate 2>&1";
pstcmd: "/usr/bin/VBoxPstcmd >> /nsr/logs/savestate 2>&1";
timeout: "12:00:00";
abort precmd with group: No;
```

The script *VBoxPrcmd* stops the virtual machine :

```
# cat /usr/bin/VBoxPrcmd
#!/bin/sh

echo "VBoxManage controlvm savestate begin"
date
VBoxManage controlvm win2003 savestate
echo "VBoxManage controlvm savestate ended"
date
```

At the end of the EMC NetWorker backup, we restart the virtual machine :

```
# cat /usr/bin/VBoxPstcmd
#!/bin/sh

echo "VBoxHeadless startvm"
date
nohup VBoxHeadless -s win2003 &
echo "VBoxPstcmd end"
date
```

The OpenSolaris client is configured with save set /vbox (which is where the win2003 vdi file resides) and with backup command savenpc.

When we do this on the old AMD server *opensolaris* we get the following results , where we have to keep in mind that /vbox is located on the internal 80 GB pci-ide disk :

```
# time savegrp -l 0 VBox
real    4m34.103s
user    0m0.024s
sys     0m0.047s

# mminfo -avr client,name,totalsize,nfiles
opensolaris /vbox                2723028596    14
opensolaris bootstrap            136012      198
opensolaris index:opensolaris    7460        5
```

The backup performs at around 9.7 MB/s (10 MB/s) which is also what we see via nsrwatch :

```
/B2D          file          B2D          writing at 9744 KB/s, 1878 MB
```

Note that there are only 14 files in the saveset /vbox, the files are the big files located under /vbox (including a /vbox/Export which is a manual export, not needed for our test) :

```
/vbox
/vbox/Harddisks
/vbox/Harddisks/win2003.vdi
/vbox/Export
/vbox/Export/win2003.vmdk
/vbox/Export/win2003.mf
/vbox/Export/win2003.ovf
/vbox/Machines
/vbox/Machines/win2003
/vbox/Machines/win2003/win2003.xml
/vbox/Machines/win2003/Logs
/vbox/Machines/win2003/Logs/VBox.log
```

```
# VBoxManage list runningvms
VirtualBox Command Line Management Interface Version 3.1.2
(C) 2005-2009 Sun Microsystems, Inc.
All rights reserved.
```

```
"win2003" {96be1e8c-854a-4dc8-a66a-c9acd57f93d8}
```

4 Image Restore on the new X4275

Instead of installing a new virtual machine on the X4275, we will try to recover the Windows 2003 machine from the SUN W2100z (AMD) to the SUN X4275 (Intel XEON).

On the X4275, we use the SSD (solid state disk) for the /vbox filesystem :

```
# zpool list
NAME      SIZE  USED  AVAIL    CAP  HEALTH  ALTROOT
rpool    416G  17.2G  399G     4%  ONLINE  -
sas      1.23T  418K  1.23T    0%  ONLINE  -
ssd      29.5G  3.46G  26.0G   11%  ONLINE  -

# zfs list ssd/vbox
NAME          USED  AVAIL  REFER  MOUNTPOINT
ssd/vbox     3.46G  25.6G  3.46G  /vbox
```

Using EMC NetWorker, we can easily recover the files from the server *opensolaris* and put them on the /vbox filesystem of the new server.

After restoring the vdi hard disk file, we simply create a new virtual machine,

```
# VBoxManage createvm --name win2003 --basefolder /vbox/Machines --ostype Windows2003
```

We attach the recovered win2003.vdi file as harddisk and the Windows 2003 system boots.

It still thinks it uses a 20GB IDE disk, but it now runs on the SSD (solid state disk) in the new server, and its network card (NIC) is now bridged to igb0 (an intel gigabit ethernet interface) instead of the bge0 card.

5 Image Backup using savenpc on the new X4275

After successfully restoring the image, we can now see how fast the savenpc method is on the new X4275 with its solid state disks and local SAS storage.

When we run a similar VBox group on the X4275 machine we get:

```
# time savegrp -l full VBox

real    0m36.772s
user    0m0.012s
sys     0m0.021s

# mminfo -avr client,name,totalsize,nfiles
vbox      /vbox                5158851044   14
vbox      index:vbox            10784       7
vbox      bootstrap            216552     268
```

So on the X4275 machine, with all its RAM and SSD and SAS disks, this backup runs (locally) very fast, we do backups to the six SAS disks with EMC NetWorker.

```
# mminfo -avr 'totalsize(3),nfiles,savetime(20),sscomp(20),name' -ot
5037 MB   14 02/ 5/10  3:11:48 PM 02/ 5/10  3:11:56 PM /vbox
   11 KB    7 02/ 5/10  3:11:56 PM 02/ 5/10  3:12:17 PM index:vbox
  212 KB  268 02/ 5/10  3:12:18 PM 02/ 5/10  3:12:18 PM bootstrap
```

It can be seen that the /vbox backup only takes 8 seconds.

During these 8 seconds we save 5037 MB, so this was at 629 MB/s.

Also during the backup we can observe these speeds using nsrwatch :

```
/B2D          file          B2D          writing at 680 MB/s, 4024 MB
```

During the backup the virtual machine is saved (and halted), and after the backup, i.e. after those 8 seconds in the case of the X4275 or after those 4-5 minutes in the case of the W2100z, we can check that it is running (and resumed) again.

6 Guest level backups with EMC NetWorker agent in the virtual machine

The disadvantage of image level backups is that we cannot easily restore individual files. So it is useful to run an agent in the virtual machine (the Windows 2003 machine).

The *postmark* tool can be used to generate small files.

Inside the Windows 2003 machine, we can open a command line and run :

```
C:\postmark_151w\release>postmark_151.exe
PostMark v1.51 : 8/14/01
pm>set location c:\small
pm>set size 100
pm>set number 1000000
pm>run
Creating files...
```

This will generate huge amounts of small files under the indicated directory.

We can define the Windows 2003 machine as a regular EMC NetWorker client, and make a backup of save set 'All'. We do this in a special group called VBoxGuest or VBoxGuestLevel.

In nsrwatch we see,

```
/B2D          file          B2D          writing at 2611 KB/s, 756 MB, 2 se
```

Or if we look a the number of files initially (before running postmark) :

```

# time savegrp -l 0 -v VBoxGuestLevel

real    7m30.415s
user    0m0.346s
sys     0m0.550s

# mminfo -avr client,name,totalsize,nfiles -qgroup=VBoxGuestLevel
  client      name                total nfiles
opensolaris  bootstrap                119812  185
opensolaris  index:smallfiles         955716  10
smallfiles   C:\                       511330412 5360
smallfiles   VSS ASR DISK:\          2865052  18

```

When we repeat the backup after creating some files with *postmark*, we get:

```

# mminfo -avr 'ssinsert(20),sscomp(20),totalsize(5),nfiles' -N "C:\\\"
  ss inserted          ss completed      total nfiles
01/ 1/03  5:03:49 AM 01/ 1/03  5:06:49 AM 764 MB 5351
01/ 1/03  6:12:21 AM 01/ 1/03  6:17:28 AM 853 MB 150712
01/ 2/03 12:32:40 AM 01/ 2/03 12:41:31 AM 1039 MB 552276

```

So initially for 5351 files, about 3 minutes of backup or 4 MB/s.

Next for about 150,000 files, about 5 minutes of backup or 2 MB/s.

And for about 500,000 files, about 9 minutes of backup or 1 MB/s.

The more small files we create, the slower the EMC NetWorker backup as it needs to update the index for individual file recovery, and it needs to read (traverse) the local filesystem which is getting slower and slower.

The hardware is important, because if we do a guest level backup on the X4275 with even more small files, we get:

```

# time savegrp -l full VBoxGuest

real    8m39.176s
user    0m0.013s
sys     0m0.022s

# mminfo -avr 'totalsize(3),nfiles'
  total nfiles
1341 MB 1100432

```

So we can see that on the virtual machine when it is running on the X4275 with its SSD disk, we get for a file level backup about 2.6 MB/s for a backup of 1,100,000 files (using VSS enabled backups in the virtual guest machine).

The big advantage of this type of backup is that file-level restore is possible.

Also, the new X4275 performs 3 times faster than the AMD machine for double the number of files (1,100,000 files compared to 500,000 files for the AMD machine).

7 Results on physical Windows 2003

We reinstalled the SUN X4275 and installed Microsoft Windows 2003 on it. Note that this requires the use of a special driver diskette (using the SUN StorageTek MSI device driver for the RAID controller, that is not supported by Windows 2003 by default).

On the RAID controller, we created two LUNs for the *postmark* benchmark : a drive *D* which is a RAID-0 stripe set with 6 SAS disks. The drive *E* is a SSD (solid state drive) volume.

Next, the results for a backup with VSS enabled are :

```
2/12/2010 2:01:25 PM 2/12/2010 2:05:04 PM 995663 508886880 D:\
2/12/2010 2:01:27 PM 2/12/2010 2:04:59 PM 976597 693861100 E:\
```

This means about 2.2 MB/s for the D SAS drive, and about 3.2 MB/s for the E SSD drive, for about 1 million files (with fewer files on the SSD disk).

For save operations set to *VSS:=off* we obtain :

```
2/12/2010 1:48:04 PM 2/12/2010 1:52:16 PM 995663 508886808 D:\
2/12/2010 1:48:03 PM 2/12/2010 1:52:21 PM 976597 693861028 E:\
```

In this case the D drive (SAS volume) worked at about 2 MB/s and the E drive at about 2.6 MB/s.

These numbers are very similar to the numbers in the case of the virtualised (VirtualBox) solution.

8 Conclusion

Thanks to virtualization, it is easier to perform BMR (bare metal recovery) as the virtualization layer can present on different hardwares the same virtual hardware to a virtual machine. With EMC NetWorker we can achieve high performance backups of both image backup (make a EMC NetWorker backup of the underlying disk images for the virtual machines) or work at filesystem level in the virtual machine (run the EMC NetWorker agent in the virtual machine).