



In This Issue

- Rumors, Rumors, Rumors
- Connecting the bits...
- Tips & Techniques

“a proven solution is ready to seamlessly take over from the older hardware.”

Rumors, Rumors, Rumors

The death of the VS has been talked about for years, and as Mark Twain would have said, it has been greatly exaggerated.

On July 21st Getronics announced that support for the aging Legacy Wang® VS™ systems would be officially winding down. At the same time they also announced that a proven solution is ready to seamlessly take over from the older hardware. The designated exclusive reseller of this solution is TransVirtual Systems.

As you read this, more than 70 clients, large and small, have been migrated to the new industry standard platforms. The new platforms are not only faster, they are physically smaller, greener and much more reliable than the legacy platforms. Three years and 70 systems later there has not been a single incident of downtime due to hardware failure.

All of us who have used Wang® products, be they hardware or software, know that Dr. Wang had a technological Midas touch. He produced products that were fast, reliable and elegant. Now thanks to state of the art hardware and TransVirtual Systems, the VS™ has a new lease on life and a bright future.

For more information, write to vsinfo@transvirt.com

Connecting the bits...

(Backup Strategies, Thomas Junker)

“New technologies offer greater flexibility and improved performance for all types of backups”

The New VS replaces the legacy VS hardware with a virtual VS machine running in a state of the art server under control of the Linux operating system. VS disk volumes are emulated by Linux files in byte-for-byte images of VS disks. Since the Virtual VS machine is seamlessly compatible with the legacy VS, all the usual backup utilities for VS volumes continue to work, but new devices and new means of backup are available, offering higher speed and greater reliability.

However, since the New VS runs in a Linux environment, this, introduces the more complex realities of a restore to bare metal,. The total replication of the environment in a replacement machine; the type of restore needed in the event of catastrophic failure. You not only need to back up the VS, but the Linux environment as well.

New technologies offer greater flexibility and improved performance for all types of backups, regular VS Volume backups as well as Linux backups. The new technologies include:

- 4mm tape (all generations of DDS including DAT72) (up to 36/72 GB)
- DLT and SDLT tape (up to 300/600 GB)
- LTO tape (up to 800/1600 GB)
- Disk-to-disk backup to RAID or enterprise storage
- Disk-to-disk backup to removable disk
- Integrated Virtual Tape (allows the VS to read and write virtual tapes stored on disk)
- Transfer of backup images to remote sites

“The first thing to realize is that the virtualized VS environment doesn’t change the rules of good practice in making backups, it only offers more options.”

You will note that the technologies go beyond the traditional legacy backup methodologies and expand into the world of disk copies and network backups.

VS Volume Backups

The first thing to realize is that the virtualized VS environment doesn’t change the rules of good practice in making backups, it only offers more options. We recommend that all VS sites continue to make the usual VS Volume backups as before. Only a VS backup/restore utility can list files contained on backup media and selectively restore single files or groups of files. The difference with the new VS, is that backups can be faster and more reliable.

DAT72

Most New VS systems have a built-in DAT72 4mm tape drive. It holds 36 GB of data uncompressed and up to 72 GB of compressed data. This is enough storage capacity to guarantee being able to back up the largest virtual VS volume presently supported by the VS Operating System. Maximum tape throughput is 3.2 MB/sec uncompressed and about twice that for fully compressible data. This translates to roughly 10 GB per hour native and 20 GB per hour compressed.

DLT/SDLT

Although Quantum, the owner of DLT technology, recently threw in the towel and joined the movement toward using LTO, DLT and SDLT drives are widely available and will likely continue to be so for some time.

In the late 1990s, Wang offered the DLT 8000 drive, primarily for larger systems. Unfortunately they had to recommend using an earlier generation of tape cartridge to limit throughput due to the low speed of VS SCSI. The New VS has no such limitations and commonly uses 160 MB/sec and 320 MB/sec SCSI, up to 64 times faster than legacy VS SCSI. DLT 8000 drives using DLT IV media can back up 40/80 GB at 6/12 MB/sec. This translates to roughly 20 to 40 GB per hour. As with DAT72, this capacity guarantees that even the largest 34.3 GB VS volume will fit on a single tape.

Some new VS systems have even used SDLT 320. It can back up 160/320 GB at 16/32 MB/sec. This translates to roughly 55 to 110 GB per hour.

LTO

Linear Tape Open (LTO, also called Ultrium) is the current tape technology, and LTO3 and LTO4 are the current generations. LTO3 can back up 400/800 GB at 40/80 MB/sec. This translates to 140 to 280 GB per hour. LTO4 can back up 800/1600 GB at 60/120 MB/sec. This translates to 200 to 400 GB per hour.

Why use a tape much larger than the largest VS Volume?

1. Speed of backup

An LTO4 tape drive can theoretically back up a 4.2 GB VS Volume in as little as 38 seconds. An LTO3 can do it in as little as 54 seconds. An SDLT 320 in 112 seconds. A DAT72 takes as little as 13 minutes to back up a 4.2 GB volume.

Although the tape media cost is higher as the capacity increases, the time saved doing backups may be well worth the higher media cost.

2. Multivolume backup

Although we don't recommend making multi-volume tape backups and VS BACKUP has some trouble restoring volumes from such tapes; products such as Back Burner and Ultrabac are used by some customers to make multivolume backups. This makes better use of higher capacity, higher cost tape media, but complicates the process of restoring data and may incur longer delays to position the tape. This is a customer policy decision.

"An LTO4 tape drive can theoretically back up a 4.2 GB VS Volume in as little as 38 seconds."

Disk-to-Disk

Disk-to-disk backups have long been used in VS systems. Since disks are generally faster than tapes, backup to disk can be used to reduce system downtime for backups, then tape backups are made from the disk backups after the system has been brought up again for production or nightly jobs.

The New VS can work with virtual VS Volumes anywhere the Linux system can see what looks to it like a disk drive. This includes external RAID, enterprise RAID and other forms of fast storage outside the New VS server. Backing up to dedicated external RAID safeguards the data outside the New VS server, while backing up to enterprise storage may safeguard the data in a different room or building and offers secondary means of backing up for offsite storage.

Removable Disk

Dell recently began to offer an innovative removable disk, the RD1000. It uses a ruggedized removable disk cartridge in sizes up to 500 GB. This offers a combination of fast disk-to-disk backup and removal of the backup disk for storage or transport to an offsite location.

“The New VS now has the capability to mount a disk file in place of a physical tape, both for writing and reading, seamlessly transparent to VS tape backup/restore software.”

The general strategy is to back up the logical disk containing Linux, the virtual VS software and the VS IPL Volume to tape or disk as a disk image, permitting the easiest and fastest restore to “bare metal.”

Integrated Virtual Tape [For a more detailed article about the IVT see the June newsletter.]

The New VS now has the capability to mount a disk file in place of a physical tape, both for writing and reading, seamlessly transparent to VS tape backup/restore software. This means that backups can be done to disk files that contain block-by-block images and control data describing tape blocks, and that such tapes can be read back by backup/restore software as if they were tapes.

Being disk files, such backup images may be stored in RAID or enterprise storage outside the New VS server, may be stored on removable disks, and may be transferred by network to other systems where they can be restored as if they were tapes.

Network Transfer of Backups

New VS systems are equipped with Gigabyte Ethernet and faster networking options are available. Anything that is in a file, such as virtual VS Volumes or backup tape images can be transferred over a network to other systems. Since a file is a file, volume or tape images can be sent to non-VS systems for storage, or of course can be sent to other New VS systems such as backup or Disaster Recovery (Business Continuity) systems. The main limitation in this is network speed.

While a 4.2 GB VS Volume image can be transferred via Gigabit Ethernet in as little as 42 seconds, a destination may be connected by slower networking. Some accommodation can be made for this by compressing images before transmission. Sparse VS volumes initially created to contain all zero bytes may compress as much as 90%, while moderately full volumes that have seen considerable use may still compress substantially.

TVS can help customers devise a VS Volume backup strategy to take advantage of available technology and serve their needs.

Linux Backups

As mentioned earlier, in the New VS there is another layer of operating system, the Linux shell. The Linux environment must also be backed up to permit restore to “bare metal” in the event of a catastrophic failure or disaster.

The general strategy is to back up the logical disk containing Linux, the virtual VS software and the VS IPL Volume to tape or disk as a disk image, permitting the easiest and fastest restore to “bare metal.” Secondary VS Volumes may be restored from regular VS backups, either tapes or tape images, or from disk volume images. Optionally, all secondary VS Volume images may be backed up and restored as a single Linux disk image, either using high-capacity tape technology or RAID or enterprise storage. This approach requires that Linux be configured to have two separate “disks” and file systems, one for Linux and the virtual VS software, the other to contain all secondary VS volumes. The VS IPL volume would be located in the Linux disk partition so that upon restore, a working virtual VS will be available to restore secondary VS volumes from backup tapes or tape images.

Linux itself needn't be backed up as frequently as VS volumes, but users

should not go too long between these Linux backups. In the event of a catastrophic failure, if a good Linux backup, with current IPL, volume is not available, the New VS probably will have to be completely reinstalled from scratch, taking much more time and effort. A good frequency of backup might be once a month unless USERLIST or other changes which affect the IPL volume are frequently updated, in which case more frequent backup may be advisable.

TVS can help customers devise a Linux backup strategy to take advantage of available technology and serve their needs.

It is also possible to use third party Linux backup packages, however, TVS has not certified any and makes no recommendations in this regard.

“TVS can help customers devise a Linux backup strategy to take advantage of available technology and serve their needs”

In summary, the TVS environment provides you with many different ways to backup to tape, most of which are much faster than the options on the legacy hardware. Some of these options are just "standard" tape backups but on faster devices, others are more creative, using the virtual tape and making a Linux backup of the entire environment. It's important to remember, that whatever you chose, a very important consideration will always be how to handle a "bare metal restore", a restore that takes place after you have lost all data on all drives, including the underlying Linux operating system. (Based on what we have seen, this is a highly unlikely scenario, but if you have a disaster recovery situation, this is what you will face.)

There is no single "perfect" solution, but one worth considering is to pick a backup media that is big enough to hold all of the information on all of the disks, and then use a Linux utility to back it all up. In the event of a disaster scenario, the Linux starting disk will make the recovery a simple process *if...* and *only if...* you were able to fit your entire backup onto a single volume.

We urge you to discuss this with us as you plan your system requirements. If you cover the "disaster scenario", properly configure your RAID drive and include a proper UPS, you will quite likely have years of totally uneventful service.

“One of the few unexpected side effects of running the VSOS in the Linux environment was a drift of the time of day clock”

Tips & Techniques

Time Drift

One of the few unexpected side effects of running the VSOS in the Linux environment was a drift of the time of day clock. The article explains why that happens and what short and long term actions are being taken to correct that drift.

The clock drift came as a surprise in large part because Linux has access to the Network Time Protocol (NTP). Systems using NTP for synchronization are capable of maintaining time of day clock accuracy to a sub-millisecond level. If

this is the case, why does the VS clock drift?

In order to understand this, we need to understand how the Linux clock operates and why it is different than the VS clock.

The Linux clock is pretty simple. You get the time from one or more dependable public time sources (e.g. tick.usno.navy.mil or time.nist.gov) via NTP, and you set the Linux clock accordingly. The NTP protocol is sophisticated enough to compensate for errors introduced such as network latency, etc. This mechanism is simple, but it has the potential for causing problems. Consider the simplest case, the "Daylight Savings Time" change.

In the fall, at 2AM, clocks move back to 1AM. In the spring, at 1AM, clocks jump forward to 2AM. Simple enough, but with a catch. In the fall, all times between 2AM and 1AM *occur twice*. Whereas in the spring, times between 1AM and 2AM *don't occur at all*. In other words, if you have a job scheduled to run at 1:30AM, in the fall, on the time change day, *it will run twice*, whereas in the spring, *it will not run at all*. While this symmetry pleases mathematicians, those of us who have to keep systems running are less amused.

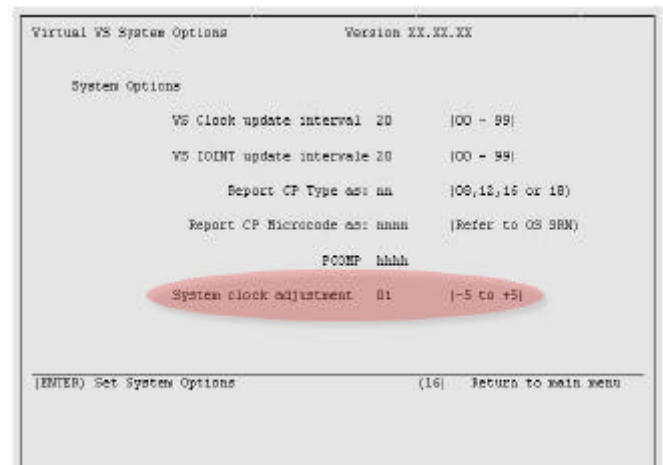
If the NTP update program is run frequently enough (e.g. hourly), most of the time, corrections introduced by NTP are quite small, usually no more than a few milliseconds.. Unfortunately, even those small changes can have similar effects *if they move the clock across a "second" boundary*. (e.g. 1:29:59 to 1:30:00) This raises the potential for the problem described above, a situation that is unacceptable for a system such as the VS.

The VS deals with this problem by creating a "monotonic inclusive" clock, a clock that moves in only one direction, and a clock where each "second" occurs once and only once. When you bring up the VS and set the time and date, VS stores that date and then starts an interval timer. On the CP18 processor (which is what the TVS product emulates) the interval timer updates 250,000 times per second, or every 4 microseconds. Every 250,000 ticks, the time of day clock is advanced one second.

This mechanism achieves the goal of a "monotonic inclusive" clock, but it is highly dependent on the accuracy of the 250,000 times per second update. Even a small error will accumulate and eventually produce a wall clock error of minutes. Unfortunately in the Linux/VS platform, the underlying clock speed of the Dell (or IBM) hardware doesn't multiply out evenly, so its not possible to have a tick precisely every 4 microseconds. Small errors creep in, and eventually, wall clock time drifts unacceptably.

The solution that we implemented is to allow a user modifiable clock adjustment. Based on the user supplied parameter, we occasionally *skip* an interval timer update or perform an *extra* interval timer update. The clock is still monotonic and inclusive of every instant, it's just that some instants are stretched a tiny bit and some are a bit shorter. The exact correction has to be determined empirically but an acceptable value can be quickly determined.

If you take a close look at



"The VS deals with this problem by creating a "monotonic inclusive" clock, a clock that moves in only one direction, and a clock where each "second" occurs once and only once."

"If you take a close look at the SCU System Options screen, you will see a new option, the "System Clock Adjustment".

“Future releases of the TVS product will make use of NTP to dynamically modify the clock adjustment value thus allowing the software to achieve near perfect time of day accuracy.”

the SCU System Options screen, you will see a new option, the “System Clock Adjustment”.

This allows you to modify the tick update rate of the VS clock. The bigger the number, the *more* often the update is skipped, and the wall clock “slows down”. The smaller the number, the *less* often the clock is skipped, and the wall clock becomes “faster”.

When determining the correct value, you can note the drift per hour (or day), change the value, and then note the *change* in slippage. When you finally get an acceptable slippage rate, you should shut down the VS, bring it back up, set the time accurately and go from there.

This method works quite well. The clock isn’t perfect, but the drift can be brought down to acceptable levels. Future releases of the TVS product will make use of NTP to dynamically modify the clock adjustment value thus allowing the software to achieve near perfect time of day accuracy.

<http://www.eecis.udel.edu/~mills/ntp/html/index.html>

http://en.wikipedia.org/wiki/Network_Time_Protocol

Do you have a question about system upgrades or the future of your legacy VS applications? Send it to us at vsinfo@transvirtualsystems.com.

Keeping Pace is a service of TransVirtual Systems, www.transvirtualsystems.com; contact us at vsinfo@transvirtualsystems.com. ©2008 TransVirtual Systems. All rights reserved.