



In This Issue

- Connecting the Bits
- A Conversation with...
- Quotes from the Field

“We found a base of legacy users that had a need for archival retrieval of data.”

“In the last twenty years, the cost of disk memory per byte has declined more than one thousand fold.”

Connecting the Bits

Integrated Virtual Tape

By Jim Donoghue, Senior Developer TVS

SUMMARY

The combination of reduced hardware costs and improved technology in a virtualized environment enabled us to implement the IVT (Integrated Virtual Tape). This eliminates one of the weakest links in a production environment, tape drives. All of this can be accomplished with no code changes.

INTRODUCTION

As we rolled out our new Intel based VS™ technology, one of the most frustrating problems we ran into was the plethora of antique peripherals in the hands of the existing user base. While this is a tribute to the quality and durability of the Wang hardware and software, trying to get data off a DDS-2 tape, or worse, a DDS-1 tape, became an endurance contest. Typically the installation of the TVS technology would be complete in a few hours, but the process of moving user data via these old tapes would take days. The problem wasn't the speed of the old tape drives (although that was a problem), it was that a combination of age and technology made them so unreliable.

Still, once this migration was done, the problem was history. Or so we thought.

We found a base of legacy users that had a need for archival retrieval of data. Legally mandated storage requirements for the data would range from a few years to more than a decade. The older tape technologies made this a problem. Far more than inconvenience, there were substantial risks that the data would be unretrievable.

THE WORLD CHANGES

Many of these tape drives were of mid 1980s or early 1990s vintage. The amount of data they could hold was impressive by the standards of their day, but seemed almost silly by modern standards. Today's wrist watch holds more memory than that.

In the last twenty years, the cost of disk memory per byte has declined *more than one thousand fold*. (Actually if you make an allowance for 20+ years of inflation, the decrease in cost is even more dramatic.) In addition, the advent of RAID controllers (see the FAQ section of the May issue of Keeping Pace) have greatly increased reliability despite the decrease in cost.

So the challenge we chose to accept was to leverage these lower cost larger

capacity disks of the 21st century to solve the data storage problems of the 1980s and 1990s.

Fortunately, in a virtualized environment, this isn't nearly as complicated as it sounds.

A VIRTUAL TAPE DRIVE

The Integrated Virtual Tape (IVT) device is a device that emulates a 2239V2 SCSI tape drive attached to the VS host. It is part of the 70V68(E) SCSI IOC emulation. IVT will function with fixed or variable block sizes.

“The IVT allows users to create, read, and write files on the Linux host as if they were SCSI tape devices connected to the VS.”

The IVT allows users to create, read, and write files on the Linux host as if they were SCSI tape devices connected to the VS. A utility provided as part of the IVT package allows tape images to be created from existing SCSI tape devices, outside of the VS environment. IVT stores its image files in a directory that is assigned to it from within the SCU. The names of the tape image files are taken from the six-character VS volume name. The tape volume name is obtained from the OS when a MOUNT request is issued. The CP intercepts the MOUNT SVC and extracts the volume name and device number. If the device requested is an IVT, the name is passed to the IVT handler for opening. If a file with the exact name is found, IVT opens the file. If no such file exists, IVT creates an empty file, which in turn behaves like a brand new tape volume.

Putting this in perspective, based on current technology, 500-600gb of disk dedicated to a tape library isn't unrealistic. (Bear in mind that whenever you use the phrase “based on current technology”, things often change before the toner finishes baking onto the page.) It is likely that 100-150 tapes could be stored on a machine with this configuration. Once the pain of the loading the tapes is over, tapes can be read out of the virtual tape library with no more tape handling and no program changes. (See the end of the article for more on tape capacity.)

“The IVT can also provide a way to ship data between systems without actually sending physical media.”

Then, using modern tape technology (e.g. LT05), this data can be backed up onto a small numbers of tapes for safety.

The IVT can also provide a way to ship data between systems without actually sending physical media. Once all of this happens in a virtualized environment it all becomes data on a disk. That data can be moved, copied or sent anywhere via ftp, or secure ftp, and then read in another city. All without any code changes. Of course if the IVT were told to write to *networked* disk storage, many possibilities open up.

A NOTE ON TAPE CAPACITY

The question that always comes up is “how much will this tape hold.” Those of us who don't have extensive experience with tapes don't understand why they can't get a simple answer. In brief, other than the physical length of the tape and actual writing density, there are two things that determine how much data a tape will hold.

- (1) Record length. When a tape drive writes data it leaves a gap between records. The amount of tape used by the record is proportional to the length of the record, but the size of the gap between records is fixed. The longer the data record, the more efficient the use of the tape.

- (2) Compression. Modern tape drives compress data as it is written. The effectiveness of compression algorithm depends on the redundancy of the data. Considering two extreme cases:
- Encrypted data, which ideally is as random as possible, barely compresses at all.
 - Upper case only English language text traditionally uses less than 40% out of a possible 256 combinations per byte, but is also subject to the well know distribution of letters. This kind of data may compress as much as 90-95%.

So the lack of an answer is not evasiveness, it is simply not clear cut.

? The letters A to Z, the numbers 0 through 9, space and some special characters (period, comma, dollar sign, etc.)

A Conversation with...

...Alexander Ramcke, IT Managers, Funk Gruppe, Germany

This month Keeping Pace sat down with Alexander Ramcke the IT Manager of the Funk Gruppe (www.funk-gruppe.com). Over 125 years in the business, Funk Gruppe is Germany's largest independent insurance broker, with complimentary offices in The Netherlands, Austria, Hungary, Poland, Romania, Italy and China. Offering a wide range of services, everything from risk assessment to public equity services, it's not hard to understand that this company needs a strong and reliable back end service to keep them organized and on schedule. The 1980s seemed to be an excellent vintage for the Wang, Funk Gruppe being no different, they've been using the hardware since 1984 and reached the end of the line with a VS18950.



"With more than two million lines of code and 800 users one of the biggest concerns, beyond testing, was transition."

"We started talking to tVS sometime in the beginning of 2005", Alex told us. Before then they'd researched various other solutions for the transition away from the original Wang technology. Alex noted that there were many different ways his company could migrate but between proprietary software, custom software compatibility issues and the literal time it would take, the options didn't really shine. "Altogether this would take about two years and would mean no development during this time."

With more than two million lines of code and 800 users one of the biggest concerns, beyond testing, was transition. With Funk unwilling to send their data out, TVS brought the testing hardware in. "Tom installed all the software through the internet." In October of 2005, A backup of their old VS was installed onto the TVS hardware on a Friday and on Monday users were allowed to work on the new system. None of the users were told about the switch in an effort by the IT department to see exactly how the experienced operator would react. They didn't receive any complaints during the test, however, Alex told Keeping Pace that reactions came in after the fact. "The users did notice the massive performance increase of the new system, but they didn't tell us about that. After the test period and the switch back to the old system, then we received a lot of calls and emails about the sudden bad performance."

What had originally been planned as a one day test had been extended to an entire week, only returning to the old system because legally they couldn't

operate full time on the new one yet.

When asked how the full transition went, Alex said, "We went live with the new VS on April 10th 2006. And we never thought of going back. Any problems were solved by tVS right away. It was a very busy time, but the results were worth it !"

"In the old system much of these processes used one of the most deeply flawed devices for sharing information, the human being."

The transition has allowed certain innovations in the processing, transfer and sharing of information. In the old system much of these processes used one of the most deeply flawed devices for sharing information, the human being. Because of the VS's ability to talk to its Linux host and thus the outside world, the extra and unreliable link in communication has been eliminated. It's like the difference between calling up a bookstore to find out an ISBN number and having it read to you so that you can transcribe it and going directly to a database such as amazon.com and finding the number yourself.

"We are now able to send a request in both directions and will (hopefully) get an answer." Now, the Wang can access everything in Linux like any other program, for example, an it can reach an Oracle database or access a web service being maintained by a government system. The first project put into production was to fill in the blanks from the German License Plate authority for information necessary to complete the insuring process without admin input, entirely within the two systems. Picture it as a series of automated information handoffs.

"We need to share data between Outlook/Exchange and the Wang about holidays, vacations, illness and so on".

Also, currently in development is a way for the Wang Calendar system to talk to the Outlook calendar system. As approximately 800 employees slowly transition from one program to the other, the problem of accurately sharing information came to the forefront. "We need to share data between Outlook/Exchange and the Wang about holidays, vacations, illness and so on. For this we needed to program a piece of software that would inform the Wang system about changes in Outlook." While not completed, Alex anticipates this bit of software will make it into production within days of the publication of this newsletter.

The Funk Gruppe's plans that take advantage of the way the VS operates were interesting to learn about. Increasing efficiency isn't always a matter of snatching up the newest and shiniest bit of computer paraphernalia. Upgrading a 20 year old computer and finding out the replacement is faster in almost all ways isn't a revelation in the computer industry. It's what someone does next that's fun to think about.

Keeping Pace would like to share your upgrade story with the VS community. Let us know about your experience at vsinfo@transvirtualsystems.com.

Quotes from the Field

This is a collection of letters and quotes from users who have recently completed the migration from their Legacy VS hardware to the TVS platform. We have not offered attribution for all of these remarks since some of these remarks were made "off the record".

"The new system is so fast that the field is complaining that they no longer have time to read the newspaper during data entry."

"We made a great cutover yesterday and today morning, I moved the old VS away and our users could work on the TVS platform without any problems."

—————
"I'm very glad and happy to tell you that everything is fixed at 100 %. We made a great cutover yesterday and today morning, I moved the old VS away and our users could work on the TVS platform without any problems."

"Roger made a great job here in Geneva. The beginning was very tense and I could tell you that I was a little stressed last week because we went down for an all day. [Editors note: This refers to the legacy VS, not the new system. As we see so often with computer equipment, the old system seemed *know* that it was about to be replaced and staged a last minute rebellion.] But each day I felt better and better."

Bernard Erny, Lorette SA

—————
"I used to enter a PACE query and get coffee while it was running. Now by the time I stand up the results are on the screen."

Captain Steve Hervey, Fort Lee PD

—————
The overnight jobs run so fast that we no longer need our third shift.

—————
I don't have time for coffee anymore. [Editor's note: This isn't a repetition, this keeps happening with each new install.]

—————
I am now able to be home for dinner instead of watching the end of week jobs run.

—————
The five hour batch job ran in 15 minutes.

—————
Why did we wait so long?

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.