

Minimum Dell PowerEdge configuration for the New VS

Version: 2009-0518-01 by Thomas Junker

Note: Compucom is the successor to Getronics North America, which was part of Getronics NV, the Dutch successor to Wang Global, which was a rename of Wang Laboratories. References to Compucom herein are to the holder of the VS intellectual property. TransVirtual has a contractual relationship with Compucom. Compucom supports the VS OS and all current VS system software. TransVirtual owns the VS virtualization technology and supports Linux and the virtual VS software.

Dell PowerEdge 2900 rackmount server

- Alternate packages are:
 - 1) PowerEdge 2900 tower
 - 2) PowerEdge 2950 (2U rackmount)
 - 3) PowerEdge 1950 (1U rackmount)¹
- Processor
 - Single or Dual 3.0 GHz Dual Core 5160 processor, 4 MB cache
 - Single or Dual 3.33 GHz Dual Core X5260 processor, 6 MB cache
 - Single or Dual 3.5 GHz Dual Core X5270 processor, 6 MB cache
 - Single or Dual 3.33 GHz Quad Core X5470 processor, 2x6 MB cache
 - Dual or Quad Core Xeon, 1333 MHz FSB and 667 MHz memory

Note that processor options are constantly changing. Those listed above are the options presently available in PowerEdge models that we use.

Note that the requirement is for one processor core for each New VS instance to be run concurrently. Most systems run just one instance, therefore a single Dual Core processor is sufficient: one core for the New VS and the other core for I/O threads. A two- or three-instance system would require four cores. A seven-instance system would require eight cores, which is presently the maximum available in a 2-slot PowerEdge.

- Embedded Dell Remote Access Card (DRAC 5) for TVS/Compucom support and optional customer admin
- 4 GB of physical memory configured as 2 GB mirror
- Dell PERC6/i Integrated RAID controller w/256 MB cache
 - Up to eight internal hot-swap SCSI or SAS disk drives may be configured as RAID
 - PowerEdge 1950 - Two drives in RAID1 mirror
 - PowerEdge 2950 - Two to six drives in RAID1 or RAID5
 - Limit five drives if internal DAT72 4mm tape drive specified

¹ 1U servers have a limit of two PCI slots

in PowerEdge 2950

- PowerEdge 2900 - Two to eight drives in RAID1 or RAID5
 - Our standard configuration is now at least a RAID1 of 73 GB or 146 GB, with RD1000 cartridges to match, and a secondary RAID1 or RAID5 array for any additional virtual VS volume storage, with RD1000 cartridges to match or as large as available. The primary RAID1 holds Linux, the virtual VS software and one or more VS IPL volume images. The secondary storage contains the remainder of VS volume images. The RAID1 in conjunction with the RD1000 removable disk provide the basis for bare metal backup and restore. Note that multiple PowerEdge systems for use in this manner must be precisely the same models and chipsets.
- Dell RD1000 removable disk bay with sufficient cartridges for bare metal backup
 - Dual Channel SCSI U160 or U320 card for tape, external SCSI

Optional:

- External SCSI and/or Fibre Channel for virtual VS storage in enterprise RAID or SAN as required
- Internal or external SCSI tape backup as required:
 - 1) DDS-4 4mm available for maximum Wang compatibility with DDS-1, DDS-2 and DDS-3 (not recommended due to poor product quality from Dell).
 - 2) DAT72 is a better choice than DDS-4 if the legacy VS has or can be upgraded to have current SCSI microcode.
 - 3) SDLT320 can read Wang DLT tapes if written on DLT IV media.
 - 4) LTO-3 is supported on New VS systems and is usually best housed in an external rackmount PowerVault. LTO-3 offers very high speed and capacity to 400-800 GB per tape.
 - 5) Combinations of the above can easily provide optimum combination of Wang tape compatibility and efficient backup going forward.
- Note that with our adoption of the RD1000 removable cartridge disk bay, the need for tape for backup/restore may in some cases be obviated. This is also influenced by the possible lack of bay space for the RD1000 and more than one tape drive. In some cases if tape is required, it may be housed externally.
- TransVirtual can also specify external RAID and/or removable disk solutions for dual-New-VS failover and/or disk-to-disk backup with easy removal to offsite storage.

Notes:

- 1) The 1U PE 1950 will only be approved for very small VS replacements where storage requirements are small, growth and potential performance upgrade are

extremely unlikely, and no PCI USIOC is required.

- 2) Physical memory can be traded off between VS memory and Unix/Linux memory. VS memory beyond the level that results in zero page-outs in the VS is unnecessary and only subtracts from Unix/Linux memory used for mapping and buffering file I/O. Unix/Linux use all of available memory for file mapping and buffering unless write-through is forced, resulting in significant storage subsystem throughput and efficiency. Physical memory should be sized at at least 512 MB more than what will be allocated to VS memory. The minimum of 2 GB is sufficient for most small and medium systems.
- 3) There may be a present limit to physical memory that Linux is capable of utilizing. If such a limit exists, it is no lower than 4 GB. We believe this is no longer the case but have not positively confirmed it. We do know that there are New VS systems in the field with 8, 12 and even 16 GB of memory, and they seem to be performing correctly.
- 4) Disk storage may be sized to the requirements of the particular system, but in no case shall be less than a pair of mirrored (RAID1) drives. Do not plan to use any existing Wang disk drives from the legacy system except possibly a late model Wang RAID. Be aware that even a late model Wang RAID will provide lower performance than the native PowerEdge RAID.
- 5) Advanced reliability options such as mirrored memory are available in the PowerEdge models. We now require mirrored memory in our minimum configuration.
- 6) Very critical applications should consider using one or the other of those options. Mirrored memory requires twice the physical memory that would otherwise be ordered, and it must be ordered specifically for use with the mirrored memory feature. The minimum for mirrored memory would be 4 GB, yielding 2 GB of redundant, fault-tolerant memory. Advanced disk RAID configurations can survive the loss of more than one disk drive at a time and should be considered for critical systems. RAID10 and RAID50 are combinations of RAID arrays using more disks but capable of tolerating certain multiple concurrent disk failures.
- 7) Any requirement to support legacy Wang "928" devices (coax or twisted pair such as VS workstations, printers, telecommunications interfaces, etc.) will generally require the larger PowerEdge 2900 model.

- 8) No software other than customer VS application software and what is provided by or approved in writing by TransVirtual and Compucom may be installed or run on the New VS platform. Just as the legacy VS hardware is dedicated to running VS software, so is the New VS. Any unauthorized software found on a New VS system will disqualify it for TVS and Compucom software support until the system is restored to its original condition. This provision is specifically aimed at preventing the loading or use of non-VS Linux software packages that could disrupt the operation of the New VS.

No Linux "updates" may be performed by the customer because Linux vendor updates have been shown to disable critical configuration of subsystems within Linux and make the New VS fail to operate. The New VS server is used exclusively for running the New VS software and the customer's VS software. In this regard, Linux should be viewed as an appliance provided and supported by TVS, not updated or maintained by the customer.

- 9) Support is conditioned on access to the server by TVS. We install Linux and our software; the customer does not. If the customer ignores these conditions, any remedies will be separately billable. Failure to provide access by TVS at least at times when support is requested will make support unavailable without a billable site visit.
- 10) Transfer of legacy VS volumes into the New VS has proven to be the most problematic aspect of New VS setup and install. This is because many legacy VS systems are still using thoroughly obsolete tape technologies such as DAT, DDS-1 and DDS-2. Those tape generations are not directly readable on the New VS using current 4mm DAT72 tape drives.

Be advised that the very best solution to this is to upgrade the legacy VS with a DAT72 tape drive, which is feasible if the legacy system has current SCSI microcode and an available SCSI ID for an additional 4mm drive or no necessity to retain the existing obsolete drive to read archived tapes. A DAT72 tape drive will cost approximately US\$500-1,000 plus installation. Sites with Compucom VS maintenance contracts can have this done for a per-incident fee. Other sites can have this done by a reputable third-party, or can pay Compucom for an out-of-contract incident, or may be able to do it themselves. The DAT72 drive will more than pay for itself in the ease and speed of data transfer. The drive is current technology and much, much faster than any earlier 4mm technology. The drive can also be recycled for use on the New VS or other system after the legacy system has been retired. In some cases we may have an external loaner tape drive available for retrofit to the legacy system.

The next best alternative is to retrofit the legacy VS with a Sony SDT-9000 DDS-3 4mm tape drive, which is supported by older SCSI microcode as long as the drive is optioned to identify itself correctly. We can provide Sony SDT-9000 drives, as can Compucom.

IMPORTANT NOTE: The most common source of SCSI bus problems is loss or lack of Terminator Power.

Any customer or third party maintainer doing any work on connected SCSI devices ***MUST OBSERVE*** the requirement that power must be ***completely removed*** before any SCSI cable is disconnected ***or*** reconnected. On large VS systems this means bringing the system down to ***DC power off***. On large VS systems having a SCSI converter, it also means dropping power to the converter. On VS5000/6000 systems it further means ***disconnecting the power cord from the system unit***, since the VS5000/6000 ***RCU is always powered*** even when the system is powered off via the keyswitch.

The reason for removing all power is that there is a risk of blowing the VS SCSI Terminator Power Fuse in the IOC if any cable is connected or disconnected while power is up. This becomes even more complicated if there is no LED-equipped external SCSI terminator on the bus to show the presence or loss of Terminator Power. If the fuse blows and Terminator Power disappears, all manner of bizarre and variable problems may ensue, and may result in failure of all attempts to port data out of the legacy system to the New VS. It is even possible for lack of Terminator Power to allow the SCSI bus to operate, but it may suffer variable failures and will certainly react badly to any change in cabling done later.

Remember, any difficulty in moving data from the legacy system to the New VS is ultimately the responsibility of the customer, who has usually failed to keep the legacy system up to date and under maintenance. It may involve some modest expense and if our requirements are not met or our instructions not followed, it may involve billable time to straighten it out.

11) FYI: Our standard setup methodology.

For U.S. And Canada systems TVS procures the server, has it shipped to Houston, TX, Tewksbury, MA or Staten Island, NY, where it is fully tested, loaded, and reshipped to the customer site.

It is standard that we require a good, complete backup of the legacy IPL volume and full information about the other volumes present on the legacy system. We create and restore the IPL volume in the New VS and make necessary configuration changes. We create the other volumes, leaving them ready to be restored by the customer or during onsite install. If requested, we can restore further volumes, but this is only reasonable for static volumes or entirely static archival retrieval-only VS systems.

For overseas systems other than Canadian, we do not generally ship servers from the U.S. for a variety of reasons, not the least of which is the likelihood of an

exclusive right to import such systems held by the server vendor's local subsidiary or dealer. In these cases we provide the customer with precise specifications for the server, to be procured locally by the customer or by the customer's IT vendor. We provide instructions for configuring the remote access facility in the server, and for forwarding ports in the customer's firewall. Nothing further can happen until those two things have been done correctly and we are able to connect to the server.

Once we can connect to the server we verify that the server in fact complies with our minimum requirements. If it doesn't, the customer or his IT vendor must remedy the noncompliance before setup and install can continue. We configure anything that was not correctly configured by the server vendor, such as mirroring the memory and setting up the RAID arrays and partitions.

We remotely install and configure Linux. In no cases shall the customer install Linux.

We remotely install and configure the New VS virtual VS software and related servers the customer may require, such as virtual LS gateways, async, WSN, RSF, and Integrated Virtual Tape. In no cases shall the customer install the New VS software.

We create and configure all the virtual disk volumes needed to duplicate the legacy system and any additional volumes that may be needed.

We restore and configure the IPL volume and create the remaining volumes for restore either by us or by the customer. We either assist the customer in restoring the other volumes or we do it remotely with the customer inserting tapes as necessary.

When we release the New VS to the customer for test and acceptance, it will be a fully functioning, IPL-able VS containing all the necessary customer configuration, servers such as LS gateway(s), async, WSN, RSF, etc., and will be ready for use. If the customer wishes to do some of the volume restore work, we will consult and assist. Usually acceptance is expected shortly after we have released the completed system to the customer.

12) Onsite Installation

We can send a CE to perform onsite installation in almost any country. This is separately billable and should be negotiated around the time of order placement and down payment.

In most cases we will still insist on preparing the server at TVS or Compucom except for overseas sites, where we will still generally insist on doing remote setup and configuration. The onsite install is mainly to handle and resolve any

issues that arise due to site conditions, but in some cases can consist of complete setup and installation to a bare server.
