

White Paper

CentricStor Virtual Tape: the Swiss Army Knife for data protection

September 2, 2008

**Josh
Krischer**

White Paper

CentricStor Virtual Tape: the Swiss Army Knife for data protection

Table of Contents

| | |
|---|---|
| Table of Contents | 1 |
| Introduction, Short Description | 3 |
| Operating Systems | 4 |
| Features and Innovations | 5 |
| High availability and Disaster Resilience | 6 |
| Serviceability | 6 |
| Security and compliance | 7 |
| Future developments | 7 |
| References | 7 |
| Market shares, Market outlook | 8 |
| Conclusions | 8 |

September 2008

CentricStor Virtual Tape: the Swiss Army Knife for data protection

Josh Krischer

Josh Krischer is an expert IT advisor with 37 years of experience in high-end computing, storage, disaster recovery, and data centre consolidation. Currently working as an independent analyst at Krischer & Associates GmbH, he was formerly a Research Vice President at Gartner, covering enterprise servers and storage from 1998 until 2007. During his career at Gartner he was responsible for high-end storage-subsystems, co-authored disks and virtual tape Magic Quadrants and spoke on these topics at a multitude of worldwide IT events, including Gartner conferences and symposia, industry and educational conferences, and major vendor events.

For more than a decade, virtual tape libraries have been popular in the mainframe domain, but they have been rarely used in Unix and Windows environments even though it is impossible to find mainframe-only infrastructures today. The first virtual tape, the VTS, was launched by IBM in 1996 as a part of the "Seascape" architecture, which, at the time, was based on standard components such as the RS6000. It was followed two years later by StorageTek's VSM; both of these products supported MVS only. Currently, IBM's TS7700 support mainframes and the TS75x0 support Windows and major Unix operating systems. Sun StorageTek VSM continues to support z/OS mainframes only. In 1999 Fujitsu Siemens Computers (FSC) launched its CentricStor virtual tape solution, which became the most open, flexible virtual tape library supporting most existing mainframe and major Unix and Windows operating systems and major tape libraries. CentricStor was conceptually designed as an appliance to be used with all types of tape software, servers, tape libraries and tape technology as well. Initially it was released by FSC to support the BS2000/OSD, MVS mainframe platforms and ADIC tape libraries; however, over the years with constant certifications for major operating systems, backup applications, tape library systems and tape drive technologies it evolved into the most versatile VTL appliance, which can be deployed in almost any environment while protecting previous investments in tape automation. The recent CentricStor Virtual Tape Version 4.0 was announced in autumn 2007.

Short Description

CentricStor is currently available in seven Virtual Tape Appliance models (from the entry model VTA 500 to the VTA 5000 top model) providing sufficient granularity to fit any organization's size requirements. At the low end the product series also includes the CentricStor Virtual Tape Controller (VTC) as an inexpensive platform for the entry into virtual tape technology and the CentricStor Smart Backup Unit (SBU) which is offered as a complete solution package.

The modular CentricStor Virtual Tape Appliance models consist of independent building blocks: A front-end handles diverse heterogeneous host server connections, and a back-end which handles the communication and data movements to the physical tape libraries. A high

scalable internal RAID subsystem serves as cache for the backup data written to physical tape and also for temporary or permanent virtual tape volume storage.

Interfaces: CentricStor supports a huge variety of different platforms, currently up to 32 host channels of different types, including standard Fibre Channel (FC), ESCON and FICON¹ for connection to IBM mainframes, and Type S, which is the Fujitsu Siemens Computers mainframe fibre-optic connection. Different channel types can be mixed in a single system.

Operating Systems:

Mainframes

- z/OS (OS/390)
- i5/OS (OS/400)
- MSP
- z/VM – z/VSE
- VME (ICL)
- XSP
- BS2000/OSD
- GCOS 8

Open systems

- Solaris
- LINUX (SUSE, Red Hat, z/Linux)
- AIX
- HP-UX
- Reliant UNIX
- Windows NT, 2000, 2003, 2008

As opposed to the Sun StorageTek's VSM, CentricStor works as a host independent appliance which does not have any virtual tape software component on the host and is therefore transparent to the users. CentricStor's True Tape Virtualization (TTV) provides a transparent virtualization layer which looks like a tape library to the user with up to 1024 virtual tapes drives. It supports up to 1,500,000 virtual tape volumes with variable volume sizes and allocates resources to servers on a just-in-time basis, similar to the "thin provisioning" concept.

The broad platform support is one of the strongest advantages of CentricStor. In addition to its proven record of supporting non-mainframe operating systems, it is the only virtual tape system supporting almost all mainframes (IBM and other vendors) and even IBM's System i. Consequently CentricStor supports the major mainframe and non-mainframe backup solutions. It is also the only such system supporting heterogeneous platforms, which translates to less Capital Expenditures (CapEx) for procurement of multiple subsystems, lower Operational Expenditures (OpEx) due to simplified operation, and lower energy consumption.

Tapes: As an integrated and automated disk-to-disk-to-tape solution, the appliance autonomously offloads (based on attributes and usage) the tape volumes from its cache to real tapes. CentricStor's True Tape Virtualization (TTV) has full control of the attached tape libraries and tape drives. Unlike IBM and Sun StorageTek, which only support their own libraries and drives, CentricStor supports tape libraries from IBM, Sun StorageTek, Quantum/ADIC and others, also mixed, and up to 112 physical tape drives with diverse technologies which can also be mixed. The back-end tape technology e.g. LTO4, TS1120 or T10000 is completely independent from the emulation at the front-end. By supporting almost all enterprise tape libraries and all standard tape formats, CentricStor ensures investment protection and avoids additional capital and operational expenditures.

¹ FICON channels are not supported by the models VTA-500 and VTA-1000

Tape data “healthcare”, automatic media cloning and data migration to new tape technology e.g. from one LTO generation to the next to achieve double tape library capacity is done automatically, which prevents host data modifications and follows compliance regulations for long term archiving requirements. These migrations cycles on the back-end have minimal impact on performance due to CentricStor’s large cache capacity.

The user benefits from:

- Fewer physical tape drives
- Fewer slots in the library
- Better usage of tape media (in particular for mainframes)
- Improved performance (disk-to-disk data transfer)
- No waits for a “free” drive due to a large pool of virtual tape drives.
- Ease of migration to new tape technology
- Better security
- Ability for application to support “old” tape formats on state-of-the art tapes.
- Lower storage management costs

Cache: Usually, the disks in mainframe VTLs merely function as a cache buffer between the host and the library, retaining the tape data during usage. One of the major differences between CentricStor and TS7700 or VSM is the highly scalable RAID disk cache which can include FC (up to 228 TByte) or SATA (up to 1PetaByte) disk drives. The data written to the disk cache in tape format is already compressed by the front-end modules, which allows the user to keep even more data on the disk. In addition the compression mode of the attached tape drives can be switched on.

Field experience shows that a majority of CentricStor customers use a larger cache size to hold the backup data from the previous 3-5 days inside the cache despite the fact that the files have already been moved to physical tape. This allows for much faster restoration of the latest data without the need to load and read from tape. This tape data on disk offers additional protection in the event of physical tape failures or library maintenance. In addition to enhanced data availability, this feature can translate into lower OpEx because users can choose less expensive maintenance contracts with longer response times for the tape library.

Features and Innovations

CentricStor automatically manages tape data based on user attributes (such as Dual Save, Cache Resident, encryption, etc.) corresponding to different SLA requirements. Usage patterns are collected and analyzed in order to optimize the utilization of the CentricStor resources. These features and the ability to use different types of disks (FC for performance and SATA for capacity) allow users to build internal multi-tiered storage, or, in other words, ILM “in a box”. CentricStor’s large cache capability enables several very innovative and unique functions as well.

Innovations announced with the CentricStor Virtual Tape Version 4.0 in December 2007 are “Triple Save”, as well as “Data-on-Disk-only”. Dual Save (available before) writes the data on two tapes in local or dispersed locations, Triple Save enhances this capability straight forward. The main usage scenario is to have two images online for highest data availability and to export one additional image for disaster recovery purposes. The Data-on-Disk-only feature retains the data on the disk, in effect, using the VTL as a disk library. Another feature which was deployed before Version 4 is “Cache Residency”, which allows locking files inside the cache after having them de-staged automatically to tape. “Save Delay” allows defining a

time delay for de-staging files thus saving resources if the virtual tapes are overwritten / appended frequently. Especially in mainframe environments, the high cache capacity and intelligent management of tape data in-cache accelerates restore cycles, optimizes the utilization of both tape drives and cartridges, and thus saves on purchasing additional physical tape drives and tape cartridges.

High availability and Disaster Resilience

FSC implemented many availability and redundancy features to ensure CentricStor's place in enterprise infrastructures. All components (building blocks) and power supplies are redundant to guarantee local availability.

The Dual Save and Triple Save options mentioned above allow local and remote copies on tape. This feature protects against failures or loss of a single physical cartridge, undetected failures of a tape drive during writing, and against local disasters concerning the loss of one of the tape libraries. By using the Dual / Triple Save option, the CentricStor produces the redundant copy without any host overhead. These features allow autonomous actions in the event of detected media errors or equipment failure redirecting CentricStor to read from the redundant copy and later replacing the failed copy with a new one.

Tape libraries can be located hundreds of kilometres from the CentricStor system. Writing to the remote tape drives is achieved by interconnecting special network components in between data centres and utilising the tape pipeline mechanism. (Reference: Postbank Systems: Two data centres in Bonn and Frankfurt with distant remote tape libraries for one another)

Partial or complete cache mirroring synchronously mirrors virtual volumes to remote locations over FC, which, in the event of local outages, allows immediate restores from the secondary cache (Reference: Swisscom with two data centers in Bern and Luzern – cache mirror over 120km).

All meta-data remains inside CentricStor. It is also internally mirrored, as well as written to the tapes. This feature ensures that even in the event of a complete system loss, the database can be rebuilt solely from tape. Additionally there is a CentricStor Add-On tool available which can manage backup and recovery for the CentricStor internal Data Base. Due to the full compatibility among all models, even a single CentricStor VTC model can be used to recover data from vault tapes.

Serviceability

FSC service engineers grew up in mainframe cultures and have collected substantial knowledge and experience in all supported platforms. Sophisticated proactive services such as “Live Monitoring” and “Health Check” help to increase the availability of CentricStor and maintain zero downtime: If an error is detected, rule-based actions can be implemented automatically. Non-disruptive planned maintenance and upgrades are supported by the system.

With the Solution Contract, FSC extended existing maintenance services by offering service management for a complete data protection solution including servers and tape libraries to speed up corrective maintenance. For customers, this means they only need to make one call to solve a problem: Fujitsu Siemens Computers takes responsibility for the uninterrupted operation of all the maintenance in a customer's system – from taking the call through to identifying the cause, possibly also involving other service providers, in order to solve the problem and avoid “finger pointing”.

Security and compliance

Authorization and Administration: CentricStor supports the strictest security requirements and is capable of passing comprehensive security audits. It allocates different user rules to control access level, and extensively logs every action.

Encryption: Decru and NeoScale encryption appliances are certified for system-independent encryption of physical tapes at the back-end, the support of tape drive embedded encryption is planned. Especially for mainframe customers this is a less expensive alternative to the integration of encryption machines in between the FICON channel and the tape controller, thus reducing investment and operation costs.

Future developments

Data-on-Disk-only will allow the installation of embedded de-duplication, which may dramatically reduce the required capacity. In anticipation of the accelerated growth of unstructured file-based data, FSC is currently developing a NAS front-end connection. These two enhancements will be very important steps to entrench CentricStor's position as an ILM "in a box". Another development in progress is towards a "cascading" architecture in order to fulfill requirements for a very long distance disaster recovery infrastructure.

References

At the beginning of 2008, FSC had more than 600 systems installed at more than 300 customers (EMEA, APAC, Japan, US, South America) storing a total of more than 260 PBytes of data. 15 percent are supporting mainframe environments, 50 percent other platforms, and 35 percent are installed in mixed environments. Some of the many references:

- **Postbank Bonn, Frankfurt** - z/OS, AIX, HP-UX, Windows NT, Novell and Linux servers as well. TSM (Tivoli Storage Manager).
 - Bonn** data centre: 4x CentricStor VTA with cache mirror and automatic failover, 6x PowderHorn 9310 tape libraries, Dual Save over two redundant DWDM disaster recovery connections to the Frankfurt backup data centre.
 - Frankfurt** backup data centre: 1x CentricStor VTA, 1x PowderHorn 9310 tape library.
- **Sparda-Datenverarbeitung eG, Germany** - CentricStor VTA 4000 split system configuration in two geographically dispersed data centres. Supports z/OS, z/Linux, HP-UX, AIX, Microsoft Windows and TSM. Dual Save to two SL 8500 tape libraries with T10000 tape drives.
- **AOK Health Insurance Bremen, Germany** - Consolidated tape storage with CentricStor VTA for z/OS, BS2000, Solaris, AIX, Windows, and two Scalar i2000 tape libraries.
- **Informatikzentrum Niedersachsen, Germany** - CentricStor Virtual Tape Appliance in two data centers, BS2000/OSD, OS/390 and z/OS mainframes as well as Open Systems with Tivoli Storage Manager (TSM), 2 x ADIC Scalar 10k
- **Orange County (Florida) Public Schools, USA** – CentricStor VTA supporting z/OS, AIX and Windows. Backup applications include TLMS for the mainframe, Tivoli Storage Manager for AIX, and Veritas (Symantec) Backup Exec for Windows. StorageTek L700 tape library with LTO Ultrium tape drives.

Market shares, Market outlook

According to different market research companies, at the end of 2007 there were about 10,000 Virtual tape Libraries installed worldwide, the vast majority in z/OS environments. Sun StorageTek's VSM only supports z/OS and IBM has very few non-mainframe installations. FSC was the latest to enter this market, and despite the first MVS installation's completion in 1999, it only managed to win 5 percent of the market based on installed units count, mainly due to a lack of powerful distribution channels. It is important to note that it's 5 percent of the total market represent almost 100% of the non z/OS market. Since 2003, Fujitsu has been marketing an OEM version of CentricStor as Eternus VT600 for their MSP mainframes in Japan. For the South East Asia region, Fujitsu launched the FSC-branded CentricStor in 2004.

IBM's TS7700 Virtualization Engine was launched in 2007 with a completely different architecture based on nodes and cluster; it currently supports one node per cluster and three nodes cluster within a grid. A single cluster is able to support up to 16 physical and 256 logical tape drives, 1,000,000 logical volumes and up to 12 FICON channels. The maximum write data sustain rate is about 450 MBps per node which is substantially less than the CentricStor (maximum 4,000 MBps). It supports all of IBM's recent mainframe operating systems, but not UNIX, Windows or System i. The TS75x0 supports Windows and major UNIX operating systems.

Sun StorageTek's VSM5 differs in architecture from CentricStor and TS7700 because the virtual tape logic is resident on the z/OS mainframe. This is one of the reasons that StorageTek has not developed support for other systems. The disk cache is provided by a StorageTek SVA, which used to be the company's high-end storage but is no longer developed because Sun is reseller of Hitachi USP high-end storage. VSM users are concerned about the future of their product and are evaluating other options for their aging VSMs. The CentricStor platform can be an ideal replacement for VSM, particularly in a heterogeneous environment.

CentricStor is the perfect virtual tape solution for hybrid mainframes using z/OS and Linux on System z or any installation using other major operating systems in addition to z/OS. Another opportunity for CentricStor is in the non-z/OS market, which includes non-IBM mainframes, System i, UNIX, and Windows. I predict much better acceptance of virtual tape libraries outside the IBM Mainframe world as technology with a proven track record for IBM mainframes should be suitable for other platforms and heterogeneous infrastructures as well. CentricStor is the only virtual tape library supporting these systems with a proven record and solid references.

Conclusions

Fujitsu Siemens Computers entered the mainframe virtual tape market after IBM and StorageTek, but today it has the most promising technology of all. As an analogy, I could compare CentricStor to a Swiss Army Knife of virtual tapes libraries. It is the most flexible virtual tape solution available today, the richest in features, the fastest, and the only one supporting every major operating system and every enterprise type tape technology. It is the only such subsystem supporting heterogeneous platforms which fits the current and future data centre infrastructures. CentricStor deployments may reduce operating expenses, increase security, and improve service levels. Organizations that do not own virtual tape or consider replacing older subsystems should put the CentricStor on their shortlist of products to evaluate