

Postbank Systems AG relies on CentricStor virtual tape technology for backups

Issue August 2007

Pages 4

A fail-safe, multi-site backup solution

Postbank Systems AG, Deutsche Postbank AG's IT services arm, deployed its first two production CentricStor solutions in early 2003 to backup data of mainframe systems at its data center in Bonn. The company now has five CentricStor Virtual Tape Appliances in operation, providing reliable backup and recovery services for the entire IT environment, including open systems. It's most recent project centered on using the storage solution to backup systems at its Frankfurt data center, too, and on deploying a disaster recovery link back to the data center in Bonn.

With more than 14.6 million customers, around 22,000 employees, and total assets of around EUR190 billion, Postbank Group is one of the largest retail banks in Germany. Postbank Systems AG, a wholly owned Postbank subsidiary, is responsible for planning, developing, implementing and operating all of the Group's IT systems.

The Databases and Storage department at Postbank Systems' Operations division is in charge of designing and implementing the DB2 and Oracle database systems used to plan, provision and operate disk storage systems and data backup hardware and software for the Group.

The problem: Plenty of capacity for little content

When in the late 1990s Postbank consolidated its six geographically dispersed data centers and centralized its IT in its Bonn data center, decision-makers and IT professionals at Postbank Systems' Databases and Storage unit saw a problem beginning to emerge. Until then, data from classic mainframe applications had been backed up on tape cartridges with a capacity of 800MB. The entire tape archive comprised 12 PowderHorn 9310 tape libraries with around 66,000 tape cartridges – too many, particularly as the quantity of data to be backed up was expected to grow more rapidly over time, creating space problems in the long run. Furthermore, just a small percentage of the tapes' capacity was actually being used during backups, because mainframe applications were only writing small data volumes to the tapes. Essentially, there was massive surplus tape storage capacity for relatively little content.



Postbank headquarters in Bonn

The solution: Virtualization

In early 2000, Postbank Systems' specialists realized that the capacity could be used much more efficiently if the data to be backed up was initially saved to hard disk and then written to tape in its entirety. Based on a five-year planning horizon, they found that deploying virtual tape technology not only represented a future-proof investment for the IT services unit, it marked a significant technological breakthrough, too.

But the requirements the solution had to meet were tough. The kind of tape virtualization system they were looking for had to be server hardware and application-agnostic and capable of interoperating with the existing StorageTek PowderHorn tape libraries. In addition to BS2000/OSD systems, it had to support a new IBM z/OS mainframe environment that was under planning as well as open systems running Unix, Novell and Windows server operating systems. Given the anticipated growth in the mainframe environment's scope and data volumes, it was also essential that the front end, back end and cache be capable of scaling flexibly. Oliver Heuser, platform manager for databases and storage, recalls, "Following a thorough review of the available options, it turned out that the only solution capable of meeting all our requirements was the CentricStor Virtual Tape Appliance from Fujitsu Siemens Computers."

The decision: CentricStor

CentricStor is a virtual tape solution that provides fully integrated "disk-to-disk-to-tape" backup capabilities. This means that hosts write data to a large number of virtual tape drives; the data is stored in an integrated disk cache and then written by CentricStor completely autonomously, in a rules-based process, to physical tape libraries attached to the CentricStor system. This approach utilizes almost all of the physically available tape capacity by combining several smaller logical volumes in CentricStor to create new volumes in a process known as volume stacking. Besides using almost all of the tapes' available physical capacity, it also greatly reduces the frequency of errors because data is streamed to the tape drives continuously. From a server's perspective, CentricStor and its virtual drives behave just like physical drives. And this means that existing backup and restore jobs can continue to be used without modification.

To ensure that recent data can be restored rapidly if need be, data is held in the cache for some time after it has been written to tape. This also makes it possible to carry out maintenance on the physical tape libraries while still continuing to fulfill restore requests. In addition, CentricStor regularly monitors the quality of the tape media and auto

matically copies data to new media when tapes produce errors. CentricStor also has the capability to automatically refresh long-term storage tapes.

Following intensive testing, the first two production CentricStor systems, located in separate fire cells, began backing up BS2000 systems in the first quarter of 2003. In each backup cell, the mainframes were connected to the CentricStor front end, each with 128 virtual drives. At the back end, the data was replicated simultaneously in dual save mode to two tape libraries, again located in separate cells. Just 16 physical tape drives were needed instead of 120 previously. The switch to the new technology also greatly reduced the number of tape cartridges needed – to fewer than 4,000. For Postbank it was a major benefit to be able to continue using its existing StorageTek PowderHorn tape libraries.

"This is where CentricStor's open concept comes into its own," says Guido Spallek, responsible for mainframe backup management. "We're free to choose how we configure the tape robot and the drives." To meet more rigorous backup and restore requirements associated with IBM DB2 and SAP systems, the company purchased a third CentricStor system with a much larger cache and new tape drive technology for the back end.

IBM compatibility remained an issue for CentricStor. In a departure from previous practice, the Postbank data center chose to equip its z/OS servers with FICON channel connections rather than ESCON because of certain technical advantages. Because Postbank Systems was eager to continue enjoying the benefits of CentricStor, this meant bringing CentricStor into line with the change in technology. "If that hadn't worked out, we would have had to rethink our backup strategy," Guido Spallek recalls. But this challenge, too, was mastered successfully and in 2005, a fourth CentricStor system, this time equipped with FICON channels, was installed. By this time, CentricStor's virtualization technology was so advanced that Postbank Systems was able to run backups with cache mirroring and automatic failover spanning two separate locations for enhanced resilience.



Cash service for Postbank customers

CentricStor for open system environments

Based on its positive experiences with CentricStor in a heterogeneous mainframe environment, Postbank Systems in 2006 decided to hook up its open systems too. This was an option that had already been trialed prior to installing the first CentricStor systems. The challenges involved in integrating an IT environment with scores of AIX, HP-UX, Windows NT, Novell and Linux servers proved to be wide-ranging and diverse. Initially, data from three cells was backed up natively to tape via the company-wide SAN using the strategic storage management solution Tivoli Storage Management (TSM). However, the 20 physical drives were busy running backups 24 hours a day, leaving no time for restore operations. At the same time, the TSM and conventional disk storage pool setup reached the limits of its capabilities. The copy storage pool, with its six million additional entries, simply proved too much for the TSM database to handle. In addition, the 4TB Oracle database, used for SAP/R3 business warehousing, was too large to be buffered by the disk storage pool. And TSM had regular tape management tasks to take care of on top of this.

Upgrading the TSM infrastructure by expanding the disk storage pool and purchasing additional physical tape drives would have radically increased costs. High-speed backup

to SATA hard drives rather than tape, decision-makers felt, was not sufficiently reliable, particularly when you took into account the added benefits that CentricStor had to offer, with its dual save capability, cache mirroring and high availability. Moreover, Postbank Systems preferred to – and had to – continue backing up to tape. Not least in light of the statutory data retention period of ten years, backup to hard drives would have involved considerable outlays for procurement, power consumption, and room cooling systems. These were sound reasons for the IT services unit to stay with tape, a reliable and inexpensive storage medium, and at the same time to leverage its positive experiences with CentricStor and mainframe backups to create a resilient backup infrastructure for its open systems environments.

Here, too, the CentricStor solution's openness paid off: "We were able to connect our TSM clients to the virtual tape drives without any problem," says Christoph Zink, the TSM environment's supervisor, impressed by the storage solution's strengths. "Backup time is down to just 5 hours from 24 previously, which means we no longer run into time clashes when we need to restore data." TSM now backs up the Oracle database directly to eight CentricStor virtual tape drives. Now that the cache has been expanded to around 10TB, database backups from the past three days are held at


the ready in it for rapid access. CentricStor automatically creates the legally required copies of physical data storage media automatically using dual save. The remainder of the disk storage pool is now used solely for restore operations for the large number of smaller file servers.

Long-distance high availability

Four CentricStor systems located in three cells are currently in operation at the data center in Bonn. The most recently deployed system not only performs dual save backups, it also provides cache mirroring and automatic failover in a fault-tolerant configuration logically spanning two cells. This setup enabled Postbank Systems to benefit from the CentricStor Grid Architecture, too, which extends system functionality to service two locations. In Postbank Systems' case, the solution connects up two data centers in Frankfurt and Bonn, a considerable distance apart.

One of the design requirements that the new backup solution had to meet was the ability to back up 30TB of uncompressed data every day within a period of 12 hours. It also had to create redundant backups in dual save mode, and the two tape libraries had to be located in separate fire cells. In addition, it was clear that the new CentricStor system in Frankfurt, the fifth in total, would have to have a redundant configuration. Plans to install a second PowderHorn 9310 tape library at the Frankfurt data center were shelved because the available building space meant it was impossible to site the system in a location that afforded adequate disaster protection. As an alternative, decision-makers considered running redundant backups over a dense wave division multiplexing (DWDM) line between Frankfurt and Bonn.

As the IT services unit was only too aware, this model involved serious challenges. Above all, the transmission time and the throughput required to stream data to the drives located in Bonn were difficult issues to resolve. Initial calculations showed that, using standard procedures, the throughput on the DWDM line would not be sufficient to complete error-free tape backups in the allotted time. However, a solution from McDATA (today, Brocade) designed to increase the efficiency of data transmissions over long distances provided an answer.



Postbank Online-Brokerage

Their patented tape pipelining technology streams data to tape using buffering and error recovery to ensure a reliable and efficient flow of data across large distances. Using Edge 3000 routers, tests showed that a data rate of 30 MB/s could be achieved at the tape drives, allowing 10TB of data (30TB uncompressed) to be backed up in a period of 12 hours using a set of eight drives. Following successful testing, this solution began production operation. The traffic between Frankfurt and Bonn runs over two DWDM lines on separate geographic routes covering distances of 170 km and 380 km respectively. Each line feeds four 9940B tape drives, each with a throughput of 30 MB/s. Postbank Systems is currently also considering the option of using the same technology to write duplicates from Bonn to Frankfurt.

Outlook: Consolidating even further

Guido Spallek's verdict: "CentricStor's strategy is really convincing. Above all, the outstanding scalability and the interoperability with all kinds of front-end and back-end systems offer enormous advantages. This assures us maximum flexibility with regard to future tape processing interfaces and formats." Postbank Systems' latest plans mark a logical extension of the backup infrastructure already deployed: First, another system in Bonn is to be expanded to support cache mirroring across two cells. Second, with maintenance support for the first generation of systems soon to be discontinued, Postbank Systems is planning to consolidate its CentricStor environment. With newer, more powerful hardware and greater cache capacity, the IT services unit will be able to sustain robust and efficient data backup operations for Deutsche Postbank AG in the longer term.

Postbank Systems AG

Postbank Systems AG is Deutsche Postbank AG's IT services subsidiary. With more than 1,400 employees, Postbank Systems supplies and operates a comprehensive range of IT systems and implements IT projects.

With its IT services, Postbank Systems drives Postbank's strategic objectives by leveraging opportunities afforded by forward-looking, competitive technologies to extend the bank's quality and cost leadership in the area of standard products through greater economies of scale, process standardization and automation.

Postbank AG



With 14.6 million customers, Deutsche Postbank AG is one of the largest banks in Germany. Although its primary focus is on retail banking for private individuals, it also does business with corporate customers, and its Transaction Banking segment provides back-office services for other financial services organizations. Deutsche Postbank AG went public in June 2004 and has been DAX-listed since September 2006. Its majority shareholder is Deutsche Post AG.

Postbank's head office is in Bonn. The Postbank Group also has 11 locations in Germany and regional offices in Luxembourg, Brussels, London, Verona, New York and New Delhi. The bank serves its customers through a network of 850 branches, several thousand Deutsche Post branches, and over 800 customer advice centers.