

Teradactyl® Provides Scalable Backup for Distributed Data Environments

The Problem

When the True incremental Backup System® (TiBS) Release 2.0 began shipping in January of 2001, the School of Computer Science at Carnegie Mellon University (CMU-SCS) wanted to know how it might help them improve their data backup and recovery. Patent-pending TeraMerge® technology offered to reduce network loads created by their current backup solution. "It was important to minimize the impact on network load", said Daryl Clevenger, Supervisor of Maintenance Programming at CMU-SCS, "The feature of taking daily incremental backups and accumulating changes over time was an attractive way to minimize the impact on our network." TiBS generates new full backups using data already present on the backup server. At the time, six backup servers provided support for about three terabytes of data distributed over hundreds of workstations and servers. Of primary concern to CMU-SCS were:

- Backup impact on networks
- Security
- Restore reliability
- Cost

To limit the network load backups were generating, CMU-SCS had adopted a four level backup schedule. Full backups (level 0) were run once each year along with periodic monthly (level 1), weekly (level 2) and daily (level 3) backups. This schedule results in a reasonable restore tape count (never more than 4) along with a lower network impact than more aggressive backup schedules. At the time, TeraMerge only supported two backup levels, full and incremental.

Leveraging Technology

Continued research showed that the merging process could generate multiple backup levels that simulate the way that traditional level backups function. Data on the backup server could be reused to produce any number of lower level backups. Most importantly, the impact of the backup server on the network would remain the same, regardless of the number of backup levels configured. Within a few months, multiple-level TeraMerge support was available for testing by CMU-SCS. The results showed that CMU-SCS would be able to deploy a single TiBS server to replace their existing six backup servers. The single server would generate much lower network traffic in a reduced backup window by only backing up daily file changes over the network. The additional processing required to produce four levels of backup would be completely handled by the backup server with no network impact.

Beating the Competition

CMU-SCS evaluated several top competitor products. Expensive pricing models based on the number of clients and an unwillingness to solve specific needs made the choice easy. The operating system pricing model from Teradactyl controls costs and is easy to maintain. Custom engineering services produced a Kerberos version of TiBS that integrated nicely into the customer's authentication and security model, something the

competition wouldn't even consider, "We highly value security", said Dale Moore, Associate Director of Computing Facilities at CMU-SCS, "Some competing methods for implementing backups just did not meet our requirements. Teradactyl was willing to add the authentication mechanisms and security that were important to us."

"We have consolidated from 6 servers to 1 while increasing the number and capacity of the machines that we are backing up"

Results, Then and Now

TiBS was deployed in the summer of 2001 on a Sun Ultra 5 with 3 DLT 8000 tape devices and three 70 GB disk drives for the backup server cache. "With TiBS, we were able to create a much simpler server configuration to meet our backup needs," said Daryl, "We have consolidated from 6 servers to 1 while increasing the number and capacity of the machines that we are backing up." Teradactyl has continued to refine the multi-merge process, and today, over a year and a half later, the same hardware configuration is supporting almost six terabytes of data. Only additional disk cache space has been added to keep up with the increased data load. "We process several restores each month with TiBS" said Daryl, "and they have worked smoothly and reliably. The file lookup database make it easy to find the desired version of the data." The TeraMerge process provides built in redundancy that validates all lower level backups. Failed tape volumes are detected and recreated.

Ongoing Data Growth Management

Improvements in backup hardware including CPU, memory, bus and device speeds will allow the CMU-SCS department to continue to scale using the same software licenses they originally purchased. "Working with CMU-SCS, we've shown that our software (TiBS) can support large amounts of data with minimal backup hardware", adds Kristen J. Webb, President and CTO of Teradactyl, "Refinements in TeraMerge will continue to improve efficiencies in our use of backup hardware. Advances in backup hardware will provide our customers with possibly the lowest TCO backup and recovery solution."

For more information about TiBS and TeraMerge® or to work with Teradactyl, visit www.teradactyl.com or call 1 505-242-1091



Copyright © 2002-2004 Teradactyl LLC. All rights reserved. Teradactyl, TeraMerge, and True incremental Backup System are registered trademarks of Teradactyl LLC.