

UNITED STATES AIR FORCE CAPABILITIES INTEGRATION ENVIRONMENT

EMC Data Domain deduplication storage earns its wings at U.S. Air Force Capabilities Integration Environment



ESSENTIALS

Challenges

- Full backups were taking up to four days to complete
- Data retention was restricted to two weeks due to capacity limitations
- Difficulties managing a costly VTL system
- Time-consuming and inefficient recoveries

Solutions

- Two EMC Data Domain DD880 deduplication storage systems
- Two EMC Data Domain DD690 gateway appliances
- EMC Data Domain Replicator software

Key benefits

- Time to successfully complete full backups improved by more than 75 percent
- Flexible data retention periods based on customer requirements
- EMC Data Domain provided almost immediate ROI by eliminating the need for the purchase of additional storage
- Recoveries were simplified and are now accomplished in a few minutes
- Dramatically reduced staff time spent managing the backup processes

The Global Combat Support Systems (GCSS) of the United States Air Force implements a vision of modern Air Force operations through a broad set of enterprise services. The GCSS supports over 50 major systems including several Oracle Enterprise Resource Planning (ERP) implementations, training simulators, IBM WebSphere, Oracle databases supporting Oracle Recovery Manager (RMAN), Microsoft SQL Server, Microsoft BizTalk, IBM Tivoli Access Manager, and the Microsoft .Net platform.

In support of GCSS and its customers, test and development (T&D) takes place in a unit known as the Capabilities Integration Environment (CIE). To give some idea of the sheer scope of USAF IT operations overall, the Montgomery, Alabama, data center supporting the CIE is managed by a 20-person infrastructure staff and hosts 1,400 servers with half a petabyte (PB) of storage organized in two tiers. About 320 terabytes (TB) sits on Tier One Fibre Channel (FC) disk arrays, while 180 TB on Tier Two is shared between EMC® Data Domain® deduplication storage systems and conventional disk arrays using SATA drives.

With approximately 50 USAF customers to support in T&D at any one time, backing up and restoring CIE systems efficiently and within necessary backup windows was a constant struggle. By implementing Data Domain systems, backup and recovery challenges have been eliminated, deduplication has solved storage capacity restraints, and the CIE unit has received almost immediate ROI benefit.

UNABATED DATA GROWTH DEMANDS NEXT-GENERATION BACKUP

Prior to implementing EMC Data Domain, CIE experienced difficulties completing its full weekly backups using existing technologies. Each weekend, the backup would begin on Friday evening and would sometimes continue until well into Tuesday. This tied up personnel resources for half the week and impacted mission-critical systems which were denied computing resources and bandwidth. Further, the organization suffered from a perennial problem—the explosion of information being generated meant that no matter how many new storage systems were added, they seemed to always be running out of capacity.

CIE initially attempted to resolve these issues via the purchase of a Virtual Tape Library (VTL). After one year, however, that project remained stalled. Software and integration bugs plagued the system. To make matters worse, the addition of a VTL only postponed the inevitable—the VTL filled up quickly and the CIE team was back to square one, needing more capacity with an ever increasing backup window. It was fortunate that the team only needed to retain two or three weeks of data onsite before a purge.

DEDUPLICATION PROOF-OF-CONCEPT SUCCESSES

When the CIE team researched the benefits of deduplication, they immediately understood what was happening in their environment. They realized that retaining multiple copies of the same data was greatly increasing CIE storage needs. By implementing the right deduplication technology, the unit could slash backup times, simplify restores, and minimize the necessity to purchase more disk arrays. The CIE team thoroughly tested the Data Domain solution and was immediately impressed with the data reduction and performance numbers they saw in their environment.

The team's due diligence reinforced EMC Data Domain as the best solution for CIE due to its market leadership with the most mature deduplication technology. Another factor in EMC's favor—the availability at the time of Data Domain gateway appliances which could be harnessed in tandem with existing storage and which integrated seamlessly with Symantec NetBackup software. For the CIE team, the choice of EMC Data Domain became a “no-brainer.”

Soon after, CIE received a request from a customer who needed a lot more storage to host an important application under development. This project also required reliable replication capabilities. Based on the success of the DD690 gateway appliances, the unit purchased two Data Domain DD880 systems which provided more than enough capacity for the customer. By deploying these units in separate buildings, any data backed up on one would automatically and efficiently be replicated to the other.

The CIE team found that the original VTL implementation still had problems after a year; in comparison, the EMC Data Domain DD880 systems were up and running smoothly in just over a day. In addition, after six months in production, the team reports that EMC support is also on the ball—resolving a software code issue in less than a day.

EMC DATA DOMAIN DEMONSTRATES BUSINESS AND OPERATIONAL VALUE

On an immediate basis, CIE fulfilled the demanding requirements of its USAF customer by adding two DD880 systems. As this customer has very specific requirements that are quite different from most CIE users, it called for a longer retention period. The combination of deduplication and replication between the Data Domain DD880 systems enable the CIE team to retain customer data as long as their customers require. Today, the Data Domain systems are achieving a deduplication rate of 13 times across protected applications, which translates to a tremendous cost and space savings.

The 13 times deduplication rate equates to 130 TB of data being compressed into only 10 TB. The resulting benefits to the unit include better service to its customers and much more flexibility on what were previously stringent retention policies. With more space available, it is no longer a necessity to purge its data every two weeks. These days, CIE can easily support backup retention for more than a month. Even then, its most utilized DD880 is only 60 percent full.

In addition, restores are no longer a concern. In the past, restore processes were cumbersome and time consuming. In contrast, CIE recently restored an LDAP server in 15 minutes. And the unit's backup window has been reduced from about four days to just over one day—a 75 percent reduction.

Prior to deploying Data Domain systems, CIE lacked the wherewithal to directly back up its Oracle databases. Instead, RMAN data was transmitted to a NAS head which was then backed up by the VTL. Now, CIE writes that database directly to the deduplication system. Oracle RMAN dumps, which currently stand at 430 TB, are being stored on just 31 TB of physical storage, for an overall data compression rate of 92.6 percent. And the team can accomplish this task with far less time managing the process.

Within a month of implementing its original Data Domain purchase, the CIE team had completely eliminated the VTL system from service, and its storage capacity was reused elsewhere in Tier Two. By purchasing EMC Data Domain deduplication storage, CIE was able to reclaim 180 TB of VTL-based storage, and the money saved in future storage acquisitions more than paid for the Data Domain DD880 systems CIE subsequently purchased.

The CIE Team Lead's advice to other IT shops is to use hardware and labor costs to calculate annual savings to see if a data deduplication solution is the right choice.

According to the CIE team, the Data Domain solution that replaced the VTL cost only a third of the price of the VTL, yet provided an extra 100 TB due to deduplication benefits. The 100 TB could potentially replace nearly a petabyte of backup storage based on the deduplication ratios the systems have achieved. This equates to literally millions of dollars in savings on storage capacity for which the CIE team will no longer need to budget.

CONTACT US

To learn more about how EMC products, services, and solutions can help solve your business and IT challenges, contact your local representative or authorized reseller, or visit us at www.EMC.com.

EMC², EMC, Data Domain, and the EMC logo are registered trademarks or trademarks of EMC Corporation in the United States and other countries. All other trademarks used herein are the property of their respective owners. © Copyright 2011 EMC Corporation. All rights reserved. Published in the USA. 07/11 Customer Profile H8825