A Large Financial Institution Migrates Datacenters with No Downtime Using HPE Shadowbase ZDM

Keith B. Evans >> **Shadowbase Product Management** >> Gravic, Inc.
Paul J. Holenstein >> **Executive Vice President** >> Gravic, Inc.

HPE Shadowbase Zero Downtime Migration (ZDM) Overview

How is HPE Shadowbase ZDM able to migrate systems with no application downtime? An overview of the ZDM process is shown in Figure 1 through Figure 4. (Note that these figures do not reflect the subject company's actual system configuration, but demonstrate the general principles involved.)

Step 1: Set up the new system (Figure 1).

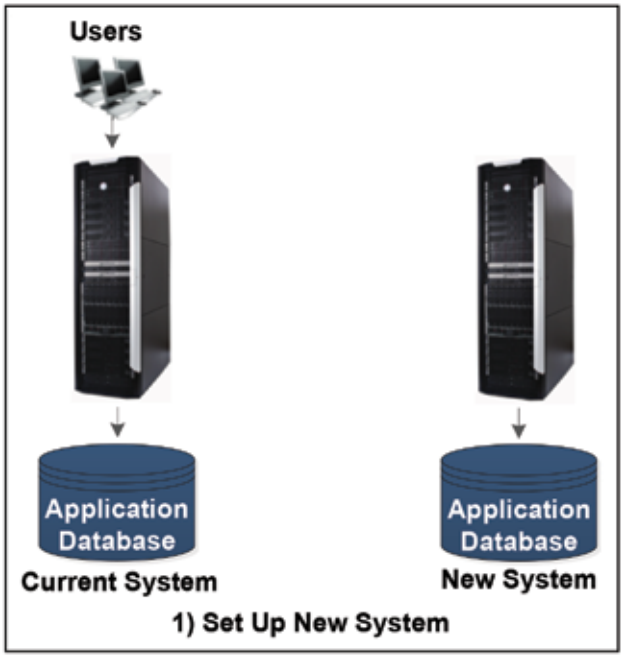

Figure 1 – HPE Shadowbase ZDM Step 1

Step 2: Load the existing source application database onto the new (target) system (Figure 2).

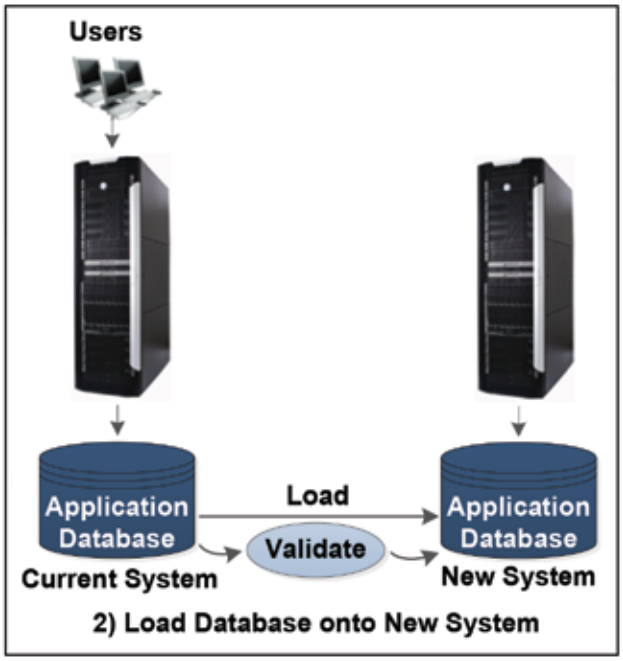

Figure 2 – HPE Shadowbase ZDM Step 2

Step 3: Once the load is complete, it is good practice to validate that the target database is an exact replica of the source database before proceeding. (The HPE Shadowbase Compare product may be used to perform this verification.) Uni-directional data replication is enabled so that the database of the new system will remain synchronized with the current system while the current

system continues to process transactions. Step 3 allows the testers to work on a current copy of the production database to verify proper functioning with 'real' data.

Step 4: Test the new system (Figure 3).

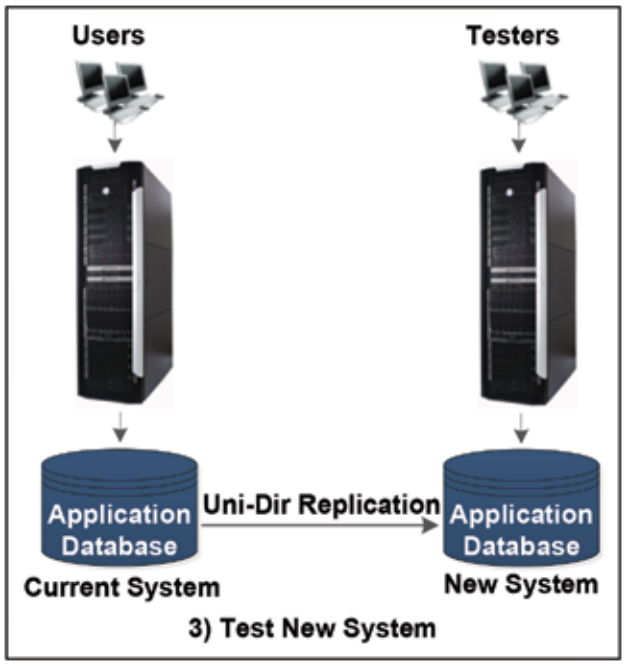

Figure 3 – HPE Shadowbase ZDM Step 3

Finally, when the testing of the new system is complete, bi-directional replication is enabled in Step 4 and all users are moved to the new system (Figure 4). Using bi-directional replication allows the original system's database to remain current when first using the new system, which greatly speeds a failback to the original system without losing any data.

Step 4: Enable bi-directional replication and move users to the new system.

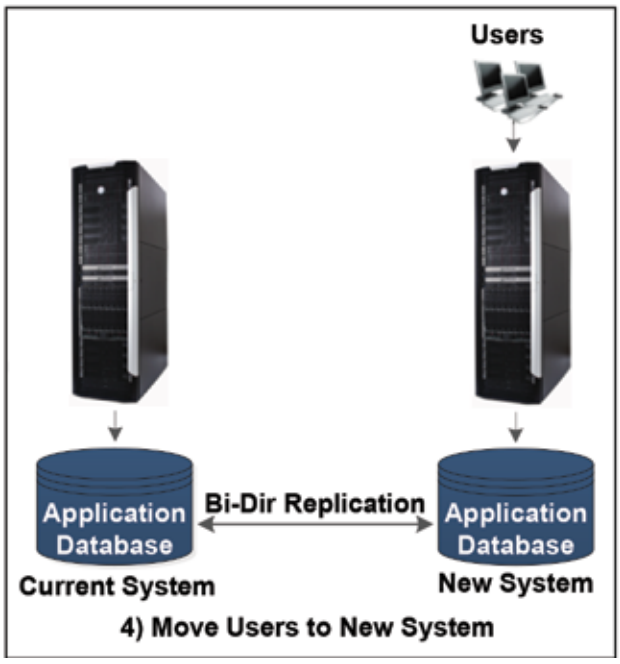

Figure 4 – HPE Shadowbase ZDM Step 4

After an appropriate period of time passes (to ensure the new system is properly functioning), the original system can be taken down (perhaps for its own upgrade).

Migration of users to the new system is therefore accomplished with no application downtime. Furthermore, the normal risks of migration are eliminated. Users are moved to a known properly functioning system, and the existing, unchanged system is always available, providing service. If there is a problem, users can quickly be returned to the original system until the problem is corrected. And, with bi-directional replication being deployed, any data modified on the new system is reverse-replicated to the original system so that no data will be lost if a failback occurs.[1]

In summary, even the most challenging of migrations (hardware, software, and location) can be undertaken efficiently and with low risk using HPE Shadowbase ZDM technology and methodology.

The Company's Migration Objectives

The company's original active/active system included a twelve-processor, four-core NonStop NB54000 system, DC1 (in Datacenter 1), and a similar system, DC2 (in Datacenter 2). The two systems were interconnected via the company's existing bi-directional data replication engine, called "Old Bi-Dir Replication," as shown in Figure 5.

The company wanted to accomplish several objectives with its migration:

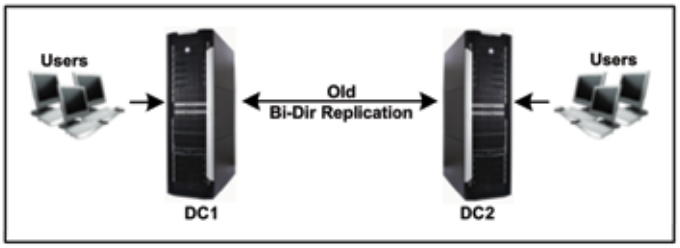

Figure 5 – The Company's Original Active/Active System

1. Retire DC1, a twelve-processor four-core NonStop NB54000 system, located in Datacenter 1, replacing it with a new twelve-processor four-core NonStop NB56000 system, called DC3, installed at a new datacenter location (Datacenter 3).
2. Move the DC2 system from Datacenter 2 to a new datacenter location (Datacenter 4), and rename the system DC4.
3. Replace its current data replication product ("Old Bi-Dir Replication") with a HPE Shadowbase active/active solution. (Its original data replication product was expensive, functionally stable and not being enhanced.)
4. Train its entire systems and operations center staff on HPE Shadowbase using a hands-on approach while the project occurred.
5. Achieve all of these goals with no application outage and no decrease in business continuity protection during the entire process. At all times, there must be at least two copies of the application and database available on two separate nodes.
6. Complete the project in a few short weeks as leases were expiring.

Figure 6 shows the company's final active/active system configuration:

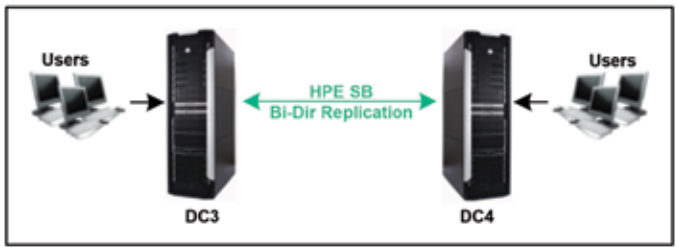

Figure 6 – The Company's Final Active/Active System

[1] For a more detailed description of HPE Shadowbase ZDM, please see the Gravic white paper, *Using HPE Shadowbase to Eliminate Planned Downtime via Zero Downtime Migration*.

All of these objectives were met with the use of HPE Shadowbase ZDM. The migration proceeded as follows:

Step 1: Install and Load the DC3 System
Step 2: Migrate Users from DC2 to DC3
Step 3: Shutdown and Move the DC2 System to Datacenter 4 (Rename it to DC4)
Step 4: Migrate Users from DC1 to DC4
Step 5: Retire the DC1 System

The following sections drill down into the details of each step.

Step 1: Install and Load the DC3 System in Datacenter 3

The company's first step was to obtain its new NonStop NB56000 (DC3) and install it in Datacenter 3 (Figure 7). The company timestamped the Audit Trail in the DC1 system and loaded the DC3 database from the DC1 database with data up to the timestamp. The company next enabled HPE Shadowbase bi-directional replication between DC1 and DC3 and began replication to DC3 from the DC1 timestamp, keeping the DC3 database synchronized with the DC1 database. The DC3 system was then exhaustively tested and verified before being put into service.

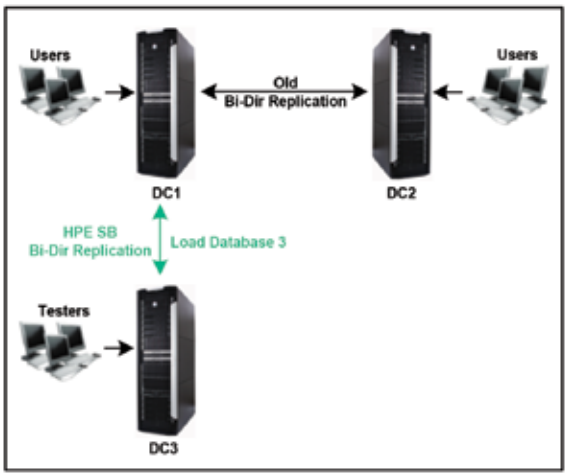

Figure 7 – Add the DC3 Node

Step 2: Migrate Users from DC2 to DC3, Isolating DC2 for the Shutdown/Move

When the test of the DC3 system was successfully completed, the migration of users from the DC2 system to the DC3 system began (Figure 8).

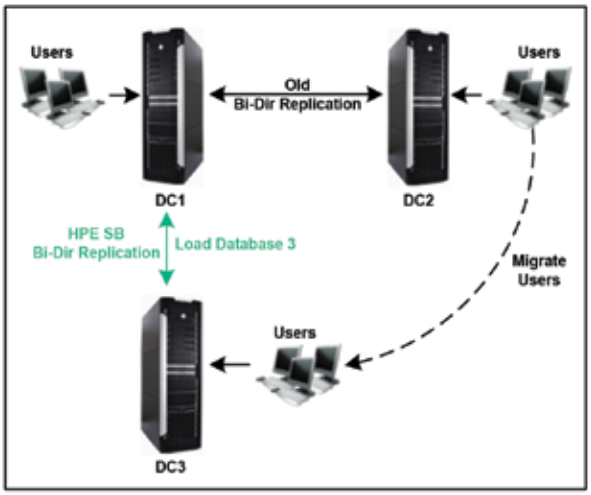

Figure 8 – Migrate Users from DC2 to DC3

Once the migration was completed, the DC2 system no longer had any users attached to it. At this point, the old data replication product that had kept the DC2 database synchronized with the DC1 database was decommissioned, and bi-directional replication between DC3 and DC2 was established using HPE Shadowbase (Figure 9). During this step, the Shadowbase configuration was verified before moving the DC2 system.

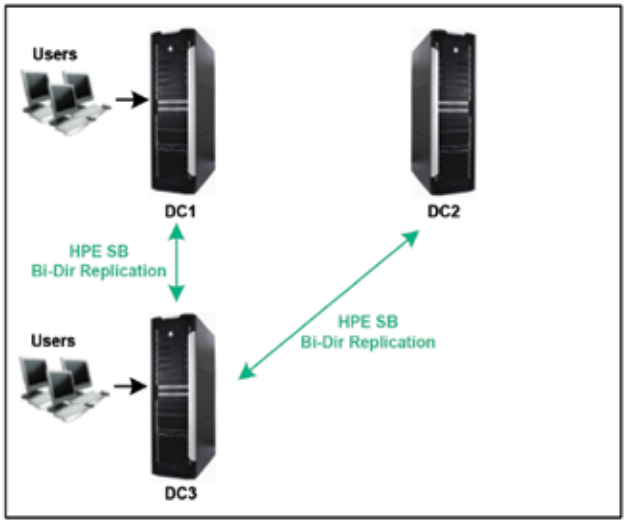

Figure 9 – Migrate Users from DC2 to DC3

Step 3: Shutdown and Move the DC2 System to Datacenter 4 (Rename it to DC4)

Replication between the DC3 and DC2 systems continued until the DC2 system move was ready, and then was paused with a timestamp in order to create a replication restart point. The DC2 system was then shut down and moved to Datacenter 4 (Figure 10).

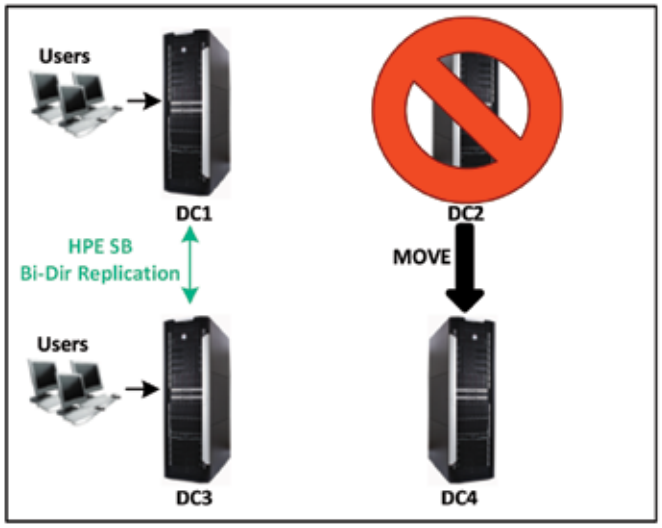

Figure 10 – Shutdown, Move, and Rename DC2 to DC4

The previously validated HPE Shadowbase data replication configuration was then deployed between the DC3 system and the DC4 (previously DC2) system (Figure 11)

The DC3 queue of database changes that had built during the move was then flushed to the DC4 system by the HPE Shadowbase replication engine from its timestamp replication restart point, thereby making the DC4 database current. The two systems were continuously kept synchronized via the HPE Shadowbase data replication engine.

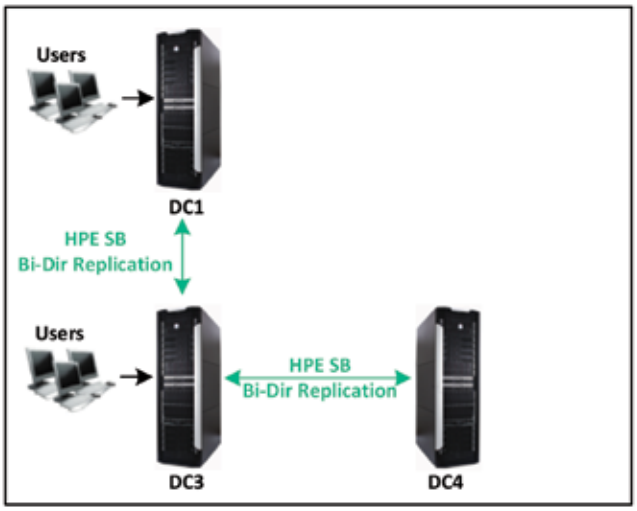

Figure 11 – Synchronize DC4 with DC3

Step 4: Migrate Users from DC1 to DC4 to Isolate the DC1 System

Next, the users served by the DC1 system were incrementally moved to the DC4 system (Figure 12).

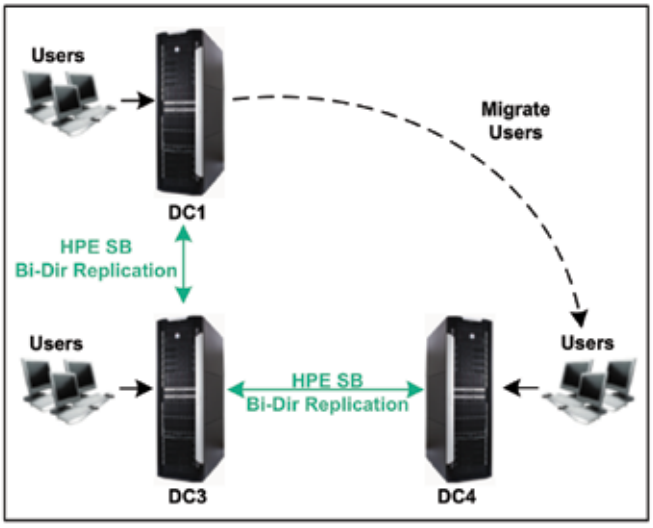

Figure 12 – Migrate Users from DC1 to DC4

Step 5: Retire the DC1 System to Create the Final DC3/DC4 Active/Active Solution

After all the users were migrated to the DC4 system, no users were left on the DC1 system. So, the HPE Shadowbase data replication engine connecting the DC1 and DC3 systems was decommissioned, and the DC1 system was retired, resulting in the final configuration (Figure 13).

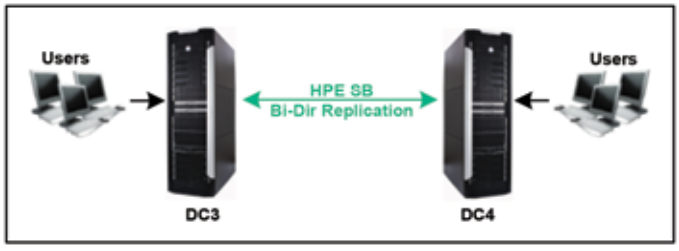

Figure 13 – Retire the DC1 System, Final Configuration

Summary

This migration presented a number of challenging parallel requirements:

- Complete within a few weeks
- No application downtime
- Full business continuity protection maintenance at all times (meaning at least two systems continuously available in the event that one system failed)
- Physical replacement (upgrade) of one hardware system
- Replacement of the existing data replication software product with a new product (HPE Shadowbase)
- Relocation of all systems and networks to new datacenters in other cities
- Minimal risk, only proceed to the next step once the current step is tested and proven
- Easily accomplish failback of each step if necessary

These requirements were undeniably challenging, but the fact that they all were met is testament to the power of the HPE Shadowbase ZDM approach, which allowed for the migration of systems with no application downtime. The company used HPE Shadowbase ZDM to change its data replication software product, to move its active/active system to two new cities, and to perform a system upgrade in the process. During all of these activities, the company suffered no application downtime.

Keith B. Evans works on Shadowbase business development and product management for Shadowbase synchronous replication products, a significant and unique differentiating technology. Asynchronous data replication suffers from certain limitations such as data loss when outages occur, and data collisions in an active/active architecture. Synchronous replication removes these limitations, resulting in zero data loss when outages occur, and no possibility of data collisions in an active/active environment. Shadowbase synchronous replication can therefore be used for the most demanding of mission-critical applications, where the costs associated with any amount of downtime or lost data cannot be tolerated. For more information and the availability of Shadowbase synchronous replication, please email sbproductmanagement@gravic.com.

Paul J. Holenstein is Executive Vice President, Gravic, Inc. He has direct responsibility for the Gravic, Inc. Shadowbase Products Group and is a Senior Fellow at Gravic Labs, the company's intellectual property group. He has previously held various positions in technology consulting companies, from software engineer through technical management to business development, beginning his career as a Tandem (HPE NonStop) developer in 1980. His technical areas of expertise include high availability designs and architectures, data replication technologies, heterogeneous application and data integration, and communications and performance analysis. Mr. Holenstein holds many patents in the field of data replication and synchronization, writes extensively on high and continuous availability topics, and co-authored Breaking the Availability Barrier, a three-volume book series. He received his BSCE from Bucknell University, a MSCS from Villanova University, and is an HPE Master Accredited Systems Engineer (MASE). To contact the author, please email: SBProductManagement@gravic.com . Hewlett Packard Enterprise directly sells and supports HPE Shadowbase Solutions (www.ShadowbaseSoftware.com); please contact your local HPE account team.