


HPE Shadowbase Solutions & HPE Pathway Domains – Perfect Together!



The Pathway Framework

HPE NonStop Pathway is an application framework that is a mainstay of today's mission-critical Online Transaction Processing (OLTP) systems. It supports a requester (client) – server model, and delivers application scalability by distributing requester workloads across dynamic pools of application server processes (serverclasses), that in turn are spread across multiple HPE NonStop server CPUs. It also offers high levels of application availability by monitoring and automatically restarting server processes in the event of failure.



The Pathway model delivers excellent application availability within a single HPE NonStop server, providing local fault tolerance. However, this model still necessitates application outages for certain planned system maintenance activities, and it offers no protection against unplanned system outages.

These limitations are addressed by HPE Pathway Domains (PDs), which allow the serverclasses to be distributed across multiple NonStop server systems. This capability improves application scalability and availability, protecting against some server failures. However, even though application availability is improved when using Pathway Domains, if the database being used by those applications becomes unavailable, then the application is still effectively down and unavailable.

Providing continuous data availability requires having at least two nodes, each being capable of hosting the application database. As an application makes changes (inserts, updates, and deletes) to its local database (the source database), those changes are sent immediately over a communication channel to the target system, where they are applied to the target database, keeping the databases synchronized. The target database typically resides on another independent node that may be hundreds or even thousands of miles away. If the server with the database fails, then applications can continue to run using the surviving database. HPE Shadowbase data replication provides all of the capabilities necessary to keep multiple database copies synchronized, ensuring current data remains available in the event of any outage, both planned and unplanned.

Technologies Converged: HPE Shadowbase and HPE Pathway Domains

When technologies converge, their applicability, usefulness, and effectiveness increase exponentially. The combination of Shadowbase data replication and Pathway Domains across nodes offers a powerful business continuity architecture to meet continuous application *and* data availability requirements. Working together, they provide geographically dispersed application and database services. The data is made available on other nodes by Shadowbase data replication, and the application services are made available on other nodes by Pathway Domains.

When a node outage occurs, for whatever reason, Pathway Domains automatically detects the outage and switches all user request traffic to the surviving node(s) (using the Pathway Automatic Weights Reconfiguration (AWR) feature). They can continue processing because the data on those nodes has already been synchronized with the failed node by Shadowbase data replication.

Thus, all the necessary components (applications and data) are in place to maintain business services with minimal downtime. In summary, Pathway Domains and Shadowbase data replication complement each other perfectly – together they provide scalability *and* protection from any single point of failure from both your application services and databases.

Multiple Business Continuity Architectures with Multiple Benefits

HPE Pathway Domains and HPE Shadowbase replication can be used together for great advantages in many business continuity architectures.

Active/Passive and Sizzling-Hot-Takeover (SZT) – What’s the Difference?

The central concept of Active/Passive (AP) and Sizzling-Hot-Takeover (SZT) architectures is that the application is only modifying the database on one node. And, with Shadowbase software, the application can be up and running on the target at all times with full read/write access to the database, ready to take over immediately. The main difference is that A/P architectures are configured for uni-directional replication and SZT architectures are configured for bi-directional replication (please see Figure 1). In practice, SZT architectures automatically facilitate re-synchronizing the databases when the downed system is recovered, reducing RTO from days/hours to minutes/seconds.

In these architectures (Figure 2), a Pathway Domain is configured across both the active and passive nodes, with the application servers up and running on both nodes (and when using Shadowbase software, with the passive node’s local database open for read/write access, ready to take over on a moment’s notice). User connections and Pathway requesters are connected and running on both nodes, but all server requests are routed by Pathway (ACS) to servers on the active node. (This routing is achieved by setting the Pathway Domain configuration weighting parameter for the PATHMONs on the passive node to zero).

Pathway Automatic Weights Reconfiguration (AWR)

The Pathway AWR feature assists in application recovery when an outage occurs (whether planned or unplanned). ACS automatically detects the outage and begins routing all requests to servers on the passive node, with no reconfiguration or operator intervention required.

Since Shadowbase replication has kept the active and passive databases synchronized, after an outage, servers on the passive (now active) node are using the most current data.

In addition, the Shadowbase configuration allows the servers on the passive node to have the local database already open for read/write (R/W) access. There is no manual user intervention required to start servers or switch them from read-only (R/O) to R/W access at takeover, which further reduces takeover times and the “failover fault” risk that they won’t start for some reason.

The Benefits of Combining Pathway Domains with Shadowbase AP/SZT Architecture in a Single Solution

- Half of the users see no outage at all; the other half can be quickly and easily switched over
- The passive application is already up and running and can be processing read-only transactions (a known-working “passive” system minimizes the likelihood of failover faults)
- Passive servers have the local database already open R/W, so no need to switch from R/O on takeover
- Reduced takeover times (better recovery time objective, or RTO)

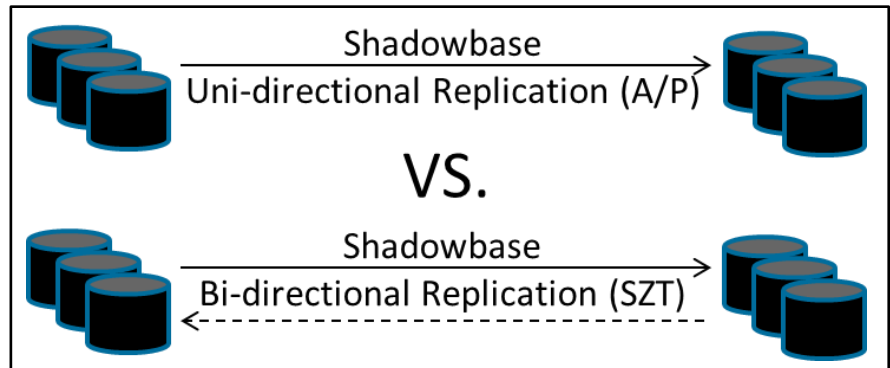


Figure 1 – Shadowbase Uni-directional (A/P) vs. Bi-directional (SZT) Replication Architectures

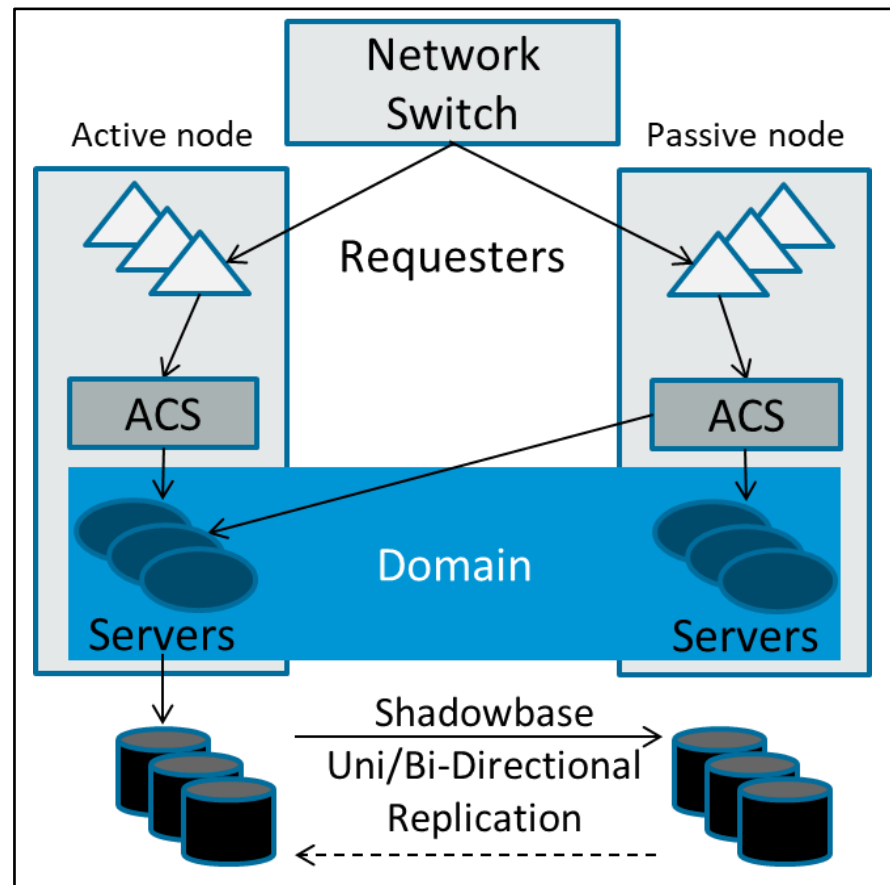


Figure 2 – Pathway Domains in A/P and SZT Architectures

- More system headroom (scalability) to handle peak loads
- Passive system capacity better utilized

Active/Active Systems

An active/active architecture improves on the active/passive or SZT approach by allowing all nodes to actively process requests against their own local copy of the database. This configuration differs from active/passive since the Pathway domain weighting parameter can be set to allow servers on both nodes to process transactions (Figure 3). The Shadowbase architecture is configured for bi-directional replication, thereby keeping both databases synchronized with changes made on either system. Data collisions may occur which will be detected and automatically resolved by Shadowbase replication (via customized processing). When an outage occurs (planned or unplanned), as for the active/passive case, ACS will automatically detect it and begin routing all requests to servers on the remaining node(s), with no reconfiguration or operator intervention required.

The Benefits of Pathway Domains with a Shadowbase Active/Active Architecture Include all of the Benefits of an Active/Passive or SZT Architecture, Plus:

- Less data loss (lower recovery point objective, or RPO), since only half of the in-flight transactions are affected (versus 100% for active/passive and SZT)
- Even more confidence in the serviceability of the backup system since it is running the entire application (including server components and database updates)
- No latency issues with long distance PATHSENDS or Network TMF transactions (the entire transaction can be processed locally against a local copy of the database)
- Each active/active system has lower average utilization
- Both servers are performing productive work
- Better balanced workload distribution across the nodes

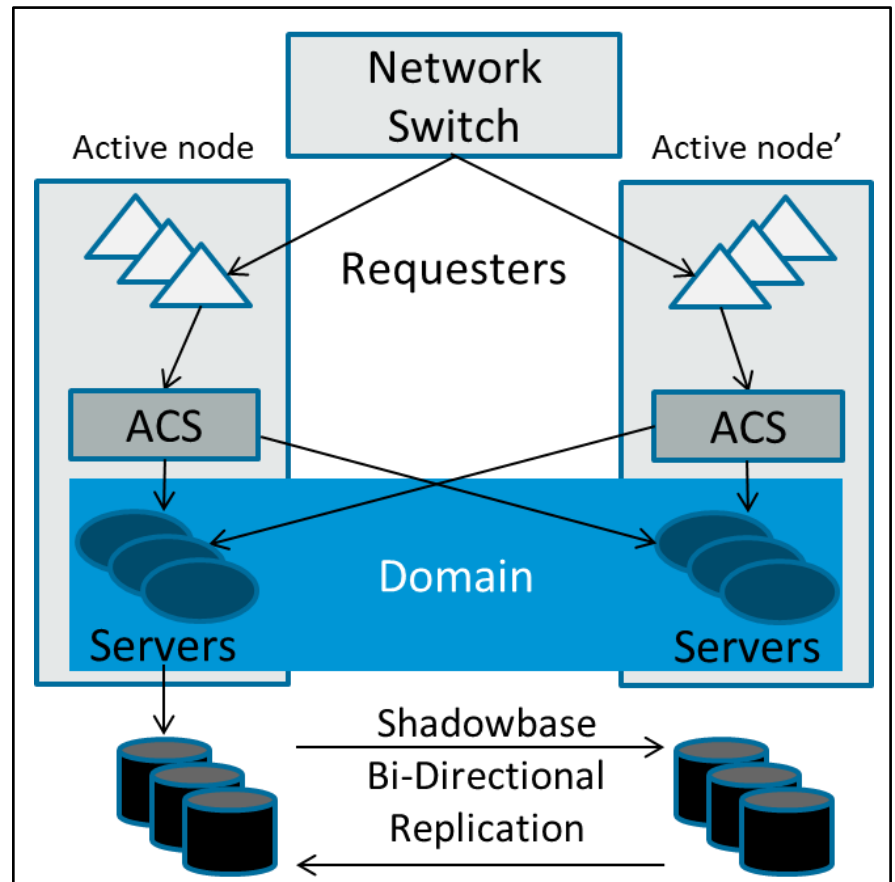


Figure 3 – Pathway Domains in A/A Architectures

Zero Downtime Migrations (ZDM) and Maintaining Business Continuity Protection

The use of Pathway Domains and HPE Shadowbase data replication to avoid unplanned outages is equally applicable to avoiding planned outages. During system maintenance or an upgrade, the online workload is routed by ACS from the downed system to another system in the domain to maintain business services while the maintenance is carried out.

In a two-node architecture, while one system is offline for maintenance, a failure of the other system will bring business services to a halt, which is unacceptable for many applications. Companies increasingly are deploying multi-node configurations to prevent this situation. In such configurations, the multiple systems capable of taking over the business services are located in different data centers.

In a Pathway domain, there can be up to four such Pathway and database environments, spread across up to four distinct data centers. Figure 4 depicts a common architecture with four domains spread across two data centers. Shadowbase bi-directional replication is used to keep all of the active databases synchronized. By having four available systems distributed across multiple data centers, planned or unplanned downtime of a system or an entire data center will always leave at least two systems available, thereby maintaining continued (geographically-distributed) business continuity protection.

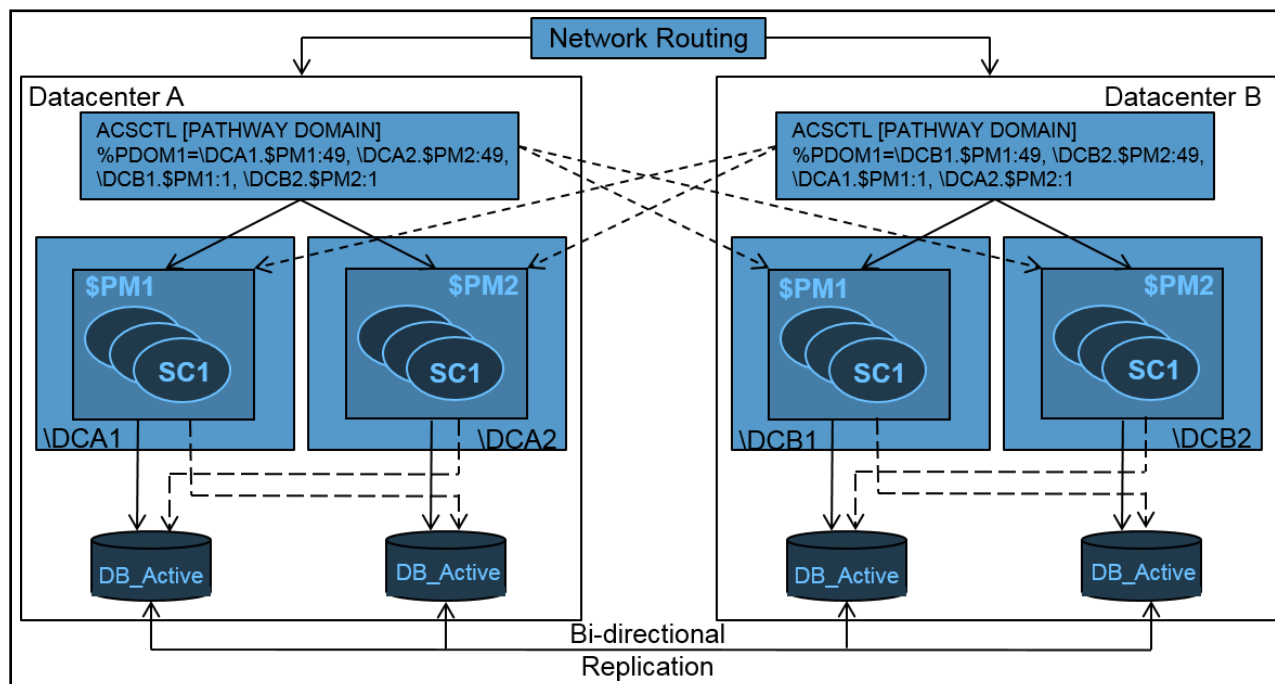


Figure 4 – Planned Downtime While Maintaining BC Protection

Summary

It takes two to tango! It is of no use having data available on alternate/backup/standby/takeover systems if the applications required to provide the business services are not also available when needed. HPE Shadowbase data replication solutions provide all the capabilities necessary to keep databases replicated and synchronized across multiple systems, whether it is a uni-directional (active/passive), bi-directional (SZT, active/active), heterogeneous, or homogeneous environment. The introduction of Pathway Domains across nodes for HPE NonStop servers also helps to ensure that applications remain available, even when catastrophic failures take out a system, an entire data center, or a geographic region. By configuring applications within a Pathway domain across multiple HPE NonStop nodes, in the event of planned or unplanned outages, user requests are automatically routed to the remaining systems, thereby preserving application availability. In addition, application scalability is increased by automatic workload distribution across those systems, and the removal of Pathway single system limits. Since Shadowbase technology is making sure the data is available on those remaining systems, the combination of Shadowbase software and Pathway Domains provides a firm basis to achieve increased levels of availability for your business services. In short, Pathway Domains and HPE Shadowbase data replication solutions are perfect together!

Hewlett Packard Enterprise globally sells and supports Shadowbase solutions under the name HPE Shadowbase. For more information, please contact your local HPE Shadowbase representative or visit our website. For additional information, please view our Shadowbase solution videos: <https://vimeo.com/shadowbasesoftware>.

Learn more:

shadowbasesoftware.com
hpe.com

Contact us:

Gravic, Inc.
 17 General Warren Blvd
 Malvern, PA 19355-1245 USA
 Tel: +1.610.647.6250
 Fax: +1.610.647.7958
 Email Sales: shadowbase@gravic.com
 Email Support: sbsupport@gravic.com

Please follow:

