Achieving Century Uptimes with HPE Shadowbase Active/Active Technology

A Gravic, Inc. White Paper

Executive Overview

Stuff happens. Unavoidable events such as fire, power outages, hardware faults, natural disasters, and more can take production IT systems offline. One study[1] shows that the average number of complete datacenter outages is one per year, with an average duration of 91 minutes. Other reports[2] estimate that the average business revenue lost per hour of downtime is $1.4M. Thus, the average annual business costs from unplanned outages are about $2.1M. For some industries the costs will be significantly higher. Further, the U.S. Bureau of Labor found that 93% of companies that suffer a significant data loss are out of business within five years. The net of these chilling statistics is that serious IT outages are not rare, and their consequences to the business can be dire. They are also completely avoidable via the implementation of an appropriate business continuity plan.

Business continuity encompasses those activities that an enterprise performs to maintain consistency and recoverability of its data, operations, and services. Application availability depends upon the ability of IT services to survive any fault, whether it is a server failure, a network fault, or a datacenter disaster. Data availability depends on the existence of up-to-date backup data copies. An enabling technology for achieving high or continuous availability for application services and the timely backup of important data is *data replication*. Here is where HPE Shadowbase solutions from Gravic, Inc. step in.

By replicating data in real-time between systems, HPE Shadowbase solutions enable recovery from unplanned outages in times ranging from minutes (high availability, disaster recovery) to immediate (continuous availability, disaster tolerance), with minimal data loss (measured in subseconds). HPE Shadowbase solutions support active/passive*, sizzling-hot-takeover* (SZT), and active/active business continuity architectures.

These same HPE Shadowbase solutions can also be used to eliminate planned downtime. By keeping backup systems current with data updates performed by production systems, backup systems can quickly be brought up to production status, allowing maintenance to be performed with no loss of business services.

For the most mission-critical services and data, those for which any outage is unacceptable, only an active/active business continuity architecture will suffice to protect the business against the costs of planned and unplanned outages. Using HPE Shadowbase replication to implement an active/active architecture is the focus of this white paper. Several case studies are provided that illustrate how several Fortune 500 companies are already using Shadowbase active/active technology to protect their businesses.

[1]Source: “Ponemon Institute – Study on Datacenter Outages.”
[2]Sources: Network Computing, The Meta Group, Contingency Planning Research.

Table of Contents

Table of Figures

What is an Active/Active System?

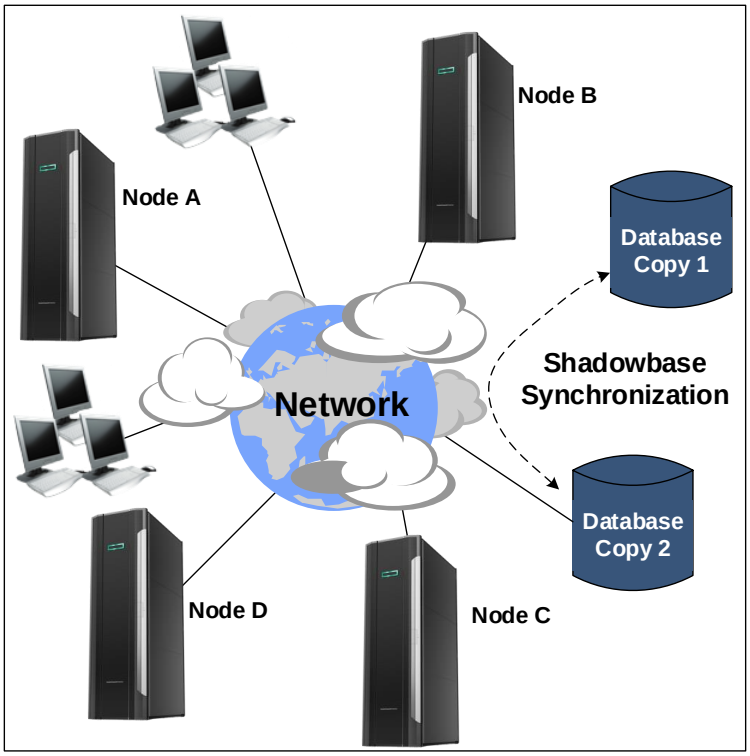

Figure 1 – An Active/Active System

As shown in Figure 1, an active/active system is a network of independent processing nodes, each having access to a common replicated database, such that all nodes can participate in a common application. In the most general case, the nodes are completely symmetric. Any transaction can be routed within the application network to any node which can read or update any set of data items in the database. This approach provides the most flexibility and maximizes system investment as requests can be load-balanced across all available processing capacity. If a node fails, users at the other nodes are unaffected. Also, the users at the failed node can be quickly switched to surviving nodes, thus restoring their services in seconds or less.

An active/active network contains at least two copies of the application database. All database copies are kept synchronized so that any copy can be used for a transaction. If a database copy fails, all transactions are routed to a surviving copy.

Provided that the nodes and database copies are geographically distributed, active/active systems provide continuous availability for little or no additional cost when compared with active/passive configurations. If a disaster takes out a node or a database copy, there are others in the network immediately available to take their place.

Why Does an Active/Active System Work?

The availability of a system is determined by the amount of time that it is operational and providing application services (the system *uptime*) as compared to the amount of time that the application services are being denied to one or more users (the system *downtime*).

Although certain techniques can be used to improve the uptime of an individual system, such as increased operator training and using fault-tolerant components, a single system can never provide the levels of uptime necessary for critical business services. Even if the system hardware, software, and operations are 100% reliable (which is impossible), a single local event such as a fire, power outage, or flood, will cause an outage. Active/active technology provides the necessary redundancy to reduce downtime by orders of magnitude.

As shown in Figure 2, if a node fails, users at that node can be switched to another operable node immediately. If a database fails, there is another consistent copy in the network that can be used. If a network component

fails, alternate routes are provided. Using technology available today, failure recovery can be achieved in seconds or less. In short, *let it fail (because it surely will), but fix it fast.*

Regardless of the type of failure, in an active/active configuration far fewer users are affected when a node or database fails than with other disaster tolerant architectures. For example, in an active/backup (classic disaster recovery) architecture, any failure and subsequent failover affects all users. Consequently the decision to failover (rather than try and recover the failed system) usually involves the approval of upper level management, which may be hard to quickly obtain, further delaying recovery. In an active/active system, these types of failures only affect the users on the failed node or database, not the entire user population. Since other known-working[3] nodes exist in the network, these users can be quickly switched to an alternate node.

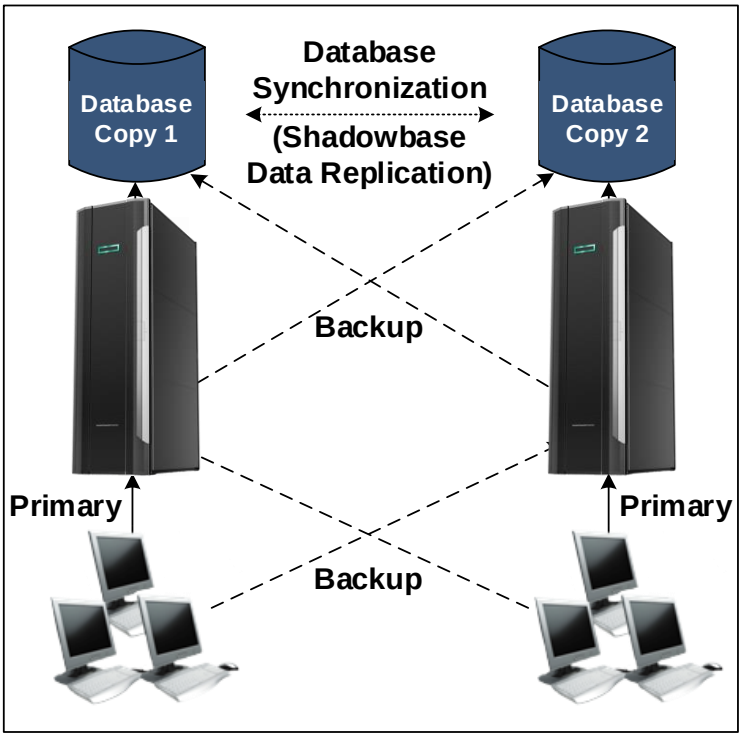

Figure 2 – How an Active/Active System Works

Active/active systems eliminate the uncertainty that always exists when an active/backup approach is in place. Such uncertainty as to whether the failover will be successful[4] often results in indecision, which further extends the outage duration. In an active/active system, when a failure occurs there is no massive leap-of-faith surrounding the failover to a backup system; all nodes in an active/active network are always known to be working, performing real work, at all times. For the active/active system, one only needs to re-route the users that were attached to the failed node to a surviving node, and this operation can often be masked from the users by network switching/routing software. Because of this extremely fast recovery, active/active systems provide continuous availability (recovery in seconds or subseconds), as opposed to active/passive systems which provide only high availability (recovery in hours or minutes).
Active/active architectures also allow for all purchased compute capacity to be fully utilized, actively working on satisfying user requests of any type (e.g., read and/or update). There is no backup (passive standby) system sitting idly by waiting for another component to fail. All nodes are actively performing real work, all of the time.

Locating the distributed nodes near user populations is another advantage of active/active systems. For example, if there are East Coast and West Coast (USA) sets of user populations, locating one node near the East Coast user population reduces their system access communications latency, likewise locating another node near the West Coast user population also reduced their communications latency. Note that an

[3]Since all nodes in an active/active architecture are up and running and processing transactions, they are known to be working correctly, which is not the case in an active/passive configuration, where the passive (backup) system must first be brought online and made operational.
[4]Because testing a failover in an active/passive environment requires taking an outage of the active (production) node – something no one likes to do – such testing is often not performed frequently or to completion, leading to the uncertainty of whether a failover will actually succeed.

active/passive architecture will suffer communication lag, if all users are routed to a single node, which may be thousands of miles away.

HPE Shadowbase Database Synchronization

A key requirement for implementing an active/active system is the synchronization of the databases. Each database copy must always be in a consistent state and must reflect the current state of the application. The Shadowbase solution accomplishes this task by automatically replicating changes made to each database copy to all other copies in the application network. The Shadowbase technology contains a powerful database replication engine that provides bi- or multi-directional replication between the database copies and guarantees that all copies remain in a consistent and correct state.

A concern that must be addressed in active/active database synchronization is that of data collisions. A data collision occurs when the application in each of the nodes make a change to the same row in its database copy at substantially the same time. Each will replicate its change to the other database copy, thus overwriting the change made there. As a result, the database copies are both inconsistent and incorrect. The Shadowbase software can detect collisions and automatically resolve them in many cases. For those cases where HPE Shadowbase replication cannot automatically resolve a collision, it supports embedding customer business logic into the replication engine to take whatever action is necessary to resolve the collision. There are also techniques to avoid data collisions in the first place, for example, by application or data partitioning (as discussed in one of the case studies below).

Another issue which arises in bi-directional data replication is *data oscillation* or *ping-ponging*. If measures are not taken to prevent this issue, a change made in one database is replicated to the other database, the change is applied, and consequently forever replicated back and forth to and from the original database. Shadowbase software contains patented algorithms to prevent the occurrence of data oscillation.

For initial target database loading, HPE Shadowbase Online Loader (SOLV) utilities are provided. These powerful and patented utilities can dynamically synchronize the target database to the source database while the applications are running on each node, updating the database. SOLV utilities combine the data being loaded with the data changes being replicated, so that no large queue of replication changes needs to be drained and applied after the initial load completes. Coupled with the SOLV utilities, the HPE Shadowbase tool set provides the critical elements necessary to convert an existing application into an active/active architecture.

Eliminate Planned Downtime

Along with eliminating unplanned outages, with Shadowbase active/active systems, there is no longer any need to take down the production application and deny service to users during planned downtime (for example, to upgrade hardware, software, the database, or the application itself). All that is necessary is to switch users from the node to be upgraded to another node, perform the upgrade, test it, and then return the users to the upgraded node. In this way, upgrades can be rolled through all nodes in the application network without ever denying service to any user. In addition, the system's capacity can be easily expanded by adding new nodes in a similar fashion.

Shadowbase technology provides the tools necessary for these zero downtime migrations (ZDMs). Its online copy facilities recreate the current state of the database on the new or upgraded node. Its data replication engine then keeps this database copy current until users are returned to that node. HPE Shadowbase replication also provides the facilities to move users to other nodes without losing their data if a problem occurs with the new node (i.e., in case there is a need to fail back).

Why Choose HPE Shadowbase Technology?

Shadowbase technology is a leader in data replication for active/active systems. Its avoidance of disk queuing points makes it fast and highly efficient, able to scale and handle the most demanding workloads. Fast replication means fewer data collisions and less potential data loss following a node failure.[5] Its efficient replication path also means less utilization of computing resources for replication.

[5] Contact Gravic, Inc. for details regarding Shadowbase synchronous replication product availability.

Shadowbase software is easy to manage and is itself fault tolerant. Of utmost importance is attention to referential integrity; it ensures that database updates are consistent and correct. Aiding this is Shadowbase technology's superior capabilities to automatically detect and resolve data collisions, and avoid data oscillation.

The Shadowbase high-speed SOLV utilities allow an active database to be copied while it is being actively updated. Any portion of the database that has been copied is immediately consistent and correct and can be used for application processing. These utilities simplify the implementation of an active/active architecture, moving either from an existing standalone or active/passive system configuration.

Shadowbase software is fully heterogeneous. Its active/active technology can be used on today's popular systems including UNIX, Linux, and Windows running Oracle, SQL Server, and Sybase databases, as well as HPE NonStop servers using NonStop SQL (/MP and /MX) and Enscribe databases.[6] Nodes in the Shadowbase active/active network can be similar, or completely different. There is no need for each node to be configured the same and to run the same software versions, as Shadowbase architecture handles interconnecting differing versions.

In summary, a Shadowbase active/active architecture's key benefits include:

- Extremely high system availability with uptimes measured in centuries
- Only a subset of users, rather than all, affected upon a node failure
- Restoration of service to users connected to a failed node in seconds
- Little data loss, if any, as the result of a failure
- Elimination of planned downtime
- Greatly improved disaster tolerance
- Support for dynamic load balancing (expandable)
- Utilization of all purchased system capacity for productive work (no idle standby systems needed)
- Elimination of the uncertainty and indecision associated with passive backup systems at takeover time (if the primary fails, will the backup system take over and work *properly*?)
- Accomplishment of all these benefits at *little* or *no* additional cost

[6] Contact Gravic, Inc. for the availability of additional commercial offerings, such as Db2®.

Case Studies

Shadowbase active/active technology is currently being used across many industries by some of the largest Fortune 500 companies. Some examples include regional bank service bureaus, international cell phone service providers, travel agencies, and many others.

Regional Bank Service Bureau

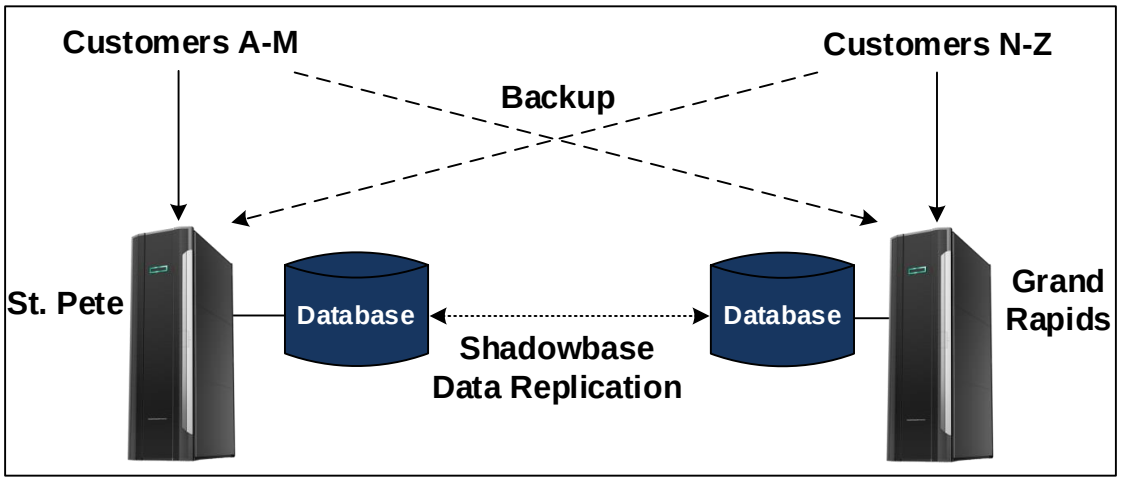

Figure 3 – Regional Bank Service Bureau

A large bank service bureau needed to provide financial services for other subscriber banks without any interruption of service. Focused on the need for continuous availability, the bureau turned to an active/active HPE NonStop system with one node in St. Petersburg, Florida, and one in Grand Rapids, Michigan. Shadowbase software is used to keep the database copies at each node synchronized as shown in Figure 3.

To avoid data collisions the database is logically partitioned, and each customer bank is assigned to one of the nodes as its primary node. All updates that are made to the database for a particular bank are always made at one node, and since each bank also has its own database partition data collisions are thereby avoided. The service bureau can dynamically reassign customers between nodes to provide load balancing.

If a node fails, the banks serviced by that node are automatically connected to the surviving node. In fact, during hurricane season, the service bureau routinely shuts down the St. Petersburg node whenever Florida is threatened by severe weather. All customer banks are serviced during this time by the node in Grand Rapids.

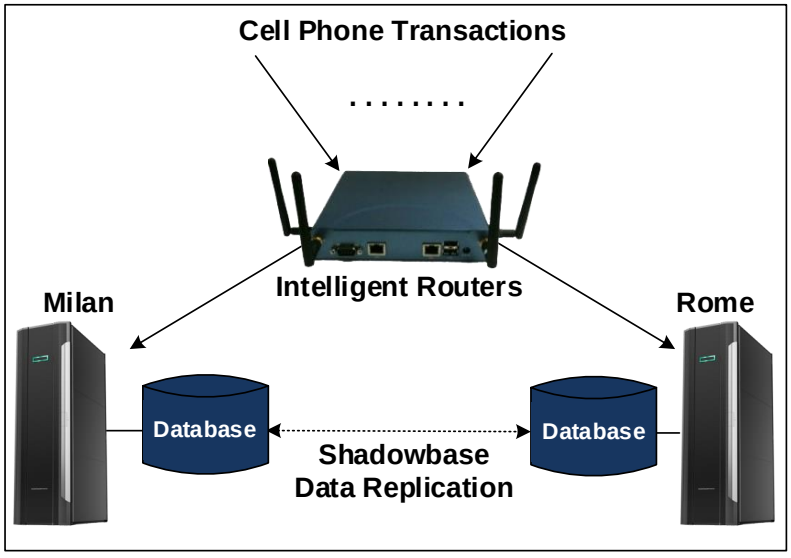

Figure 4 – International Cell Phone Service Provider

International Cell Phone Service Provider

An international cell phone service provider must support its subscribers without interruption. For disaster tolerance purposes, this company maintains one node in Milan, Italy, and one in Rome, Italy. With the variability in call activity, it needs to be able to instantaneously balance the load between these two nodes. To accomplish these goals, it implemented an active/active HPE NonStop system using the HPE Shadowbase product suite. Shadowbase software is used to keep the Milan and Rome databases synchronized as shown in Figure 4.

Transactions are routed according to a load balancing algorithm. When a transaction is received by the system, an intelligent router decides to which node the transaction should be routed based on the current load of each node. Either node may process any transaction. The system handles over 1,000 transactions per second without overloading either node.

It is possible that both nodes could receive transactions for the same cell phone at substantially the same time. In this case, a data collision might occur. Consequently, each node could end up with a different state for that cell phone, and both would be wrong.

These collisions are all numeric. For instance, one node may be adding minutes to the phone's SIM card while the other node is subtracting minutes currently being used. Shadowbase software resolves these data collisions via *relative* replication. With relative replication, the absolute (final or ending) value of the record is not replicated. Rather, only the numeric operation affecting the record is replicated (for instance, add 50 and subtract 3). By applying only these relative changes, the HPE Shadowbase replication engine ensures the consistency of the database in the presence of numeric collisions.

Large Travel Agency

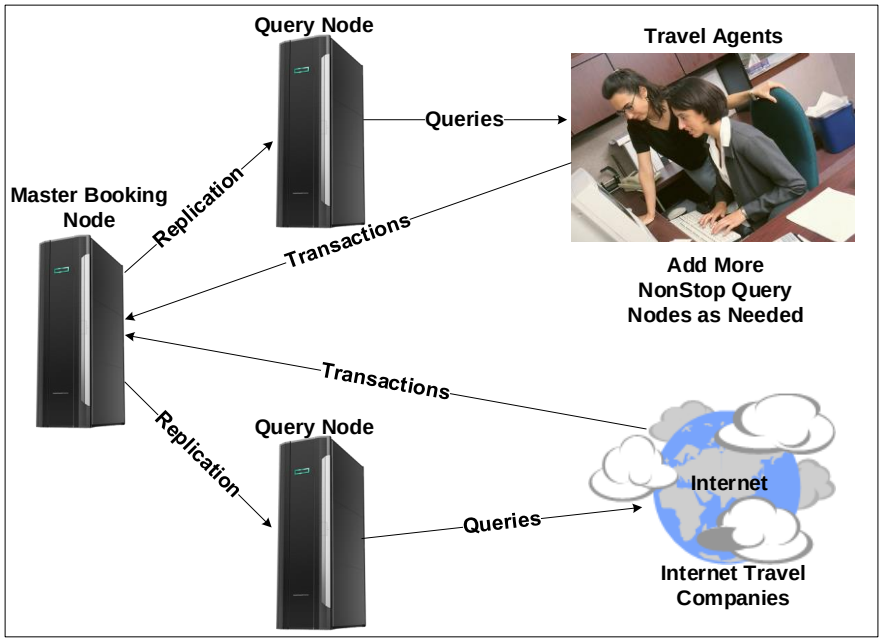

Figure 5 – Large Travel Agency

A major travel agency found that its HPE NonStop system activity could increase by up to two orders of magnitude or more during peak travel periods as customers searched for the best rates and availability. Most of the system load was devoted to large queries with relatively small update activity occurring when travel arrangements were finally booked.

To accommodate these wide swings in system activity, the travel agency turned to active/active technology to offload query processing from the booking system by creating an HPE NonStop master booking node and a series of HPE NonStop and Windows query nodes. The master booking node maintains the master database. Booking transactions are routed directly to the master booking node, which updates its master database. These changes are then replicated to the query nodes via the Shadowbase technology to support the extensive query processing as shown in Figure 5.

Query nodes can be added or removed as load dictates. If a query node fails, the other query nodes can still support all of the users. Also, recovery from a master booking node failure can be achieved by promoting one of the HPE NonStop query nodes to master. This architecture provides disaster tolerance “for free” as any of the query nodes has a full copy of the database and can be promoted to be the new master at any time.

Summary

In today’s always-on business world, access to real-time online transactional data is a competitive advantage. To realize the advantage, this data must be available at any time, all the time, and it must be current. The corollary to this idea is that the inability to access or update this current data carries a significant business cost, possibly measured in many thousands of dollars per second. These requirements necessitate an IT infrastructure that is continuously available, able to provide services throughout planned and unplanned outages. This environment is likely to be heterogeneous, with many different platform types and databases. These requirements can only be met by an active/active database replication architecture.

Gravic, Inc. is a world leader in providing innovative data replication solutions. HPE Shadowbase technology provides the means to meet these requirements, via reliable low-latency real-time data replication and distribution across heterogeneous systems and applications. HPE Shadowbase solutions include the necessary capabilities to implement continuously available active/active IT architectures, with features such as data collision prevention, detection, and resolution. With these powerful capabilities, as demonstrated by the above case studies, HPE Shadowbase products provide businesses with the tools needed to realize the competitive advantage of continuous access to real-time transactional data across the enterprise, and to avoid the significant costs of system and data unavailability, planned or unplanned.

International Partner Information

Global

Hewlett Packard Enterprise
1501 Page Mill Road #1429
Palo Alto, CA 94304-1100
USA
Tel: +1.408.691.2908
www.hpe.com

Japan

High Availability Systems Co. Ltd
MS Shibaura Bldg.
4-13-23 Shibaura
Minato-ku, Tokyo 108-0023
Japan
Tel: +81 3 5730 8870
Fax: +81 3 5730 8629
www.ha-sys.co.jp

Gravic, Inc. Contact Information

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

Hewlett Packard Enterprise Business Partner Information

Hewlett Packard Enterprise directly sells and supports Shadowbase Solutions under the name ***HPE Shadowbase***. For more information, please contact your local HPE account team or visit our website.