Disasters Happen, Will You Be Ready?

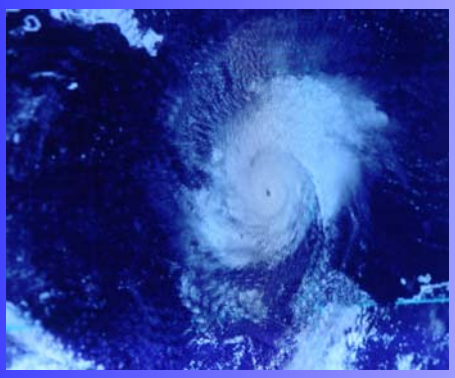



Achieving Century Uptimes with Shadowbase® Active/Active Technology

White Paper

Gravic, Inc.
301 Lindenwood Drive
Suite 100
Malvern, PA 19355-1758
+1.610.647.6250
Shadowbase@gravic.com
www.gravic.com

Disasters Happen, Will You Be Ready?

Achieving Century Uptimes with Shadowbase Active/Active Technology White Paper

Introduction

At the last minute the Category 4 hurricane veers to the west and misses the bank service bureau's main data center. After the shock of the near-miss wears off, management dictates that it must support its customer banks around the clock with no interruption whatsoever in service, be it from disaster, system failure, operator error, nor scheduled downtime. To achieve these demanding goals, the bureau turned to an active/active database replication solution powered by Shadowbase technology, available from Gravic, Inc., Malvern, Pennsylvania, USA (www.gravic.com).

Key benefits realized by the bureau's use of an active/active Shadowbase architecture include:

- An extremely high availability with uptimes measured in centuries,
- Only a subset of users are affected upon a failure,
- Service is restored to those users in seconds,
- Little loss, if any, in the data as the result of a failure,
- Planned downtime is eliminated,
- Greatly improved disaster tolerance,
- Expandable and support for dynamic load balancing,
- Accomplish all of this *at little or no additional cost*.

You, too, can realize the benefits of an active/active architecture by configuring your current or new systems to take advantage of Shadowbase technology. The following white paper explains how.

What is an Active/Active System?

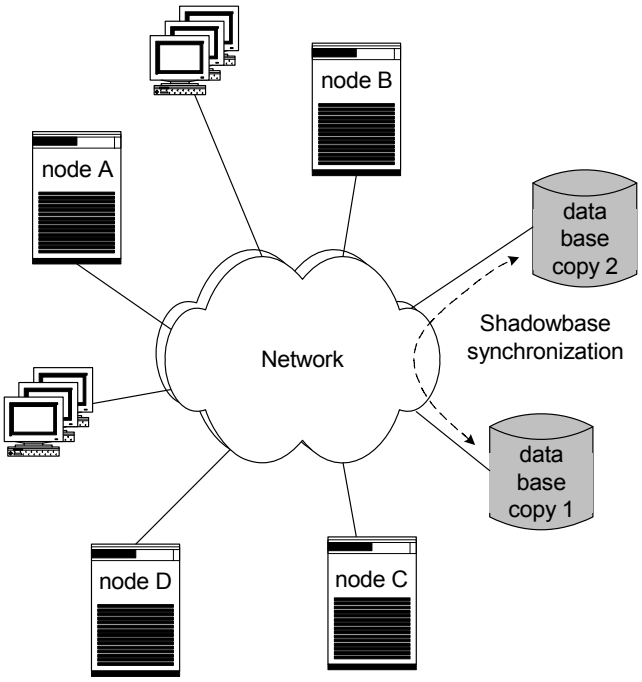

An active/active system is a network of independent processing nodes, each having access to a common replicated database, such that all nodes can participate in a common application.

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.

Should a node fail, users at the other nodes are unaffected. Furthermore, 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 in synchronism so that any copy can be used for a transaction. Should a database copy fail, all transactions are routed to a surviving copy.

Providing that the nodes and database copies are geographically distributed, active/active systems provide disaster recovery for little or no cost. Should a disaster take out a node or a database copy, there are others in the network to take their place.

Why Does Active/Active Work?

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

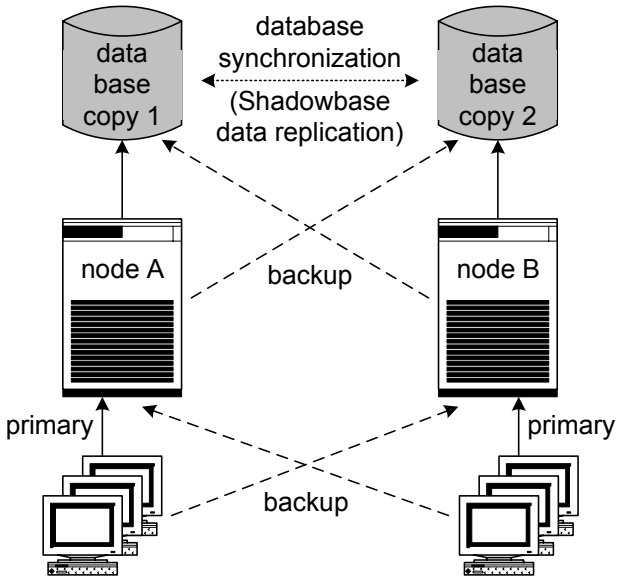

Although certain techniques can be used to improve the uptime of an individual system, such as increased operator training and using fault-tolerant components, there comes a point where there is not much more that can be done to add the necessary orders of magnitude to the uptime into the extreme reliability realm we are talking about –it is ultimately determined by hardware and software reliability, operator accuracy, and so on. What active/active technology does is reduce downtime by orders of magnitude.

If a node fails, users at that node can be switched to another operable node. 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, but fix it fast.*

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 replicating changes made to each database copy to all other copies in the application network. The Shadowbase technology contains a powerful database replication engine which provides bi-directional replication between the database copies and guarantees that all copies remain in a consistent and correct state.

Coupled with the Shadowbase AutoLoader, a powerful online loading utility, the Shadowbase tool set provides everything necessary to convert an existing application into an active/active architecture.

A concern that must be addressed in active/active database synchronization is that of data collisions. A data collision occurs when two nodes make a change to the same row in the database 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 different and both are wrong. The Shadowbase software can detect collisions and can automatically resolve them in many cases.

Eliminate Planned Downtime

With Shadowbase active/active systems, there is no longer any need to take the production application down and deny service to your users 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 a user.

In addition, the system's capacity can be easily expanded by adding new nodes in a similar fashion.

The Shadowbase technology provides the tools necessary for these zero downtime migrations. Its online copy facility, the Shadowbase AutoLoader, recreates 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. The Shadowbase technology provides the facilities to move the users to other nodes without losing their data should a problem occur with the new node.

Why Choose Shadowbase Technology?

The Shadowbase technology is the recognized leader in data replication for active/active systems. Its avoidance of disk queuing points makes it the fastest and most efficient active/active data replication engine in the marketplace. Fast replication means fewer data collisions and less potential

data loss following a node failure. Furthermore, its efficient replication path means less utilization of computing resources for replication.

The Shadowbase software is easy to manage and is itself fault tolerant. Of utmost importance is the Shadowbase software's attention to referential integrity. It guarantees that every database copy is always consistent and correct. Aiding this is the Shadowbase technology's superior capabilities to automatically detect and resolve data collisions.

Furthermore, the Shadowbase technology's online copy utility, the Shadowbase AutoLoader, allows 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.

The Shadowbase technology is fully heterogeneous. The Shadowbase active/active technology can be used on today's popular systems including UNIX, Linux and Windows running Oracle and SQL Server databases, as well as NonStop servers using NonStop SQL and Enscribe databases. Contact us for the availability of additional commercial offerings, such as Sybase, MySQL, and 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

The large bank service bureau mentioned in the introduction needed to provide financial services for other banks without any interruption of service. Focused on the need for disaster tolerance, the bureau turned to an active/active 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.

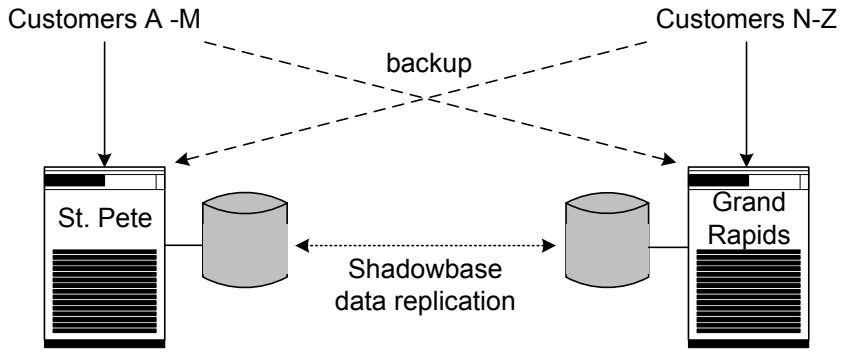

To avoid collisions, the database is logically partitioned as follows. Each customer bank is assigned to one of the nodes as its primary node. Therefore, all updates that are made to the database for a

particular bank are always made at one node, thus avoiding data collisions. The service bureau can dynamically reassign customers to nodes to provide load balancing.

Should a node fail, 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.

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. Furthermore, with the variability in call activity, it needs to be able to instantaneously balance the load between its two nodes. To accomplish these goals, it implemented an active/active NonStop system using the Shadowbase tool set. Shadowbase software is used to keep the Milan and Rome databases synchronized.

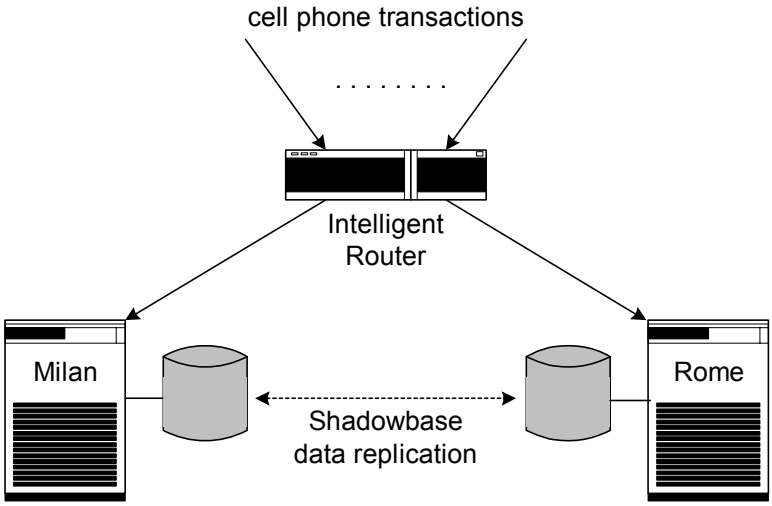

Transactions are routed according to a load balancing algorithm. When a transaction is received by the system, an intelligent router decides to which node to route the transaction based on the current load of each node. Either node may process any transaction. The system handles over 1,000 transactions 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 Shadowbase replication engine ensures the integrity of the database in the presence of numeric collisions.

Large Travel Agency

A major travel agency found that its 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.

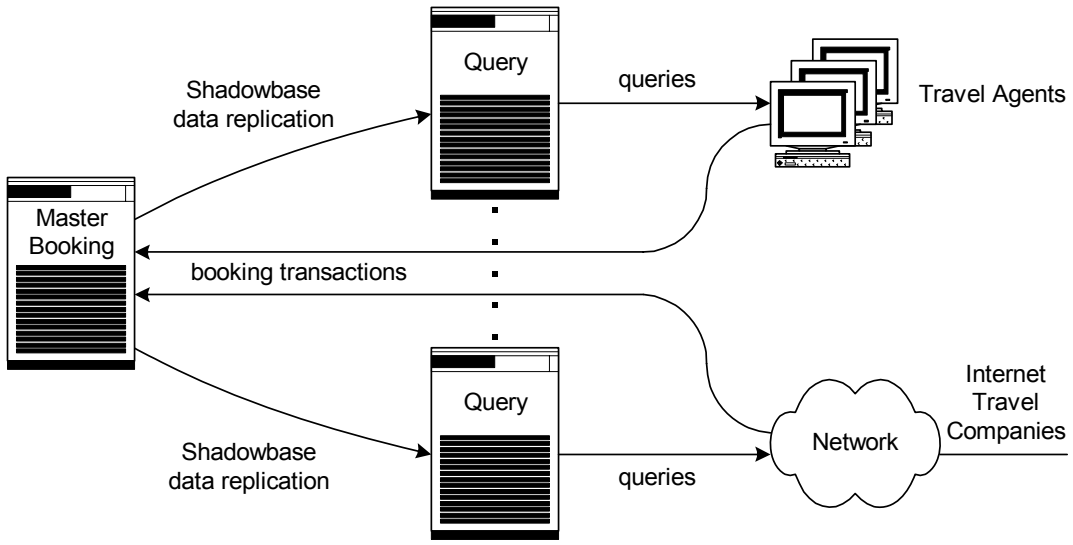

To accommodate these wide swings in system activity, the travel agency turned to active/active technology to offload query processing from the booking system by creating a NonStop master booking node and a series of NonStop and Windows query nodes. The master booking node maintains the master database. 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.

Query nodes can be added or removed as load dictates. Should a query node fail, the other query nodes can still support all of the users. Furthermore, recovery from a master booking node failure can be achieved by promoting one of the NonStop query nodes to master.

Want to Learn More?

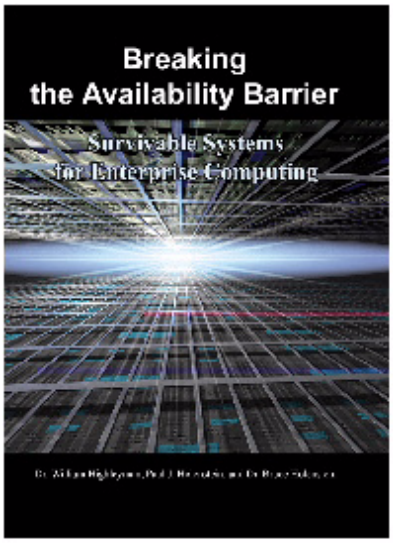

The underlying theory supporting the extreme availability of active/active systems and the techniques for building these systems is extensively covered in the book "*Breaking the Availability Barrier: Survivable Systems for Enterprise Computing*," published by AuthorHouse. Look for the second volume in this series, "*Breaking the Availability Barrier: Achieving Century Uptimes with Active/Active Systems*," to be available in the summer of 2006. These books can be ordered from www.authorhouse.com or www.amazon.com.

For further information about Shadowbase solutions and how they enable active/active replication architectures or are used to solve other business problems such as feeding a data warehouse with real-time data, contact Gravic at +1.610.647.6250, Shadowbase@gravic.com, or visit our website at www.gravic.com.