



**Achieving Century Uptimes with  
Shadowbase<sup>®</sup> Active/Active Technology  
White Paper**

Gravic, Inc.  
301 Lindenwood Drive, Suite 100  
Malvern, PA 19355-1772, U.S.A.  
+1.610.647.6250  
[Shadowbase@gravic.com](mailto:Shadowbase@gravic.com)  
[www.gravic.com/shadowbase](http://www.gravic.com/shadowbase)

## 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, or 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/shadowbase](http://www.gravic.com/shadowbase)).

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, rather than all, of the users are affected upon a node 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,
- Utilize all purchased nodal capacity for productive work (no passive standby systems needed),
- Eliminate the uncertainty associated with passive backup systems at takeover time (if my primary fails, will the backup system take over and work *properly?*),
- 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. This 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.*

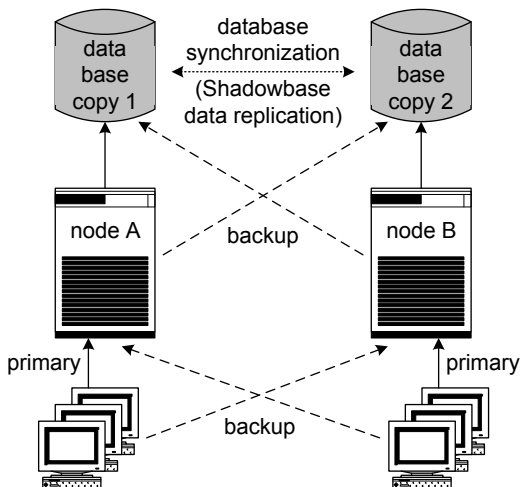
In the most general case, the nodes are completely symmetric. Any transaction can be routed within the application network to any node which can read or update any set of data items in the database. This approach provides the most flexibility and maximizes system investment as requests can be load-balanced 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 additional cost. Should a disaster take out a node or a database copy, there are others in the network 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, 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*.

Regardless of the type of failure, 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 switchover affects all users, and hence usually involves the approval of upper level management (which may be hard to quickly obtain). In an active/active system, these types of failures only affect the users on that node or database, not the entire user population. And since other nodes exist in the network that are known to be working, these users can be quickly switched to an alternate node.

Active/active systems eliminate the uncertainty that always exists when an active/backup approach is in place. In an active/active system, when a failure occurs there is no massive failover leap-of-faith surrounding the take-over on a backup system; all nodes in an active/active network are always 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.

Active/active architectures allow for all purchased nodal processing capacity to be actively working on satisfying user requests of any type (e.g. read and/or update). There is no backup, or passive standby, system sitting idly by waiting for another component to fail. All nodes are actively performing real work, all of the time.

Another advantage of active/active systems is that you can locate the nodes in the active/active network near your user populations. For example, if you have an east coast and a west coast (USA) set of user populations, you can locate one node near the east coast user population, thereby reducing their system access communications latency, and locate another node near the west coast user population, thereby reducing their system access communications latency as well.

## **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 that provides bi (or multi)-directional replication between the database copies and guarantees that all copies remain in a consistent and correct state.

For initial target database loading, Shadowbase provides the Shadowbase AutoLoader and the new Shadowbase Online Loader (SOLV). These powerful utilities can dynamically synchronize the target database to the source database while the applications are running on each node, updating the database. These patented 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 Shadowbase AutoLoader and the SOLV utilities, the Shadowbase tool set provides the critical elements 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 their 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 different and both are wrong. The Shadowbase software can detect collisions and can automatically resolve them in many cases. And, for those cases where Shadowbase cannot automatically resolve a collision, Shadowbase supports embedding customer business logic into the replication engine to take whatever action is necessary to resolve the collision.

## **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 any 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 (ZDMs). Its online copy facilities, the Shadowbase AutoLoader and SOLV, 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. 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 (i.e., in case you need to fail back).

## **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 perhaps 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 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, and the high-speed SOLV utility, 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, SQL Server, and Sybase databases, as well as NonStop servers using NonStop SQL and Enscribe databases. Nodes in the Shadowbase active/active network can be similar, or completely different. They can be running the same software, or completely different versions; there is no need for each node to be configured the same and to run the exact same versions, as Shadowbase handles interconnecting differing versions. Contact us for the availability of additional commercial offerings, such as 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 subscriber 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.

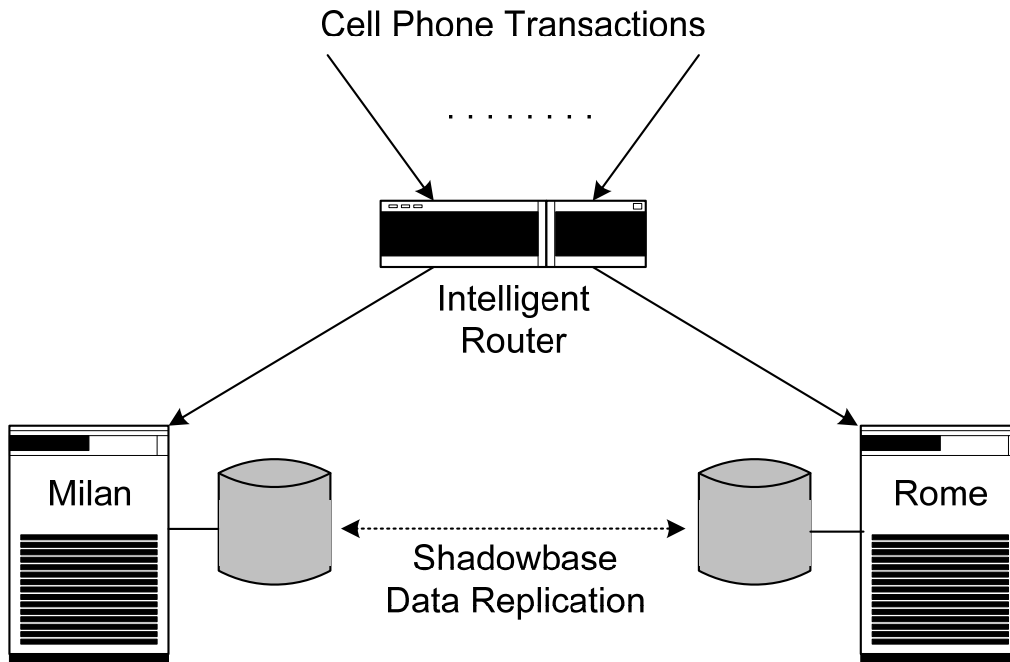


To avoid collisions, the database is logically partitioned as follows. Each customer bank is assigned to one of the nodes as its primary node. Therefore, all updates that are made to the database for a particular bank are always made at one node, thus avoiding data collisions. The service bureau can 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.



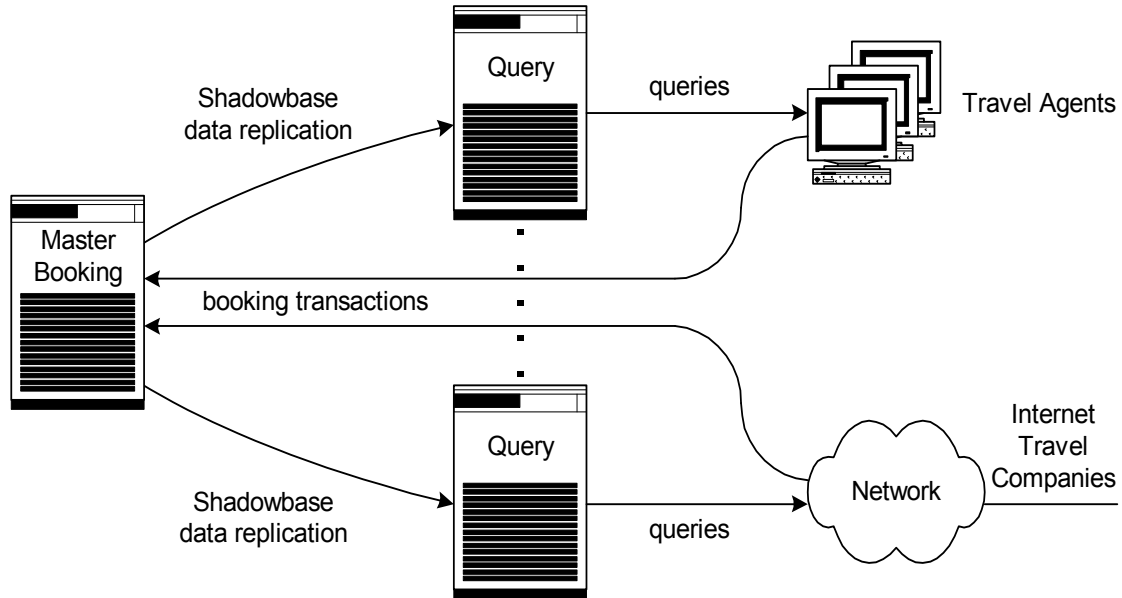
Transactions are routed according to a load balancing algorithm. When a transaction is received by the system, an intelligent router decides to which node to route the transaction based on the current load of each node. Either node may process any transaction. The system handles over 1,000 transactions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

## Want to Learn More?

### Breaking the Availability Barrier Series



[Breaking the Availability Barrier, Survivable Systems for Enterprise Computing](#). This book describes the techniques that can be used today for extending the time between system failures from years to centuries, often at little or no additional cost.

[Breaking the Availability Barrier, Volume II, Achieving Century Uptimes With Active-Active Systems](#) and [Breaking the Availability Barrier, Volume III, Active/Active Systems in Practice](#).

Volume II extends the theoretical topics covered in Volume I and covers significant issues and technologies related to active/active architectures, including availability calculations, redundant reliable networks, distributed databases, TCO, and node failures. Volume III is really a companion to Volume II and contains practical examples and case studies using Shadowbase for actually building active/active systems based upon the principles covered in Volumes I and II.

The authors of this series, Dr. Bill Highleyman, Paul J. Holenstein, and Dr. Bruce D. Holenstein, have a combined experience of over 90 years in the implementation of fault-tolerant, highly available computing systems. This experience ranges from the early days of custom redundant systems to today's fault-tolerant offerings from HP (NonStop) and Stratus. To learn more about the authors, [click here](#).

These books can be ordered from [www.authorhouse.com](http://www.authorhouse.com) or [www.amazon.com](http://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](mailto:Shadowbase@gravic.com), or visit our website at <http://www.gravic.com/shadowbase>.

## About Gravic, Inc.

Gravic, Inc. is a privately-held corporation located in Malvern, PA USA, about ten miles west of Philadelphia, PA USA. We make a number of computer software products for data collection, transformation, and distribution uses. Customers for our Enterprise products include the data centers of brokerages, travel agencies, stock exchanges, electric utilities and telecommunications companies.

---

### Copyright and Trademark Information

This document is Copyright 2008 by Gravic, Inc. Any reproduction is prohibited without the express written consent of Gravic, Inc. Gravic, Shadowbase and Total Replication Solutions are registered trademarks of Gravic, Inc.. All other brand and product names are the trademarks or registered trademarks of their respective owners. Specifications subject to change without notice.

## Appendix 1

### ***Shadowbase Supported Platforms and Databases***

This list is constantly being updated. Contact us at [SBSupport@gravic.com](mailto:SBSupport@gravic.com) for availability of additional database and operating system support.

#### **Standard Source Databases**

<b>Source Databases</b>	<b>Operating Systems</b>
HP Enscribe	HP NonStop Kernel (Guardian)
HP NonStop SQL/MP	HP NonStop Kernel (Guardian)
HP NonStop SQL/MX v1.8-2.1 (SQL/MP Tables)	HP NonStop Kernel (Guardian)
HP NonStop SQL/MX v2.2 (Native ANSI Tables)	Contact Support for Availability
Oracle 8i, 9iR2, 10g, 10gR2	Windows 2000/2003/XP Sun Solaris IBM AIX HP-UX PA-RISC & Itanium 2 HP Tru64 UNIX SuSE Linux Red Hat Enterprise Linux
Microsoft SQL Server 2000, 2005	Windows 2000/2003/XP
Sybase ASE 12.5+ (Needs Sybase Replication Server 12.6+)	Windows 2000/2003/XP Red Hat Enterprise Linux

#### **Standard Target Databases**

<b>Target Databases</b>	<b>Operating Systems</b>
HP Enscribe	HP NonStop Kernel (Guardian)
HP NonStop SQL/MP	HP NonStop Kernel (Guardian)
HP NonStop SQL/MX v1.8-2.1 (SQL/MP Tables)	HP NonStop Kernel (Guardian)
HP NonStop SQL/MX v2.2 (Native ANSI Tables)	Contact Support for Availability
Oracle 8i, 9iR2, 10g, 10gR2	Windows 2000/2003/XP Sun Solaris IBM AIX HP-UX PA-RISC & Itanium 2 HP Tru64 UNIX Open VMS SuSE Linux Red Hat Enterprise Linux v. 3-5 HP Open VMS

Microsoft SQL Sever 2000, 2005	Windows 2000/2003/XP
Sybase	Windows 2000/2003/XP Sun Solaris IBM AIX HP-UX PA-RISC & Itanium 2 HP Tru64 UNIX SuSE Linux Red Hat Linux

Contact Shadowbase Support at [SBSupport@gravic.com](mailto:SBSupport@gravic.com) for specific version compatibility for your environment or if your environment/platform is not listed above.

## Appendix 2

### Contact Information

#### Gravic, Inc.

301 Lindenwood Drive  
Suite 100  
Malvern, PA 19355-1758  
USA  
Tel: +1.610.647.6250  
Fax: +1.610.647.7958

Email Shadowbase Sales:  
[Shadowbase@gravic.com](mailto:Shadowbase@gravic.com)

Email Shadowbase Support:  
[SBSupport@gravic.com](mailto:SBSupport@gravic.com)

Website:  
<http://www.gravic.com/shadowbase>

#### International Distributors

##### Europe

#### BrightStrand International Limited

Tumberry House  
175 West George Street  
Glasgow  
United Kingdom G2 2LB  
Tel: +44 (0) 141 204 4046  
Fax: +44 (0) 7831 775091  
[mcrosby@brightstrand.com](mailto:mcrosby@brightstrand.com)  
<http://www.brightstrand.com>

#### CS Software Concepts & Solutions GmbH

Schiersteiner Strasse 31  
65187 Wiesbaden  
Germany  
Tel: +49 611 890 85 33  
Fax: +49 611 890 85 43  
[support@cs-software-gmbh.de](mailto:support@cs-software-gmbh.de)  
<http://www.cs-software-gmbh.de>

##### South Africa

#### Business Connexion (PTY) Ltd

Business Connexion Park  
789 16th Road  
Randjiespark  
Midrand  
South Africa  
Phone: +27 (0)11 653 5365  
[johannes.may@bcx.co.za](mailto:johannes.may@bcx.co.za)  
[www.bcx.co.za](http://www.bcx.co.za)

#### NSK Systems Services (Pty) Ltd

Millennia Park  
Stellentia Avenue  
Stellenbosch  
7600, South Africa  
Tel:+27 82 65 00 991  
Fax:+27 82 131 65 00 991  
[info@nsk-systems-services.com](mailto:info@nsk-systems-services.com)  
[www.nsk-systems-services.com](http://www.nsk-systems-services.com)

##### Asia Pacific

#### BRIGHTSTRAND ASIA-PACIFIC LTD.

25 Toby's Way  
Waikanae Beach  
New Zealand 6010  
Phone: +64 4 570 3330  
Fax: +64 4 570 3339  
[a.dewsnap@brightstrand-ap.com](mailto:a.dewsnap@brightstrand-ap.com)  
<http://www.brightstrand-ap.com>

#### Legato Solutions & Services Pte Ltd

101 Geylang Lorong 23  
#05-01 Prosper House  
Singapore 388 399  
Tel: +65 8463 150  
Fax: +65 8411 530  
[tg.sim@legato.com.sg](mailto:tg.sim@legato.com.sg)  
<http://www.legato.com.sg>