



Vodacom's One-Year Recovery

A Gravic, Inc. Case Study



Vodacom's One-Year Recovery

Executive Summary

Vodacom (Pty) Ltd. is one of the largest cellular telephone service providers in Africa, currently supplying voice and messaging services to over 130 million customers. Prepaid calling cards are a major Vodacom service. If the service was unavailable, it would cause much of Africa's cellular traffic to come to a halt as subscribers exhaust their cellular minutes. Therefore, Vodacom uses HPE NonStop server pairs to provide prepaid calling card services via its Prepaid Front End (PPFE). To minimize PPFE outages and to ensure multi-second recovery times, the HPE NonStop server pairs implement HPE Shadowbase bi-directional *Sizzling-Hot-Takeover* (SZT) data replication by Gravic, Inc.¹



Vodacom implements and sells the PPFE solution in Africa as a complete hardware, software, and application solution. For example, they have installed the systems and PPFE application for various entities in Lesotho, The Democratic Republic of the Congo (DRC), and Tanzania. In 2011, Vodacom installed an HPE NonStop PPFE Shadowbase SZT architecture in Tanzania. In August, 2013, faulty batteries resulted in a fire, which downed Tanzania's production HPE NonStop PPFE system. It took until the next day for Vodacom to switch PPFE operations to its backup system because of issues experienced with the Online Charging System/Intelligent Networking (OCS/IN) platform.

The explosion and resulting fire severely damaged the original production server, but Vodacom and HPE were able to restore it to service as the backup system in two weeks even though the damaged server demonstrated erratic behavior. For those first two weeks, the PPFE system ran without a backup. Vodacom then began the preparations for a new backup system in a different datacenter.

Many problems plagued Vodacom's effort. There was a protracted negotiation with the insurance company to get funding since the damaged server was still working at times. Once the new server was ordered, the air transport could not carry the NonStop cabinets standing up. The cabinets had to be dismantled and laid on their sides. Reinstalling units in the cabinets was a delicate, time-consuming job.

After the new server was brought online, the process of copying the database from the production system to the new backup system began. The database, twelve terabytes in size, primarily consisted of HPE NonStop SQL/MP tables using HPE NonStop SQL/MX aliases. Unfortunately, the versions of the operating system and SQL/MX databases were different in the two systems and led to the failure of the NonStop export/import utilities to transfer the HPE NonStop SQL/MX tables properly. Other techniques had to be employed for the SQL/MX database load. Eventually, standard SQL/MP utilities were used for the bulk of the data and for the rebuilding of index files for the twelve-terabyte database. During the setup of the new P9500 disk storage LUN configuration, issues were experienced with the cabling causing the Volume Paths to fail. HPE had to redo the LUN setup and the Naming of the Volumes exercise for the application database, which added to the serious delay of creating the tables and loading the data.

The backup system was finally brought online in July, 2014, almost one year after the original fire and explosion. The new system is now the production system, and the older system is the backup (these roles switch when a takeover is performed). The production and backup systems are being taken down one at a time to rebalance their databases to get rid of hot spots (service is maintained when the production system is taken down for planned maintenance by performing a takeover on the backup system). Fortunately, during the previous year, the fault-tolerant NonStop production server experienced no problems that required an unplanned failover to the erratic backup system.

This case study provides insight into designing advanced business continuity architectures, and the lessons learned when even this amount of extensive preparation is tested by a challenging environment and Murphy's Law² sequence.

¹For more information, please read the Gravic white paper: [Choosing a Business Continuity Solution to Match Your Business Availability Requirements](#).

²Murphy's Law is an adage or epigram that is typically stated as, "Anything that can go wrong, will go wrong."

Table of Contents

Executive Summary.....	2
Vodacom's One-Year Recovery	4
The Tanzanian PPFE Configuration.....	4
Call Handling	4
Adding Minutes	5
Continuous Availability for the PPFE	5
The Production PPFE System is Downed by an Explosion	6
Disaster Strikes	6
The Initial Recovery	7
The Long Database Recovery	7
The Erratic Backup System.....	7
Replacing the Backup System	8
Bringing Up the New Backup System	8
Summary.....	10
International Partner Information.....	11
Gravic, Inc. Contact Information	11

Table of Figures

Figure 1 – PPFE Architecture.....	4
Figure 2 – Approximate Locations of HPE NonStop PPFE Servers	6
Figure 3 – Recovering \TANZKW as the New-Hot-Standby.....	8

Vodacom's One-Year Recovery

Vodacom (Pty) Ltd. is one of the largest cellular telephone service providers in Africa. Vodacom has expanded from its roots in South Africa to include cellular telephone networks in Tanzania, Lesotho, and the Democratic Republic of the Congo (DRC). It currently offers voice and messaging services to over 130 million customers. Vodacom is owned by Vodafone, the world's second largest mobile phone company (behind China Mobile).



Prepaid calling cards, the fastest growing cellular option in Africa, are a major Vodacom service. If this service is unavailable, much of Africa's cellular capability comes to a halt as subscribers exhaust their cellular minutes. Therefore, Vodacom uses HPE NonStop server pairs to provide prepaid calling card services via its Prepaid Front End (PPFE). The PPFE, primarily a NonStop OSS application, is used by customers to recharge their accounts. If the PPFE fails, subscribers cannot add money to their cellular accounts; if subscribers exhaust their accounts, they no longer have cell phone service. To minimize PPFE outages and to ensure multi-second recovery times, the NonStop server pairs implement HPE Shadowbase bi-directional *Sizzling-Hot-Takeover* (SZT) data replication to ensure multi-second recovery times.

As within its other geographical areas of service, Vodacom provides an HPE NonStop PPFE Shadowbase SZT architecture in Tanzania. However, in August 2013, a battery explosion downed Tanzania's production NonStop PPFE server. Vodacom was able to switch PPFE operations to its backup system in just a few minutes; however, problems with the OCS/IN platform restricted full operational capability of the remote (backup system) until the next day. Even worse, a year later, the backup capabilities for this installation were still being brought back into service. This case study investigates the causes of the lengthy failure/recovery cycle as a sober lesson for other mission-critical businesses that are implementing business continuity architectures.

The Tanzanian PPFE Configuration

Call Handling

A simplified overview of how cell calls are handled is shown in Figure 1. In this figure, cell phone subscribers connect to the Vodacom cellular network via Intelligent Network (IN) systems. Each subscriber is assigned an IN based on the first digits of his telephone number. Each cell phone tower connects to an IN. Various information is collected about a cell phone user (such as the number of minutes he has left in his account, and the geographical location of the subscriber).




Figure 1 – PPFE Architecture

When a subscriber turns on his phone, his mobile signal is picked up by the closest cell phone tower. The IN to which the cell phone tower is connected determines the subscriber's assigned IN from his mobile- phone number and notifies the assigned IN of the subscriber's location. The subscriber is connected to an IN, which also obtains the subscriber's account data from his assigned IN.

When a subscriber places a call, it is handled by the connected IN, which contacts the assigned IN of the called party to determine the location of that party and then routes the call to the connected IN of the called party. The call has now been established.

When the call has been completed, both the IN handling the calling party and the IN handling the called party update their respective subscribers' account data (for instance, deducting minutes from the subscribers' accounts of prepaid minutes). The updated data is then sent to the subscribers' assigned INs to update their databases.³

Adding Minutes

In order to add airtime to a prepaid cellphone account, a subscriber has a number of choices:

- *Purchase a voucher at a store*
To activate this airtime, the subscriber has to redeem the voucher by sending the unique number printed on the voucher to the Vodacom PPFE system using Unstructured Supplementary Service Data (USSD).⁴ The PPFE system then performs a lookup in its database to get the amount of airtime associated with this unique number from the voucher and applies it to the subscriber's IN.
- *Purchase airtime from a Bank ATM*
This transaction is sent to the Vodacom PPFE system via a standard financial format message (called an ISO8583 message). The PPFE system verifies the request by using data in its database and applies the correct airtime to the subscriber's IN account.
- *Purchase airtime from a street vendor*
Street vendors have Vodacom accounts (similar to bank accounts) that are stored on the Vodacom PPFE system. The vendor uses USSD to interact with the PPFE system and the PPFE applies the airtime to the subscribers IN account when the vendor completes the transaction.

Continuous Availability for the PPFE

If the PPFE experiences an outage, subscribers will not be able to add minutes to their accounts. Cell phone service will become unavailable to subscribers once their balance is exhausted. Therefore, the PPFE is implemented as a dual NonStop server pair using bi-directional HPE Shadowbase SZT, where one system is the production system and handles all transactions (see Figure 1). However, the backup system is actively running the PPFE application and is ready to take over in seconds if the production system fails. This method is called cross region replication; its database is kept synchronized with the production system by HPE Shadowbase data replication.⁵

HPE Shadowbase data replication is accomplished in sub-second times so that the backup database is synchronized with the production database. If the production system fails, then rerouting transactions to the backup system is all that needs to be done. By configuring bi-directional replication, any changes made on the promoted system queue are delivered to the original system to resynchronize its database once it is recovered. The backup system can be assessed continually with test transactions against test accounts to ensure that it is functional and working end-to-end. This assessment is quite useful since validating the backup system's application processing can be done at any time without requiring a production application outage. This testing helps to ensure the backup system will be able to take over processing without any problems if the production system fails.

³The actual details of how calls are handled are much more complex than described here, which is just a simple overview for the purposes of this case study.

⁴USSD is a protocol used by cell phones to communicate with the service provider's computers.

⁵For more information about HPE Shadowbase and the sizzling-hot-takeover business continuity architecture, please read the Gravic white paper: [Choosing a Business Continuity Solution to Match Your Business Availability Requirements](#).

On behalf of the customer, Vodacom installed one NonStop PPFE system in Oyster Bay, Kwale Street, Tanzania, and the other in Mbezi, Tanzania (Figure 2). The two systems are about 30 kilometers apart, situated on high ground to address the primary cause of disasters in Tanzania – tsunamis. The systems are managed for the customer by Vodacom personnel located in a remote office connected to the systems via Telnet links.

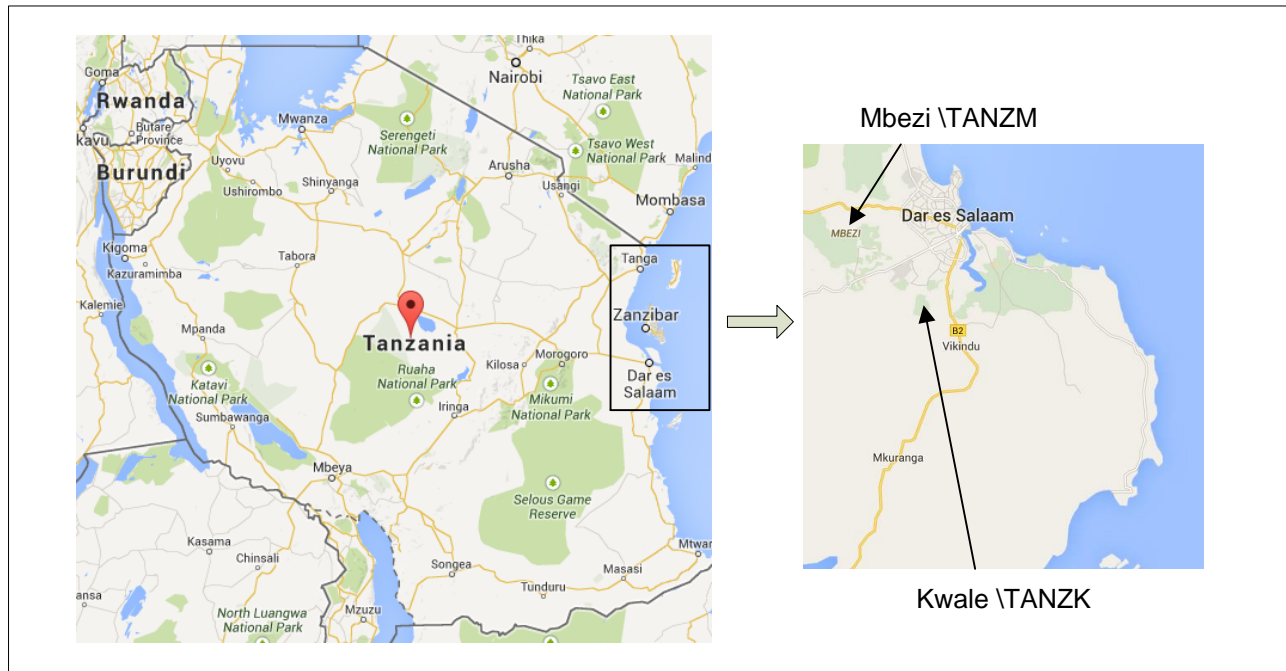


Figure 2 – Approximate Locations of HPE NonStop PPFE Servers

The node name for the HPE NonStop system in Kwale is \TANZK, and the node name for the NonStop system in Mbezi is \TANZM. Node \TANZK was the production node and HPE Shadowbase architecture replicated its database changes to its SZT backup node, \TANZM.

The IN systems at the backup site are typically down, and require a restart sequence to bring them up in the event of a failover. This means that during normal processing at Kwale, all IN systems are connected to the production node in Kwale. IN access to the backup site in Mbezi is provided by an Unstructured Supplementary Service Data (USSD) gateway.

The Production PPFE System is Downed by an Explosion

Disaster Strikes

On Friday afternoon, August 16, 2013, disaster struck. A UPS battery exploded and caught fire in the Kwale datacenter. The explosion, along with the fire-suppression system, damaged the \TANZK PPFE production node, taking it out of service. The sudden outage of \TANZK caused significant database corruption, as updates in progress could not be completed.

Even worse, the IN systems were in the same datacenter as the \TANZK PPFE node, and their local database was corrupted by a “hard down” (meaning it was inoperable) when the failure occurred. Their communication links with the backup PPFE server in Mbezi were also damaged.

Failover to the Mbezi IN systems failed, with the Mbezi IN systems being unable to complete the DB startup sequence and allow access to the IN customer database application. Later, it was learned that this was caused by a massive number of simultaneous farm node requests to connect to the Mbezi IN database coming in and overloading the environment, causing timeouts and retries.

If the IN system failover to the backup systems installed in the Mbezi datacenter had been successful, subscriber calls to redeem prepaid vouchers could have continued with little, if any, interruption. Instead, cell

phone requests to redeem prepaid vouchers for additional minutes could not be serviced by the backup PPFE system in Mbezi. Subscribers running out of minutes lost their cell phone service.

Lesson Learned: *Avoid single points of failure. All network, system, application, and data components necessary to provide service must be redundant, must exist in geographically separate locations, and need to be tested for failover periodically.*

The Initial Recovery

By the next morning, the \TANZK node was once again operational from a hardware viewpoint. However, none of its software had yet been restored. There were no startup scripts, no subsystems, and no applications running. An attempt was made to start TMF, but it immediately ran into problems. Since the database had been left in an inconsistent state from the sudden node outage, TMF could not recover the database from the Audit Trail. The database had to be recovered from its backup, provided by ETI-NET's virtual tape facility, BACKBOX.⁶

By noontime Saturday, communication between the IN systems in the Kwale datacenter was established with the PPFE backup in the Mbezi datacenter via the USSD gateway system installed in Kwale. The USSD gateway ran on a Windows server that had not been damaged by the explosion. By this time, the Tanzanian prepaid voucher redemption service was offline for over a day.

The Long Database Recovery

The database backup strategy is that a full backup (via a TMF online dump) is done over the weekend with incremental backups (via dumping the generated audit trail files) during the week. Therefore, the database had to be restored, followed by rolling forward thru nearly a full week's worth of Audit Trail change data. This recovery was accomplished via a TMF *Recover Files* operation for the corrupt database partitions, and it had to be completed before the database could be made available for application access. The database comprised HPE NonStop SQL/MP tables with SQL/MX aliases. The database recovery from the online dump was completed by Saturday afternoon. However, it took an additional two weeks to complete the Audit Trail roll forward recovery.

During this time, even though the PPFE service was restored to subscribers, the production node, now in Mbezi (the original backup node) had no backup, and a failure of that system would have again rendered the PPFE service unavailable. When database recovery was finally completed, the Tanzanian PPFE system was restored to an active/passive SZT configuration with production accomplished by the \TANZM node and backup services provided by the recovered \TANZK node.

Lessons Learned:

- 1) *Run backups more than once per week – perhaps every few hours if you are using virtual tape. Otherwise, to recover the database, you may need to restore considerable amounts of data as you roll the database forward from the last backup point, which can take a long time, extending the service outage.*
- 2) *Parallelize the backup and restore operations as much as possible (and test the recovery).*
- 3) *Consider whether a single backup system/site is sufficient, since an extended outage of the primary system/site will leave you vulnerable to another prolonged outage if the backup system/site also fails.*

The Erratic Backup System

However, even though the PPFE was now running in an active/passive SZT mode, the backup node in Kwale proved to be somewhat unstable. Vodacom made many service calls to HPE, who was very helpful, but the service calls were for a backup system and not for a production system and hence were not of the highest priority due to the Service Level Agreement (SLA) contract. The service contract allowed for longer service intervals for non-production system issues, and it sometimes took hours for the problem to make its way to the front of the queue as support “followed the sun.” However, no matter what the time of day or night, HPE was very helpful and consistently went beyond its contractual obligations under its support agreement.

⁶For more information about the BACKBOX product line, see ETINET.com.

Lesson Learned: Carefully consider the support SLA's required for backup systems for mission-critical applications.

Vodacom finally asked HPE to visit onsite and verify that the Kwale node hardware was in good condition. HPE performed exhaustive tests on the hardware and determined that the server was unusable due to smoke, water, and other fire-related damage.

Replacing the Backup System

Vodacom turned to its insurance company to obtain funding for purchasing a new backup node. The insurance company at first balked because the damaged node was still functioning, albeit erratically. However, after an extended negotiation, Vodacom convinced the insurance company to finance a new node.

Lesson Learned: If you do not have a total loss, be prepared to negotiate with the insurance company to receive a successful claim.

The new node funded by insurance was installed in a different building than the original \TANZK node. It was given the name \TANZKW. At the same time, Vodacom replaced its IN network with the IN Advantage system from Siemens. The multiple INs in the original configuration were replaced with just two of the more powerful Siemens' INs, freeing considerable space in the new datacenter.

Bringing Up the New Backup System

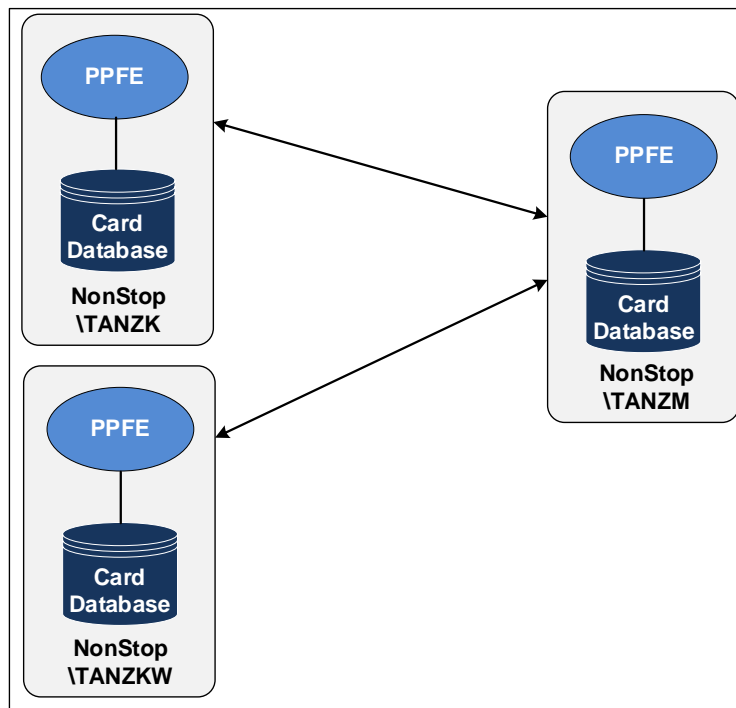


Figure 3 – Recovering \TANZKW as the New-Hot-Standby

As shown in Figure 3, bringing the \TANZKW node into service was the next task. The first problem was the shipping of the replacement HPE NonStop server to the Kwale datacenter. When the new hardware was shipped, the cabinets had to be stripped and reassembled on site because the cabinets could not be accommodated standing up on the airplane. Every single device was unracked and uninstalled from the cabinets in the factory, and the cabinets were shipped in the horizontal position. Devices such as servers, switches, and disk shelves were all separately packed and shipped on pallets. All of this shipping resulted in a much longer time to rebuild the new node and get the hardware online.

Once the \TANZKW hardware was operational, the node's subsystems and communications were installed. Vodacom personnel were onsite for six days bringing up the \TANZKW node. At this point, further restoration procedures for the node could be accomplished from the Vodacom remote office via Telnet.

Next, the support team began the migration of the PPFE application to the new node. They started by loading the \TANZKW databases from the \TANZM production system. Several challenges slowed this procedure to a crawl. The first problem was that the software versions on the older \TANZM node and the newer \TANZKW node were different. The original \TANZM and \TANZK nodes used the HPE NonStop operating system version J06.12, and the \TANZKW node used version J06.16. Furthermore, the original nodes used HPE NonStop SQL/MX version 2.3.4, and the new node used SQL/MX version 3.2. The result was that the standard NonStop export/import utilities would not copy the database directly from the source environment to the target environment.⁷

Lesson Learned: *Understand the limitations of different hardware and software versions on systems that must interoperate, and be prepared to address complexities*

The support team switched to HPE NonStop's UNPAK2 and PAK2 utilities to perform block moves via FTP for the SQL/MX tables. Macros were created to generate SQL LOAD jobs to run in parallel for the large SQL/MP tables. Because the production system, \TANZM, had to be quiesced to PAK the SQL/MX tables, the PPFE application was quiesced at midnight when it was handling a minimum load and these tables were copied during overnight periods. Each morning, before the application was restarted, HPE Shadowbase replication was initiated on the newly copied tables to keep them synchronized as the application updated them.

HPE Shadowbase SOLV Online Loader (SOLV) was employed to copy the smaller SQL/MP tables since it could replicate these tables while the application was running. It also kept these tables synchronized with the production system as the tables were being created and loaded. Thus, small-table migration proceeded continuously throughout the day.

Once the SQL/MP tables were loaded, their indices had to be built, which took an extraordinary amount of time. Because of the size of the tables and the amount of system resources available (disk and CPU), the indexing had to be undertaken one table at a time. Twelve terabytes of data had to be indexed, which took a total of 55 days. The support team worked long hours through March and into April to accomplish this task. Progress slowed considerably in April, which is vacation month in Tanzania. Progress was further gated by the lack of management authorization for the team to work overtime hours to finish the restoration task. All work had to be accomplished during normal office hours while the employees were already tasked with their regular support and systems administration jobs.

Lesson Learned: *Allocate additional budget and time for recovery operations after a failure to allow for the additional effort. Such factors should also be considered when developing or updating a business continuity plan, to ensure sufficient resource availability when it's needed, rather than after the fact when it only further slows the recovery process.*

Finally, after the database was fully loaded onto \TANZKW with all of the indices created and HPE Shadowbase replication keeping the database synchronized with the production database on \TANZM, the support team brought up the applications. To their horror, the disks immediately disconnected. It turned out that the logical unit locations for the various volumes were incorrect on ones that had previously been used on the older system's XP storage area network. The allocation on the new system was different. In addition, the team discovered that some of the disk cabling was wrong. It took another two weeks to fix these issues.

Once the database and disk storage were correctly configured, application environment issues arose with the OSS environment that affected the PPFE application, and further challenges with HPE NonStop SQL/MP and the SQL/MX aliases arose. These issues were finally resolved, and testing of the new \TANZKW node began in June 2014. Many remaining concerns were uncovered, and HPE had to return several times over a four-week period to resolve some remaining hardware/configuration issues.

⁷Although XP arrays were used on both systems to hold the database, it was not possible to directly copy the data from one XP array to the other.

Lesson Learned: *"It ain't over 'til it's over!"**In other words, even though you think the end is in sight, it often isn't...*

Finally, on July 10, 2014, nearly a year after the original explosion, \TANZKW was brought into service as the backup system for \TANZM, and \TANZK was shut down. During this extended outage window, the damaged system \TANZK had continued in an erratic fashion as the backup PPFE system. It even survived another fire that erupted in its datacenter in December, 2013. The fire caused further damage to the \TANZK system, but the support team was able to return it to service.

As of this writing, the systems have now been switched so that \TANZKW is the production system, and \TANZM is the backup system. The XP storage array used by \TANZM is currently being rebalanced to eliminate OSS hot spots that impose heavy loads on some disks in the array, which is necessary so that \TANZM will have the capacity to handle peak PPFE loads if a failover occurs.

Summary

Vodacom's PPFE is a critical system for its mobile services in Africa. The PPFE allows subscribers to top off their cell phones with purchased vouchers. Without the PPFE, cellular service will gradually grind to a halt as subscribers use up their minutes and cannot add additional minutes. To ensure continuous availability, the PPFE systems use HPE NonStop servers with HPE Shadowbase bi-directional SZT that provides data access failover times measured in seconds.

An explosion in August, 2013, nearly destroyed one of Vodacom's Tanzanian PPFE systems. It took Vodacom nearly a year to rebuild the damaged system using new hardware. Challenges included hardware problems due to smoke damage, software version issues, lengthy database loads, and limited staff hours. During this time, the PPFE system limped along as an active/passive SZT system with the damaged node providing erratic backup to the production node. Fortunately, during this period, the fault-tolerant HPE NonStop production node experienced no problems that required failover to the erratic backup system.

Though prepaid voucher services to Vodacom subscribers should have been restored in seconds by a rapid SZT failover following the explosion, PPFE services were lost for a day because the IN communication channels linking the production site and the backup site were damaged. The IN environment did not have a backup system; probably the single most important lesson of this entire saga is to eliminate single points of failure.

Lesson Learned: *If you want high availability, make sure that you can always failover to a backup environment quickly. Active/Passive environments can provide this level of availability. If you want continuous availability, the "backup" environment must be up and running and in a known-working state at all times. HPE Shadowbase Sizzling-Hot-Takeover and fully Active/Active environments can attain this level of availability. Pick your architecture to meet your business objectives.*

A series of unfortunate events caused a mission-critical production environment to fail, an outage that lasted over a day. The recovery of the failed environment took considerably longer than even the most pessimistic business continuity planner could have conceived. Failures happen — it is not a matter of *if*; it is only a matter of *when*. This Vodacom case study provides valuable insight into having a plan in place that avoids having a failure turn into a catastrophe. Consider whether any of the *Lessons Learned* might apply to your IT systems, and address them *before* disaster strikes.

International Partner Information

Global

Hewlett Packard Enterprise

6280 America Center Drive
San Jose, CA 95002
USA

Tel: +1.800.607.3567

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](#).

Copyright and Trademark Information

This document is Copyright © 2016, 2017, 2022 by 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.