ZERO NINES®
ALWAYS AVAILABLE™

ZeroNines® and CloudNines™: **Using the Cloud to Prevent Data Disasters**

ZERONINES: WHITE PAPER

www.zeronines.com

ZeroNines® and CloudNines™: Using the Cloud to Prevent Data Disasters

Ensuring Continuity and Preventing Outages Despite Disasters and Aging Infrastructure

Overview: The Need for Preventing Downtime

The cloud is well known for cutting costs and improving efficiency. But few know that with the correct cloud-aware technology and configuration it also offers an unparalleled capacity for preventing downtime by keeping applications, data, and services from being knocked offline by unfortunate events. This is particularly important regarding aging infrastructure that is at risk due to old hardware, unsupported applications, and outdated architecture.

Datacenter and cloud outages happen all the time despite the best efforts of IT departments and equipment manufacturers. Yet the demands on government and business systems are such that any outage at all can cause irreparable harm. Downtime threatens the successful handling of monetary transactions, emergencies, medical processes, just-in-time business models, and so forth, and puts both revenue and human lives at risk. Downtime can cost even a moderately sized business multiple millions of dollars per hour. Thus, even though an event like a fire, storm, or equipment failure may be the initial cause, the real disaster is often the consequent downtime among networked applications, data, and services. It's no wonder that many companies are reinvesting cloud savings into reliability and service uptime.

Logically, if datacenter crashes are unavoidable but business must continue, then a method is required for bypassing the damaged datacenter and maintaining access to networked assets. No datacenter can be allowed to become a single point of failure. The solution is to make any given datacenter or cloud node expendable, so the loss of one does not become a disaster that impairs delivery of network service and assets.

Using the Cloud to Create Expendable Datacenters℠

Cloud-friendly CloudNines™ technology that makes individual datacenters or cloud nodes expendable is commercially available. CloudNines uses patented Always Available™ technology from ZeroNines Technology, Inc. to enable any datacenter, cloud, or cloud node to go offline without causing an outage. Business can continue as normal while the damaged node is repaired and brought back into the array.

The key is that without upgrading or replacing existing infrastructure, CloudNines can add multiple cloud nodes that will "failsafe" vital transaction processing. Each node in the array processes all application transactions and data equally and simultaneously. Remove a node for any reason – software glitch, maintenance, or natural disaster – and processing simply continues on the others.

Benefits of Always Available technology include:

- Improves reliability of applications and data stored in clouds and traditional datacenters, enabling uptime in excess of 99.999% (five nines).
- Protects against monetary losses caused by outages.
- Reduces customer attrition and brand damage caused by service disruptions.
- Increases the power of your cloud by synchronizing multiple private, public or hybrid clouds at multiple vendors.
- Facilitates migration to the cloud.
- Helps meet government regulations and industry standards.
- Extends the useful life of current technology assets.

An Example From Calgary

As we saw in Calgary on July 11, 2012, a single disastrous event can lead to widespread loss of business and municipal services. On that day, a transformer explosion knocked out data services at medical centers, took an IBM datacenter offline, and crippled some city government systems for two days or more including the 911 emergency system.[1]

Less-spectacular mishaps like this happen all the time, caused by faulty power systems, software failure, hardware failure, human error, natural disasters, everyday maintenance, and many other instigating events. These are compounded by failover- and backup-based disaster recovery (DR) systems that often fail to prevent the real disaster, which is the loss of data, network transactions, and network services.

It’s All About Disaster *Prevention*, Not Recovery

At the core of the downtime problem is the IT industry’s attitude toward disasters.

Think about the meaning of “disaster recovery” or “DR”. This standard industry term describes a mission-critical IT function: picking up the pieces after a failure, and getting systems going again. But these very words reveal a grave flaw in IT thinking: Disasters must be recovered from, after the damage has been done. Sadly, prevention is not part of the standard vocabulary. By expecting disasters to happen and being content with cleaning up afterward, the IT world is subject to immense costs that threaten businesses, government services, and human lives.

The Calgary explosion provides an excellent example of the structure of most outage-causing events:

1. There is an instigating event (the explosion/fire and consequent loss of power) and
2. There is a resultant data or business disaster (the outages and loss of services.)

We contend that there are far too many different kinds of disaster-causing events to ever prevent them all. Bad things will happen and datacenters will be knocked offline, period. This view is substantiated almost daily by reports of business and government IT outages.

However, we also contend that the resultant data or business disasters CAN be prevented. The optimal way is to deploy additional processing nodes in the cloud so that if one element of your network goes offline (whether it is your legacy systems, hosting partner, or cloud node) the others continue your business processes. This is the function of Always Available technology and CloudNines.

Won’t Failover and Backup Prevent a Data Disaster?

In a word, no. They are reactive. They occur after the disaster-causing event has happened, they often don’t work, and they are often rendered ineffective by the very catastrophe they are expected to rescue us from. They frequently even cause tertiary business disasters of their own.

Since 1989, we have asked countless Fortune 500 and Global 2000 clients what their level of confidence in their DR plans and recovery strategies is. Without exception, the response is “None, but this is all we have.” Yet we continue to see these household name companies invest billions in failover- and backup-based DR strategies knowing they are unlikely to work and that the outcome will probably be disastrous.

They do this because failover and backup are still seen as the leading DR paradigms. Although these outmoded recovery techniques will eventually restore a crashed network, the cost in lost business, lost productivity and potentially lost lives (in the case of medical and security systems) makes their low reliability unacceptable. When seen in that light, their continued use in high-stakes enterprise systems is rather shocking.

Failover

Failover architecture first appeared in the 1960s before "always on" business systems became the norm. Back then, outages did not carry the high price they do today. Failover is intended to switch computing on the fly from a primary system to a secondary system. The problems with failover are legion, but the key drawbacks include lost in-flight transactions, cascading application failures, secondary sites that are crippled by the same disaster, and the risk of corrupted data. The same risks occur during cutover from the secondary back to the primary after the disaster has run its course. Failover is still used because it is mistakenly thought to be the only viable paradigm for handling a computing disaster in progress.

Tape or Optical Backup

Backup to physical media like tapes or disks is still part of the standard recovery regimen. Although this holdover from the 1970s may be acceptable and even necessary for archiving and compliance, it is extremely problematic when trying to recover from a disaster. Key weaknesses include the loss of all data from the time of the last backup to the time of the disaster, high potential for data corruption, delays as backup tapes/disks are retrieved from the storage vault, failed restoration due to out-of-sequence or damaged media, and damage to or inaccessibility of the backup media because of the same disaster.

Other DR systems are basically new variations on failover and backup. Although some of these modernized techniques can significantly reduce the actual outage to a few seconds, many of the risks of data corruption, failed failover, and failed restoration still exist. Again, just look at the news for examples. The Amazon EC2 Cloud went down twice in June 2012,[2] and we can assume they used the most effective DR methods known to them.

DR in the Cloud

Any business or government entity that moves to the cloud will be trusting their disaster recovery to the cloud provider.* This can be naïve to the point of irresponsibility, unless ample due diligence is undertaken. The lower cost structure, elimination of responsibility for the hardware, and reassuring words from the cloud provider can lull unsuspecting business managers into a false sense of security.

* By "cloud provider" we mean any of a number of permutations. It may be a commercial cloud provider like Amazon or Microsoft, an array of cloud and hosting solutions providers in the case of a hybrid cloud, or even the entity's own IT staff if they are running their own cloud internally.

Cloud outages occur with frightening regularity so customers need to know exactly what kind of DR support to expect. The cloud provider may offer a variation on the failover paradigm to try to maintain continuity. They may offer multiple availability zones as Amazon does.[3] They may urge their customers to provide their own DR system. And as some companies have found out recently, it is a mistake to assume that your cloud provider will make backups for you:

> *"Amazon doesn't make any promises to back up data... The real issue is that many users are under the impression that their data is backed up… but in fact it isn't due to mismanaged infrastructure configuration."*

This was said in June 2012 by Cameron Peron, VP Marketing at Newvem, a cloud optimization consultancy that specializes in the Amazon cloud.[4]

Even if the provider offers a good service level agreement (SLA) that promises high uptime rates, they won't be responsible for your systems' reactions to periods of downtime. Poorly architected systems hosted in the cloud can be highly vulnerable to cascading application failures and data loss if continuity is interrupted at all. Their two-second outage may lead to your two-day disaster.

Those who use the cloud need to deploy systems architected to maintain uptime within the cloud, regardless of what the cloud dishes out.

How Always Available™ Technology Uses the Cloud to Protect Applications and Data

Always Available technology and the CloudNines product function in an altogether different and far more reliable way than failover and tape-based recovery. It enables all transactions, data exchanges, and other network activities to occur equally and simultaneously on multiple clouds and other datacenters. All clouds, cloud servers, and other servers in the array are hot, and all are active. There is no hierarchy, and consequently no single point of failure.

If a cloud, cloud provider, hosting provider or datacenter goes offline for any reason, all activities continue uninterrupted via the virtual applications on other clouds and other datacenters. There is no need for failover or recovery from tapes because continuity is maintained.

Here's What Happens:

1. Server calls are originated by networked clients: computers, email, PDAs, phones, business software, databases, etc.
2. Each transaction passes through the ZeroNines Smart Switch Nodes, which route it simultaneously to all clouds and other datacenters in a one-to-many (1:m) session.
3. Processing and storage take place equally in all locations. All clouds and servers are hot, and all are active. There are no primary or secondary clouds or servers.
4. All transactions and data exchanges are recorded, verified, and subjected to security measures.
5. Each location sends its responses back through the Smart Switch Nodes.
6. The Smart Switch Nodes cooperatively eliminate duplicate responses and return a single response to the client that originated it.

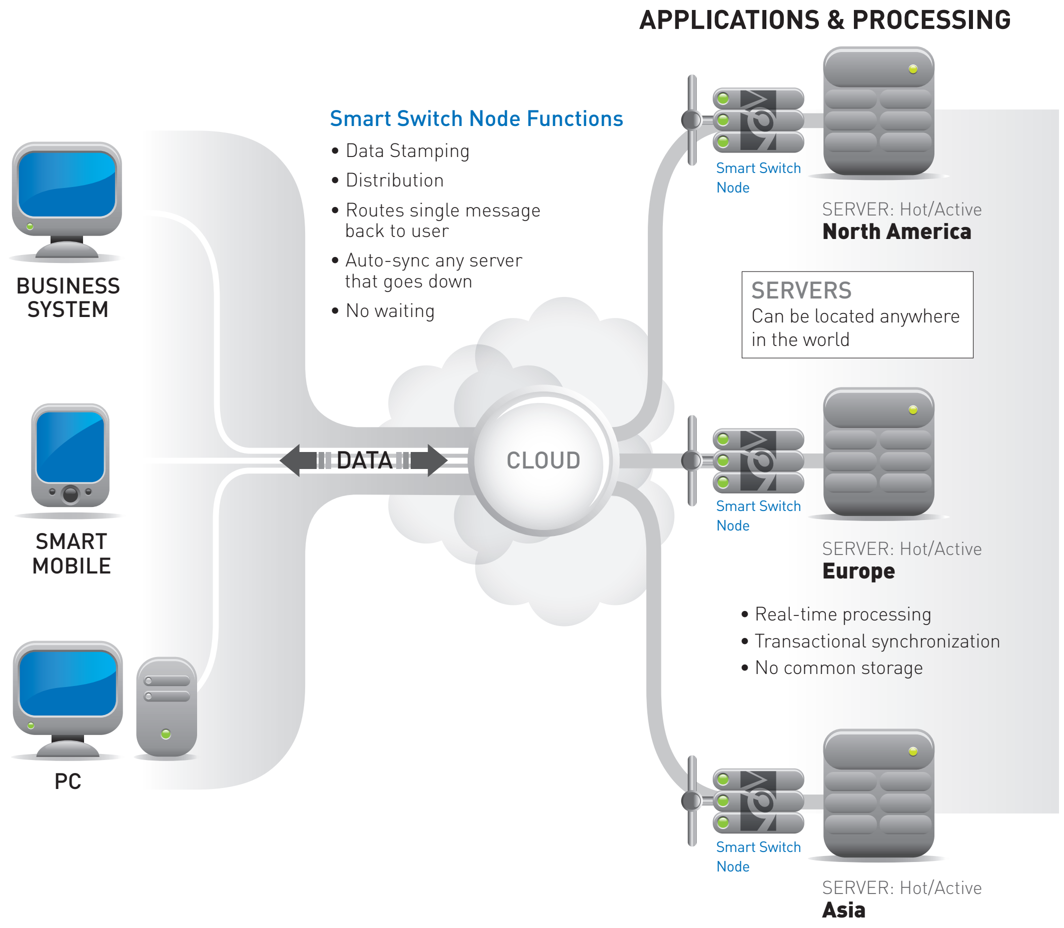

APPLICATIONS & PROCESSING

Smart Switch Node Functions

- Data Stamping
- Distribution
- Routes single message back to user
- Auto-sync any server that goes down
- No waiting

SERVERS
Can be located anywhere in the world

- Real-time processing
- Transactional synchronization
- No common storage

Safeguards are in place for guaranteed message delivery, data authorization, journaling, and synchronization to make sure that every transaction is secured, translated, completed, recorded, and communicated to the other networked clouds and servers. If one part of a network goes offline, the Journaling feature and Smart Switch Nodes automatically update it once it comes back online, enabling it to function at full capacity and to take over for the others if necessary. This effectively eliminates the need for tape or disk backup as a recovery method although it's prudent to continue to apply these methods for archival and SAS70 requirements.

Interoperability

Always Available architecture is agnostic regarding applications, operating systems, platforms, and cloud vendors, accommodating existing equipment and business methods. It is located at the transaction level within the network architecture, so there is no need to modify existing apps or data to make it work. Old and new hardware can be mixed and matched without compromising cloud integrity.

The ZeroNines Always Available architecture natively supports Web 2.0 and .NET strategies including HTTP, POP3, and SMTP protocols. A protocol interface development kit enables easy creation of interfaces for applications that use other protocols, even those unique to legacy systems.

Facilitating Cloud Migration

The typical migration to the cloud can be uncomfortably similar to failover. There comes a moment when processing switches from the older source datacenter to the target cloud. As with failover, there is the risk of incompatibilities, data corruption, and outages.

When Always Available technology is used as a migration tool there is no single do-or-die migration event. Multiple cloud nodes can be added at any time, and they can be configured, tested, and brought into the array whenever they are ready. The old source datacenter can continue to function as long as is seen fit, processing in tandem with the target cloud nodes. All are monitored for stability and proper function. If one or more nodes fail in some way, they can be removed from the array and be added back in once they are repaired. At some point the old source datacenter can be discontinued, leaving only the cloud nodes. Or the old datacenter could be retained indefinitely in a hybrid array, depending on business needs.

Case Study: Web Portal Startup ZenVault

ZeroNines client ZenVault® Medical **www.ZenVault.com** lets people store and manage their confidential personal health records online. Website reliability and uptime is of paramount importance because their customers' lives are literally at stake. ZenVault Medical launched in the cloud in September 2010 and also hosts at a colocation facility. All cloud nodes and the datacenter are part of an Always Available architecture.

Since launch, ZenVault Medical has maintained true 100% uptime, with no downtime for any reason including planned maintenance, upgrades, and other events that would have forced an ordinary website offline. When a network element fails or needs to be taken offline, ZenVault staffers remove it from the configuration, modify it as necessary, and seamlessly add it back into the mix once it is ready. ZenVault customers don't experience any interruptions.

Savings and Security Justify the Effort

The need to migrate aging or expensive systems to the cloud is becoming urgent. Unfortunately, this often conjures up a Catch-22: They must be migrated to avoid excessive costs and the risk of downtime, but the cloud itself sometimes seems as frail as the archaic hosting models it is intended to replace. IT planners have to be constantly ready for any given cloud node or datacenter to go offline.

But even with severe physical breakdowns it is unquestionably possible to prevent the data disaster that sends medical systems, business websites, emergency services, and other systems into the void. If the networked applications and data had remained fully available despite the explosion that day in Calgary – if their own outage had been prevented – then there would have been no data disaster.

In truth, the cloud is just a collection of datacenters running virtual machines. They are physical entities, subject to all the same risks as any other datacenter. Once this disappointing truth has been grasped, it is easy to see that the concerns over reliability and outages can be allayed simply by applying best practices developed specifically for the cloud.

We believe that the leading best practice for continuity in the cloud will be to dispense entirely with failover. Instead, businesses should adopt a technique that uses the cloud itself and its affordable computing capacity to virtually eliminate outages. The price of downtime events is highly unpredictable and extremely high, starting at millions of dollars per hour for even a moderately sized business. Exchanging this extreme risk for the known costs of cloud fees and disaster prevention makes a highly attractive cost-benefit equation. Many are already discarding old and inefficient hosting models, buying extra cloud capacity, and reinvesting the savings into reliability. Their strategy is to compete based on reliability and to eliminate the unknown but undoubtedly bitter cost of future outages.

[1] DatacenterDynamics http://www.datacenterdynamics.com/focus/archive/2012/07/transformer-explosion-knocks-out-hospital-ibm-data-centers-calgary

[2] Data Center Knowledge http://www.datacenterknowledge.com/archives/2012/06/29/another-outage-amazon-cloud/

[3] Wired.com http://www.wired.com/wiredenterprise/2011/11/amazon-in-orego/

[4] VentureBeat http://venturebeat.com/2012/06/21/using-cloud-services-40-of-you-arent-ready-for-the-next-outage/

ZERO NINES
ALWAYS AVAILABLE™

www.zeronines.com

For more information: info@zeronines.com

ZERO NINES
ALWAYS AVAILABLE

Corporate HQ
5445 DTC Parkway
Penthouse Four
Greenwood Village, CO 80111

T: (303) 814-8121
F: (303) 814-1495

ZeroNines® Technology, Inc. provides a new standard in network disaster recovery, shifting the paradigm from reactive recovery to proactive business continuity. Our Always Available™ information security and availability technology pushes application uptime beyond five nines (99.999%) to virtually 100% anytime, all the time – zero nines. This enables uninterrupted access to business data, applications, and transactions despite disasters or network disruptions that would otherwise cripple the enterprise. Always Available™ processes all transactions in parallel on geographically dispersed servers that are all hot and all active, eliminating single points of failure. It operates agnostically across multiple platforms, leveraging existing processing and storage infrastructure. We also offer enterprise infrastructure assessment, program management and project implementations. Founded in 2000 and based in Denver, Colorado, ZeroNines' primary target customer base includes Global 2000 companies.

Contact ZeroNines today to find out how your business can be Always Available!

info@zeronines.com
www.zeronines.com