The Disaster of Disaster Recovery

ZERO NINES

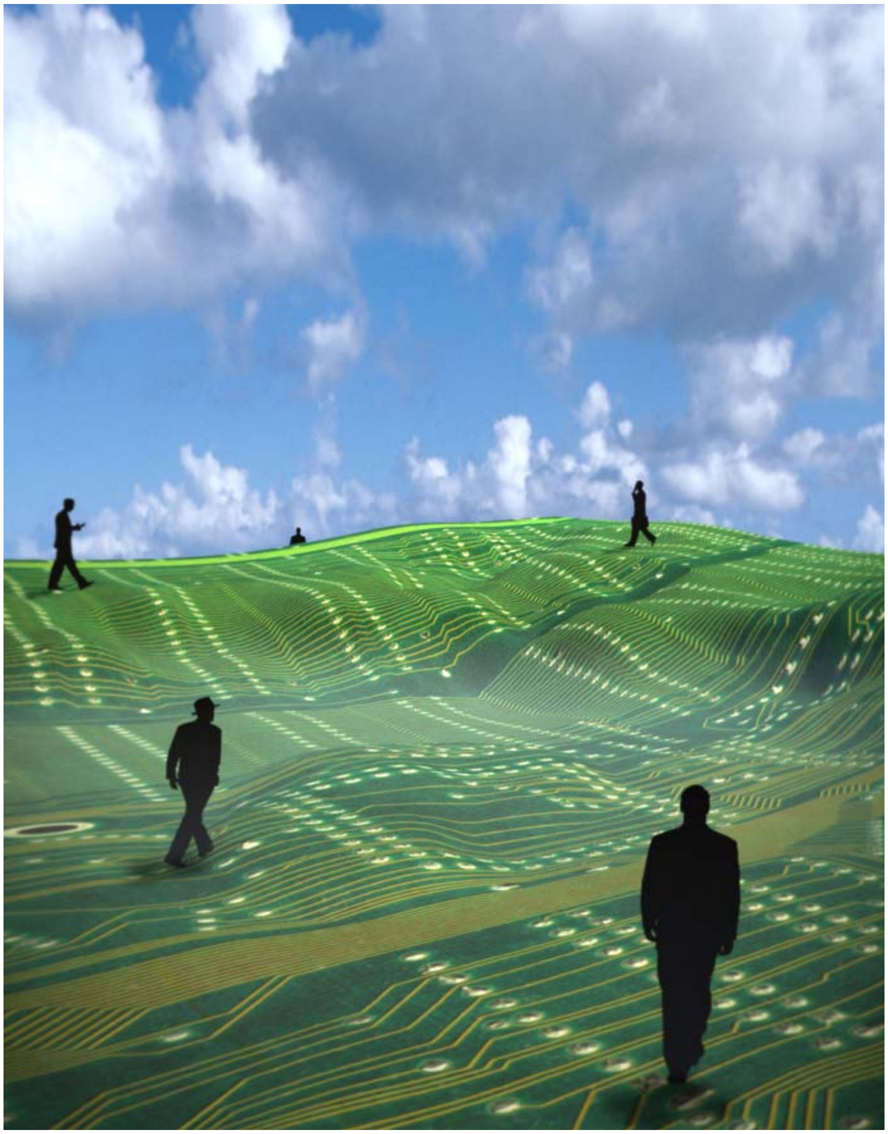

11 September 2007

This FailSafe Server Consolidation Brief (“Brief”) was prepared by the management of ZeroNines Technology Incorporated (“ZeroNines”), and is being furnished by ZeroNines, subject to the prior execution of the Confidentiality Agreement, solely for use by a limited number of third parties potentially interested in exploring business continuity solutions. ZeroNines does not make any representations as to the future performance of ZeroNines. Additionally, ZeroNines believes that the sources of the information presented herein are reliable, but there can be no assurance that such information is accurate and ZeroNines expressly disclaims any and all liability for representations or warranties, expressed or implied, contained in, or for omissions from, this Brief or any other written or oral communication transmitted or made available, except such representations and warranties as may be specifically provided in definitive contracts to be executed and delivered. Except as otherwise indicated, this Brief speaks as of the date hereof. Neither the delivery of this Brief nor any ensuing discussions conducted hereunder shall, under any circumstances, create any implication that there has been no change in the affairs of ZeroNines after the date hereof, or other specified date. This Brief is being furnished for information purposes only with the understanding that recipients will use it only to decide whether to proceed with discussions with ZeroNines management involving ZeroNines solutions. The information contained in this Brief is confidential and proprietary to ZeroNines and is being submitted solely for recipients’ confidential use with the express understanding that, without the prior express permission of ZeroNines, such persons will not release this document or discuss the information contained herein or make reproductions or use it for any purpose other than potential discussions with ZeroNines management. By accepting this Brief, the recipient reaffirms its obligations set forth in the Confidentiality Agreement entered into in connection with the receipt of the Brief and agrees: (a) to maintain in strict confidence the contents of the Brief in accordance with such Confidentiality Agreement; (b) not to copy any portion of this Brief, and (c) if the recipient of the Brief does not enter into a transaction with ZeroNines to promptly return this Brief to ZeroNines at the address below. Inquiries regarding ZeroNines should be directed as follows:

For financial matters

Mr. John C. Botdorf, Chairman
ZeroNines Technology, Inc.
Corporate Headquarters
450 East Happy Canyon Road
Castle Rock, CO 80104
+1.303.814.8121
John.Botdorf@ZeroNines.com

For all other matters

Mr. Alan Gin, President and CEO
ZeroNines Technology, Inc.
West Coast Operations
308 42nd Avenue
San Mateo, California 94403
+1.303.814.8121
Alan.Gin@ZeroNines.com

Contents

Chapter 1
The Anatomy of Disaster Recovery

Chapter 2
The Assurance of FailSafe

Introduction

The Disaster of Disaster Recovery

ZeroNines does not use a disaster recovery strategy, and does not advocate DR for our customers, for strategic and practical reasons. Our strategic reason: *recovery* is reactive, what happens *after* a disaster has *already harmed* your business. On its face, this is unsound strategy. Even if DR were strategically tenable, however, we would not rely on it because the methods available today for its implementation are riddled with failure points.

The disaster recovery architecture, which uses the synonym "failover," is based on the cutover archetype: a system's primary component fails, damaging operations; then failover to a secondary component is attempted to resume operations. The problem with the cutover archetype is that it views unplanned downtime as inevitable, acceptable, and so requires that business halt.[1]

During each cutover, either some transactions are lost or the entire system is down. This is the failure of the architecture. No amount of diligence works around it. Beyond the two principal cutovers, an additional cutover can be required. Some organizations cannot occupy a disaster recovery service provider's secondary system for the time necessary to effect primary recovery, due to oversubscribed assets of non-exclusive access contracts. In these scenarios, typically driven by resource constraints, a cutover occurs from the secondary site to a temporary site, then from the temporary site to the primary site for recovery.

An executive from EMC Corporation, the leading computer storage equipment firm, puts it this way: "failover infrastructures are failures waiting to happen."[2]

> If the boards of several publicly traded companies had any idea how much they are spending on today's disaster recovery architectures, they would realize they are paying for a fire sprinkler system that probably won't work if they have a fire.[3]

[1] We see the cutover archetype as a subtle systems design flaw that, in addition to driving unsubtle risks, also feeds the organizational learning disability known as the "fixation on events." For more information, see The Fifth Discipline, Peter Senge, Doubleday, 1990.

[2] Dorian Naveh, Director, Product Marketing, 2005.

In this view, disaster recovery enables disasters: its very design enables damage.

When market, political and regulatory expectations that drive always-on operations did not exist, DR weaknesses were not a material risk to commercial organizations or a political risk to governmental organizations.[4] Executives and IT professionals assumed that unplanned downtime was inevitable due to technology or other constraints, and with reasonable stakeholder expectations that was acceptable. Given these assumption, organizations surrendered in advance and accepted the weaknesses of the DR paradigm. But stakeholder expectations have risen and continue to rise, not only because people can be impatient, but because they pursue growth, improvement and excellence.

ZeroNines' belief in the value of business continuity exceeds our faith in disaster recovery strategy and other commercially available products and services. Our founders have seen so many organizations go down because of the limitations of widely used single-vendor DR implementations. ZeroNines has developed the patented FailSafe method and architecture to enable real business continuity.

To paraphrase Sam Nunn, former US Senator and Chairman of the Nuclear Threat Initiative: if an application outage damages our organization, what would our after-catastrophe reports say we should have changed to prevent it? So why aren't we making those changes now?

We explore these themes in this paper. We first examine the value of operational continuity, what we call business continuity, and explore rising commercial and regulatory expectations for resilience. We then survey the common exposures, technical and practical flaws of the disaster recovery strategy. We close with a description of our FailSafe solution.

[3] Conversation with ZeroNines, Benjamin Taylor, Chairman Emeritus, Disaster Recovery Institute, January 2002.

[4] Military risk from operational outages has always existed. Our focus here is rising expectations in the civilian context.

Chapter 1

The Anatomy of Disaster Recovery

Table of Contents

The Anatomy of Disaster Recovery

The business of business continuity

Continuity is valuable

How disastrous is a disaster recovery that fails? Put another way, how valuable is business continuity—and why?

Data security and business continuity are valuable because operational failures are expensive in their direct and indirect costs. A vivid example of direct cost is lost revenue. An indirect cost is a drop in the company's stock price after an operational crisis.

A study of 350 operational crises at North American and European financial institutions, in which the direct financial loss exceeded $1 million per crisis, shows shareholder loss metastasizes to 12x the direct loss over 120 working days, cutting total shareholder returns by an average of 2 percent. The average direct loss in the sample is $65 million. Less than half of the risk events in the sample are from betrayals such as embezzlement, loan fraud, deceptive sales practices, antitrust violations and noncompliance with industry regulations—leaving more than half to other categories such as natural disasters and computer system failures.[1]

Figure 1-1
Indirect vs direct losses, financial services firm crises (McKinsey)

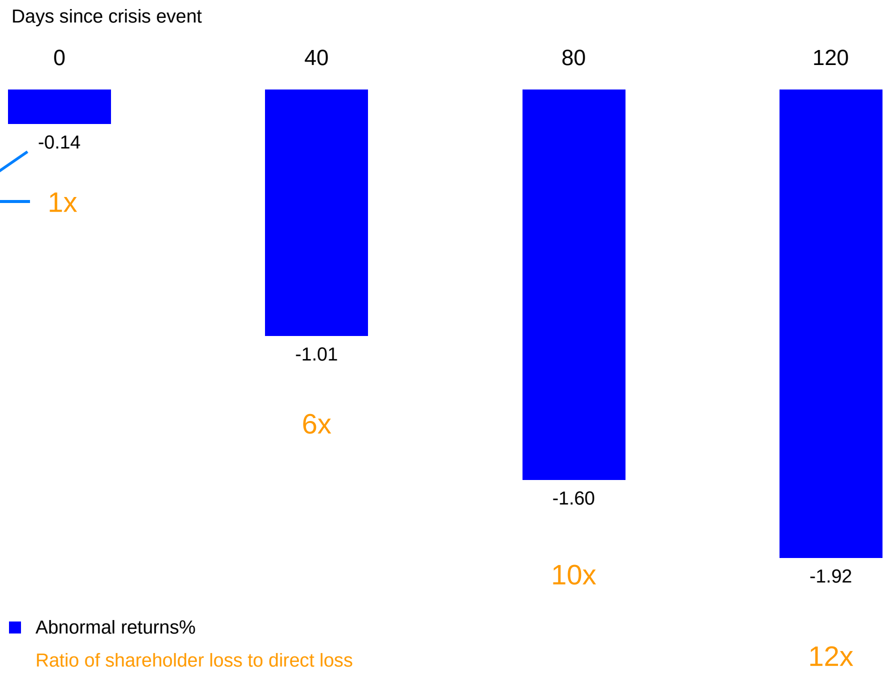

Average direct loss is equal to -0.16% of shareholder wealth, so the 0-day indirect impact of -0.14% rounds to 1x the direct impact. Indirect loss metastasizes to just under 2% of shareholder wealth over 120 working days.

[1] "Managing Operational Risk in Banking," *McKinsey Quarterly* 2005, 1.

Quantitative studies of operational failures include the following:

- Since 1982, "failover" software recovery attempts using traditional disaster recovery approaches have averaged 40 per year, primarily due to loss of electricity, hardware and fires.[2]
- Large companies forego *3.6 percent of revenue* annually due to downtime, and the leading cause of those failures is application software faults, 36 percent of the total.[3]
- Of the 350 companies in the World Trade Center before the 1993 bombing, 150 were out of business a year later because of the disruption.[4]

These are examples of *private* value of business continuity, when the wealth of one set of shareholders, or the paychecks of one set of employees, is at risk.

New expectations for resilience

Systemic risk is the value lost when the interaction of different companies or parts of the economy is disrupted. This is the conceptual space where economic damage of a disaster grows exponentially and the complexity of recovery stupefies the imagination. It is the place where companies greet regulators who are interested in uptime. We believe regulators are beginning to view firms that cannot recover quickly as imposers of economic externalities, like polluters. Appropriately or not, what has long been a private matter of competition is becoming a public matter of regulation.

As part of the Federal regulatory response to 9/11, three Federal agencies solicited financial services industry comments on draft resilience practices for the US financial system. The thrust and intent of the draft was retained in the Interagency Paper issued in April 2003. The Paper now has Final Rule status.[5]

In interpreting the Interagency Paper, ZeroNines concurs with the Evaluator Group, a consultancy:

> Every CIO and Chief Legal Officer needs to read these documents. While they apply only to their industries in the short run…, they…. will define security standards for much of the IT industry by the end of this decade.[6]

[2] CPR Research, 2005.

[3] The Costs of Enterprise Downtime", Infonectics Research, 2/11/2004.

[4] Gartner/RagingWire report cited in "Without the wires," Fabio Campagna, Disaster Recovery Journal, Winter 2002.

[5] Unless otherwise noted, what follows is based on ZeroNines analysis and "Interagency Paper on Sound Practices to Strengthen the Resilience of the U.S. Financial System." Board of Governors of the Federal Reserve System, Office of the Comptroller of the Currency, Securities and Exchange Commission. April 2003.

[6] "All aboard the new federal security rules super train," Jack Scott, TechTarget.com, 6/11/2003.

Regulators expect essential firms to recover and resume with zero data loss within two hours of a disaster (the *two-hour rule*) using a distant secondary site (the *dispersal rule*). They state that "back-up sites should not rely on the same infrastructure components (e.g. transportation, telecommunications, water supply and electrical power) used by the primary site." Regulators clearly want a failover site hundreds of miles away from the primary site so the secondary site is not disrupted by the same weapon of mass destruction, earthquake or hurricane that disrupts or destroys the primary site. When the Interagency draft was circulated for comment in August 2002, all three of these trauma scenarios were considered plausible threats.

Note ZeroNines' site diversity concept enables our customers to fulfill the requirements of the dispersal rule. The always-on nature of our MultiSynch technology enables customers to fulfill the requirements of the two-hour rule—or, for that matter, two-minute or two-second rules, if they are ever established.

Information security and business continuity standards are changing and the trend is clear. Customers are beginning to judge by the new standard of *business continuity*, virtually 100 percent accessibility. The more important your firm is to the economy—the more successful it is or the more central its role in commerce—then the more likely you face the security and continuity requirements of regulated industries. We are not saying that this degree of government involvement is appropriate or not. We state that it is expanding.

Figure 1-2 depicts IDC research indicating a 53% reduction in commercial expectations of planned + unplanned downtime through CYE2007.

Figure 1-2
Commercial operational continuity expectations (IDC)[7]

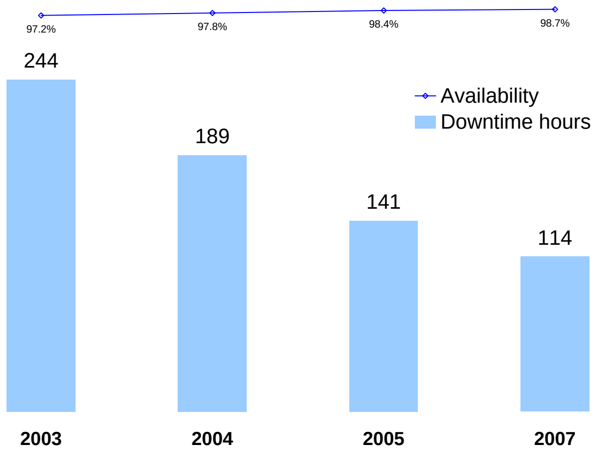

[7] The study omits 2006 data. “Optimizing Business Performance Requires Optimizing Information Availability Investments.” IDC, 2006.

The threats

Given the value of business continuity—of disaster avoidance—what threats must be recognized? We summarize the breadth of the threat universe in Table 1-1.

Table 1-1
Threats summary (ZeroNines)

Threat type	Examples
Component	• Hardware and software failures • Backup system failures • Communications component failures
Data center	• Loss of data center resources, such as electrical, networking • Fire detection or retardent systems • Man-made (accidental, cracking)
Regional	• Acts of nature such as earthquakes, storms, floods and fires • Loss of utility resources, such as electrical grid, communications, water or transportation for resources such as recovery media
Global	• Distributed denial of service attacks • Viruses, worms, etc.

A quick scan of these threats invokes Murphy's Law: if something can go wrong, it will.

Every application service protected by the ZeroNines FailSafe architecture and technology has remained available to its application clients' network 100% since implementation. There has never been a case of a FailSafe application client failing to reach its FailSafe application service across an operational network. That said, we have seen many "threats" become "facts." Mentioning them conveys the bitter flavor that challenges conventional disaster recovery architectures.

- On August 12, 2004, Hurricane Charley caused electrical grid fluctuations that drained the Orlando local exchange carrier battery backup systems, isolating our node. Our own battery system prevailed and still had a 75% charge when commercial power was reliably restored, but the site could not communicate for 16 hours because of LEC downtime.

- During the late-December 2004 Santy worm attack on phpBB code, AOL email to two of our Board members was disrupted as AOL battled the worm. Email service by our system was not disrupted.
- In December 2004, a 3-day data center move disrupted service from our Florida node. As before, email clients received uninterrupted service.

So if those are the threats, why can’t disaster recovery architectures handle them?

The cutting edge of cutover

ZeroNines believes that existing disaster recovery designs are vulnerable. These exposures aren't the fault of IT departments, but flaws propagated by proprietary vendor designs that have been present for years.

To restate a key insight: the disaster recovery architecture, which uses the synonym "failover," is based on the cutover archetype. The cutover archetype is flawed because it forces the customer to accept outages that disrupt business and might abruptly terminate careers.

The design flaw of failover is that data protection is driven by the last image backup before the threat materializes. Primary system recovery requires system downtime, data migration and replication. At least two, sometimes three, cutovers are required (Figure 1-3):

- from the primary system to the secondary system (the failure from the threat)
- from the secondary system back to the primary system (the recovery).

Figure 1-3
At least two cutovers per disaster

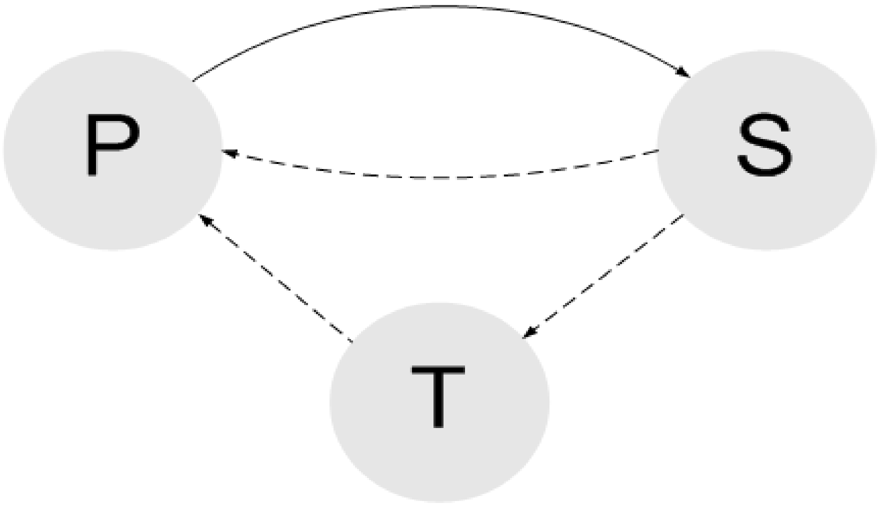

During each cutover, either some transactions are lost or the entire system is down. This is the failure of the architecture. No amount of diligence works around it.

Beyond the two principal cutovers, an additional cutover can be required. Some organizations cannot occupy a disaster recovery service provider's secondary system for the time necessary to effect primary recovery, due to oversubscribed assets of non-exclusive access contracts. In these scenarios, typically driven by resource constraints, a cutover occurs from the secondary site to a temporary site, then from the temporary site to the primary site for recovery.

Figure 1-4 depicts the central technical flaw in action, showing system events and the end-user experience. Downtime persists from when the threat becomes an event until the user session resumes on the secondary system. Downtime returns during recovery from the secondary system back to the primary. If a temporary system other than the secondary and primary is utilized, more downtime is encountered.

Figure 1-4
Why downtime is inevitable with disaster recovery architecture

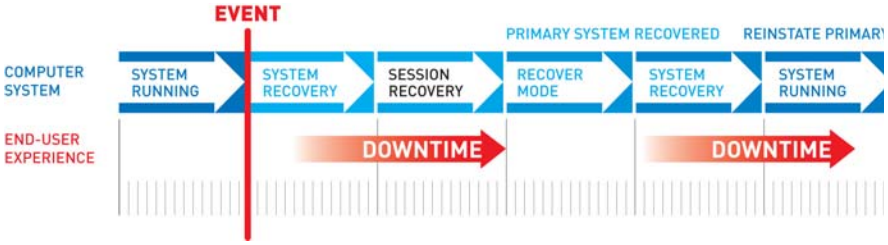

What’s recoverable from recovery?

Before 9/11—indeed, before Katrina or the always-on Web operations now expected by customers, constituents and regulators—the following disaster recovery designs were usually deemed adequate:

- Tape-based recovery on page 2-1
- Remote vaulting recovery on page 2-5
- Failover and clustering recovery on page 2-7.

We now explore each of these designs in terms of their methods, architectures and exposures. Due to their implementation of the cutover archetype, as well as other drawbacks, we believe none of these approaches provides sufficient and affordable business continuity assurance for our customers.

Tape-based recovery

Most companies use a tape-based disaster recovery strategy that was developed in the 1970s, before IT moved from the back office to become central in business. Tape-based disaster recovery uses a failover approach as depicted in Figure 1-5 and described as follows.

Figure 1-5
Tape-based recovery architecture

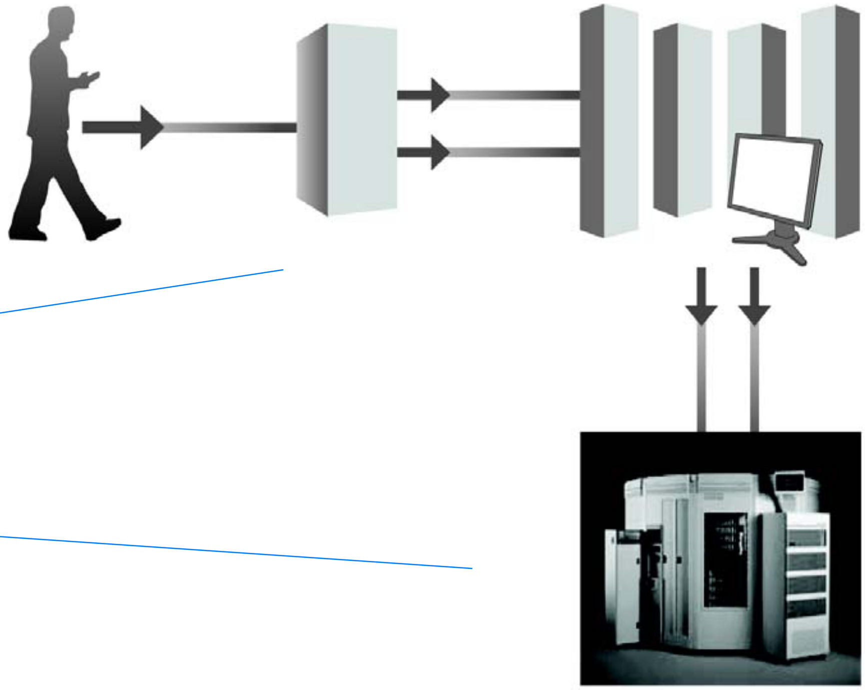

SCSI backup ~ 50 GB/hr. (also using virtual tape technology)
Fibre channel theoretical
(360 GB/hr.— serverless backup)

SCSI tape drive ~ 50 GB/hr.
80 GB tape cartridges
90 tape changes/hr. max.

1 Periodically, backup copies of essential business data are produced at the primary site and transported to an offsite storage facility. For 90% of Global 1000 firms that use failover services,[8] each backup copy utilizes myriad magnetic tape cartridges, each about the size of a paperback book.

2 The primary site fails.

[8] GartnerGroup.

3 Seeking access to a contracted secondary site run by a disaster recovery service provider (DRSP), such as IBM, Sungard or HP, the CIO meets the contractual access requirement by declaring a disaster. If the CIO is not the first to declare a disaster in a shared-resource contract, access to the secondary site is not assured.[9]

4 The most recent backup copy from Step 1 is ordered transported to the secondary site. All tapes might be included in the shipment, but perhaps one is omitted accidentally. Subsequent transit time depends on interaction between the means of transit and weather conditions.

5 Tapes are used to "restore" the data and application software to the computers at the secondary site. If a single tape is damaged, used out of sequence, or is missing, the restore operation fails and must be restarted—assuming all tapes are present.

6 Operations resume at the secondary site.

This simple example shows only one cutover, from the primary to secondary site. At least a second cutover is required, from the secondary back to the primary. As we noted on page 1-7, a third cutover might be required as well. The DRSP may eject a shared-resource customer out of an oversubscribed recovery site to make room for another customer.

A representative timeflow of a tape-based recovery attempt is as follows.

Table 1-2
Tape recovery attempt time flow

00:00	Last back-up performed
Processing continues	
09:45	Disaster strikes. Shortly thereafter, disaster is declared. Tapes are ordrered to recovery site.
10:50	Recovery starts
10:55	Backup systems brought on-line
??:??	Tape recovery starts
??:??	Users access recovered system

[9] Contracts for dedicated resources average 7x the cost of the shared-resource alternative. "Things to consider before choosing a primary site recovery approach or telecommunications vendor," Randolph Fisher, CBCP. Disaster-Resource.com.

In the first cutover, there is a tangible gap between the time the threat materializes and the tape recovery begins. Latency thereafter and in subsequent cutovers depends to some degree on tape-based data transfer rates. Table 1-3 depicts theoretical limits of widespread tape technologies. With the storage requirements that our customers describe, tape-based recovery doesn't even come close to meeting stated recovery time objectives.

Table 1-3
Tape transfer theoretical limits

Data amount (TB)	1 SCSI channel	4 SCSI channels	1 Fibre channel 1GB	1 Fibre channel 2GB
0.1	2 hrs	30 min.	< 18 min.	< 9 min.
1	20 hrs	5 hrs	< 3 hrs	< 1.5 hrs
10	> 8 days	50 hrs	< 28 hrs	< 14 hrs
36	30 days	7.5 days	> 4 days	> 2 days

Exposures and drawbacks

What are the key exposures and drawbacks of tape-based solutions?

- Any new transaction between the last tape backup-up and the threat event is potentially lost. This appears to be the central flaw.
- Tape inventory management must be flawless. A missing or out-of-sequence tape not discovered in advance ruins the first recovery attempt. A second delivery request for a missing tape delays the first recovery attempt. Tape damage jeopardizes the entire recovery.
- Tape loading is constrained by the quantity of simultaneously available tape drives.
- Travel is risky in natural disasters. Conditions at the storage site, recovery site and in between must be considered. A jet cannot deliver tapes if it cannot land. A truck cannot deliver tapes if the road is coated with ice or diced by a hurricane or earthquake.
- Under Service Level Agreement queuing, only the first customer of the recovery site to declare a disaster is contractually assured access to recovery resources.

Examples of delivery problems On August 29, 2005, the surface course of five miles of Interstate 10, the principal road access to New Orleans across the eastern edge of Lake Pontchartrain, was chopped to pieces by Hurricane Katrina and did not reopen until October 14. Both other routes across the lake, US 11 and US 90, were restricted to emergency personnel for three days. The freeway system of Los Angeles was heavily damaged by the Northridge quake.

Remote vaulting recovery

Attempting to address the weaknesses of tape-based recovery, vendors now support remote vaulting with split-mirror imaging (Figure 1-6). Vaulting has the advantage of reducing data transportation risk to practically zero by utilizing highly reliable telecommunications networks.

Figure 1-6
Remove vaulting with split-mirror imaging

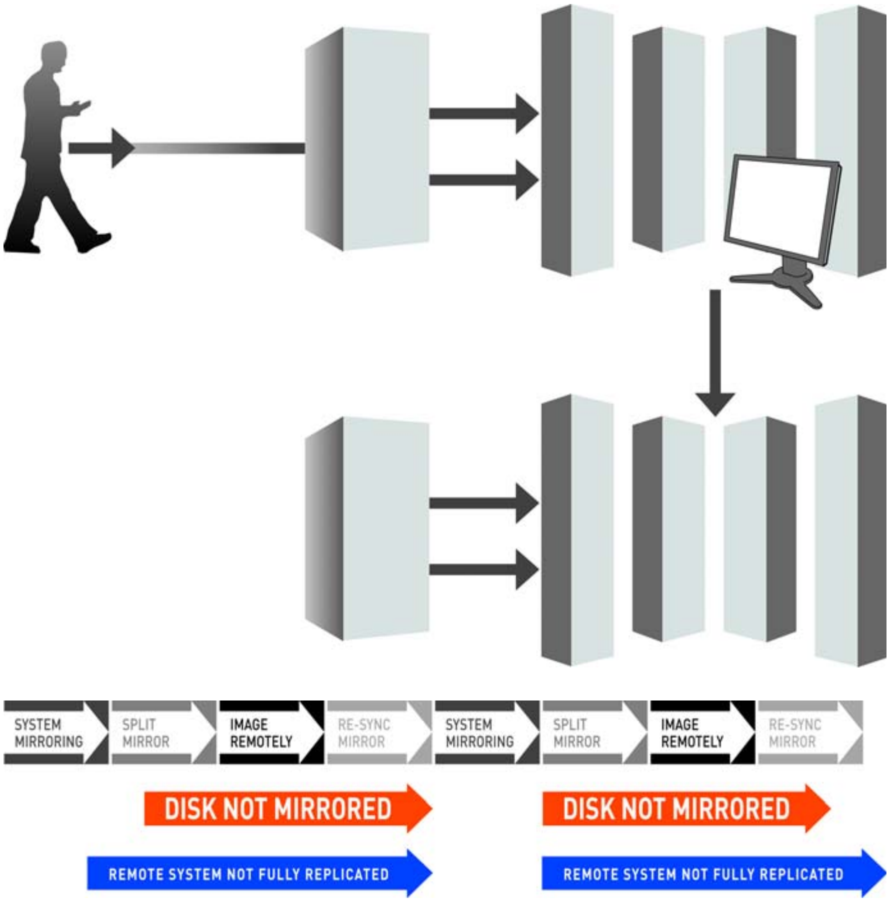

Data written to DASD on the primary system is mirrored locally.

Mirrored DASD is then split (broken) and then the modifications made to disk since the last copy are sent to the remote site utilizing remote copy function across leased lines (IBM Remote Copy, EMC SRDF).

The locally mirrored DASD is then re-established and re-synchronized.

A representative timeflow of a tape-based recovery attempt is as follows.

Table 1-4
Vaulting recovery attempt timeflow

00:00:00	Last replication
00:15:00	Next replication. Replication process continues.
09:00:00	Disaster strikes.

09:05:00	Recovery starts
09:15:00	Backup systems brought on-line
09:30:00	Essential applications brought on-line
09:40:00	Users access recovered system

Exposures and drawbacks

What are the key exposures and drawbacks of remote vaulting solutions?

- Any new transactions between the last mirror (replication) and the threat event realization is potentially lost. This appears to be the central flaw.
- Leased-line expenses are incurred, and supported distances are not adequate to ensure continuous availability.
- Due to leased-line expenses and related capacity constraints, the common practice is to protect only the "most essential of the most essential" data.
- Technologies required are proprietary to big hardware vendors and service providers, so customer negotiating leverage is difficult to achieve or maintain. Matched hardware is required, so capacity must be added in larger-than-desired chunks.
- Mirror splitting and re-establishment must be flawless or database consistency must be explicity controlled, a technical and managerial headache. Even commit loggers cannot protect in-flight transactions.

Failover and clustering recovery

Server-based failover and clustering solutions are the least bad of traditional disaster recovery architectures, but they have their own problems.

Figure 1-7
Server-based failover and clustering

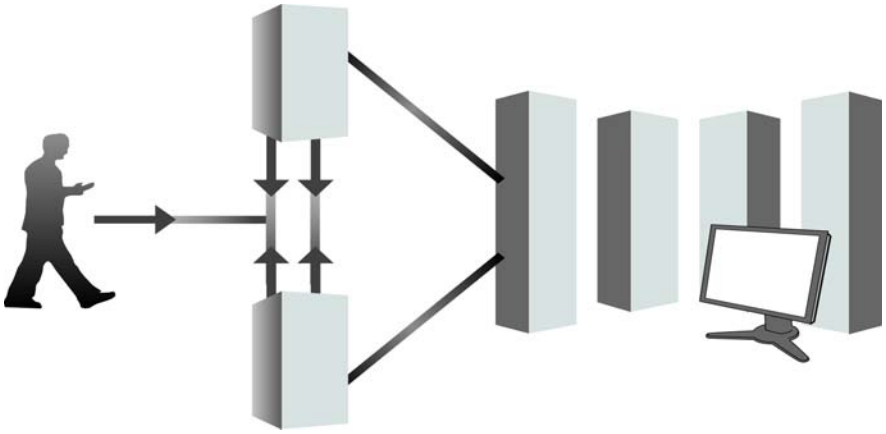

The key failover method is as follows: The secondary node monitors primary through a “heartbeat” connection. When the primary fails, secondary takes over processing. Application sessions are thus maintained. Usually the primary and secondary share disk space, and the distance between servers is usually less than 1 km.

Exposures and drawbacks

What are the key exposures and drawbacks of server-based failover and solutions?

- Latency is possible between the time a threat event is realized and the heartbeat detection triggers secondary processing. Transactions can be lost. This appears to be the central flaw.
- Supported distances are inadequate to support required site dispersal.
- Technologies required are proprietary to big hardware vendors and service providers, so customer negotiating leverage is difficult to achieve or maintain. Even more than with vaulted solutions, clustered systems tend to be among the most

expensive in the commercial computing market. Matched hardware is required, so capacity must be added in larger-than-desired chunks.

- Shared storage must be replicated carefully or it becomes a single point of failure; even then, block-rewrite issues must be addressed, increasing technical complexity (and therefore risk).
- Application compatibility with cluster operating systems has typically been more difficult to assure. Third-party software availability might be constrained, further diminishing customer negotiating leverage.

Conclusion

We conclude that disaster recovery is not strategically tenable. Extensively used disaster recovery architectures have fundamental design exposures that cannot be worked around. IT organizations cannot circumvent the weaknesses with clever and diligent implementation. Disaster recovery designs are indadequate to support continuous application availability.

The two-hour rule and the dispersal rule cannot be satisified jointly by any alternate commercial disaster recovery technology from any other leading service provider or vendor today.

Figure 1-8
Problem summary

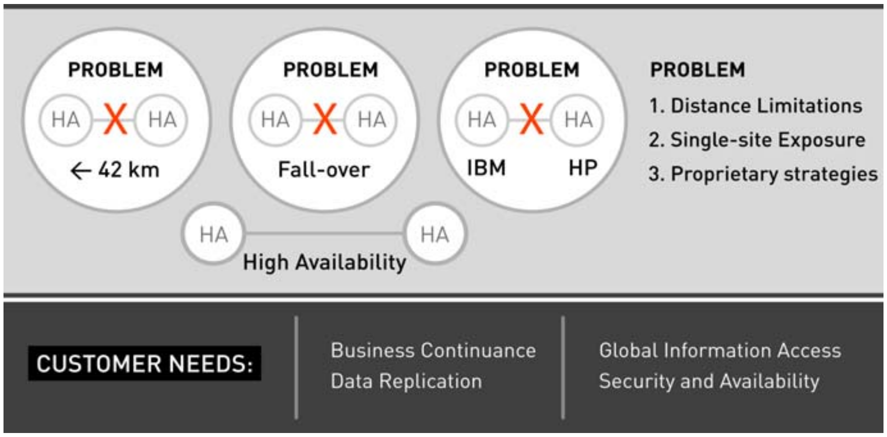

Sungard, a disaster recovery service market share leader that was taken private in August 2005, issued a press release in response to the draft Interagency guidance along these same lines. Sungard wrote:

> [A]ccelerated intra-day recovery/resumption with zero data loss, and a separation of 200-miles [sic] between primary and secondary sites, are technologically incompatible at this time….[C]yber-attacks, which represent a clear and present danger … are not sufficiently addressed by the Draft Interagency White Paper.[10]

[10] "SunGard Offers Comments on Draft Interagency White Paper on Sound Practices to Strengthen the Resilience of the U.S. Financial System." Press release, 12/18/2002. http://www.sungard.com.

Chapter 2

The Assurance of FailSafe

Introduction

The Assurance of FailSafe

Given the exposures and drawbacks of the disaster recovery architecture and technology, what requirements must a business continuity solutions address? We suggest the following.

- Mitigate regional disasters
- Leverage current assets, not requiring speed- or capacity-matched hardware
- Hardware-agnostic
- Operating-system-agnostic
- Network-agnostic
- Do not require prolonged application customization
- No loss of in-flight transactions
- Simple, elegant and cost-effective.

Our FailSafe solution meets these requirements, as summarized in Figure 2-1.

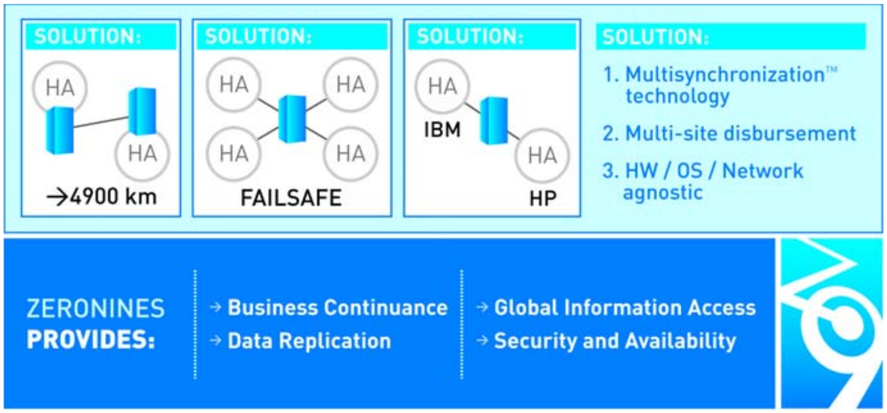
SOLUTION:
HA
HA
→4900 km
SOLUTION:
HA
HA
HA
HA
FAILSAFE
SOLUTION:
HA
IBM
HA
HP
SOLUTION:
1. Multisynchronization™ technology
2. Multi-site disbursement
3. HW / OS / Network agnostic
ZERONINES PROVIDES:
→ Business Continuance
→ Data Replication
→ Global Information Access
→ Security and Availability

Figure 2-1
FailSafe solution summary

ZeroNines' FailSafe architecture disaster-proofs an application without a wholesale application rewrite. Instead, a protected application communicates to the infrastructure through a FailSafe protocol interface (adapter).

The "application" in this sense is the user of the FailSafe architecture. Examples include but are not limited to:

- storage configurations
- databases
- transactions monitors
- email systems
- other business application software.

Application availability on a ZeroNines FailSafe configuration exceeds commercial alternatives at the same or lower cost for the same or greater uptime. Our architecture overcomes the limitations of disaster recovery architecture with novel topology and protocols. The effect is similar to assembling ordinary struts into a geodesic dome. Our architecture makes the system more reliable than its component parts, and the larger the system, the more flexible and robust it (and IT) become.

Figure 2-2
FailSafe solution topology

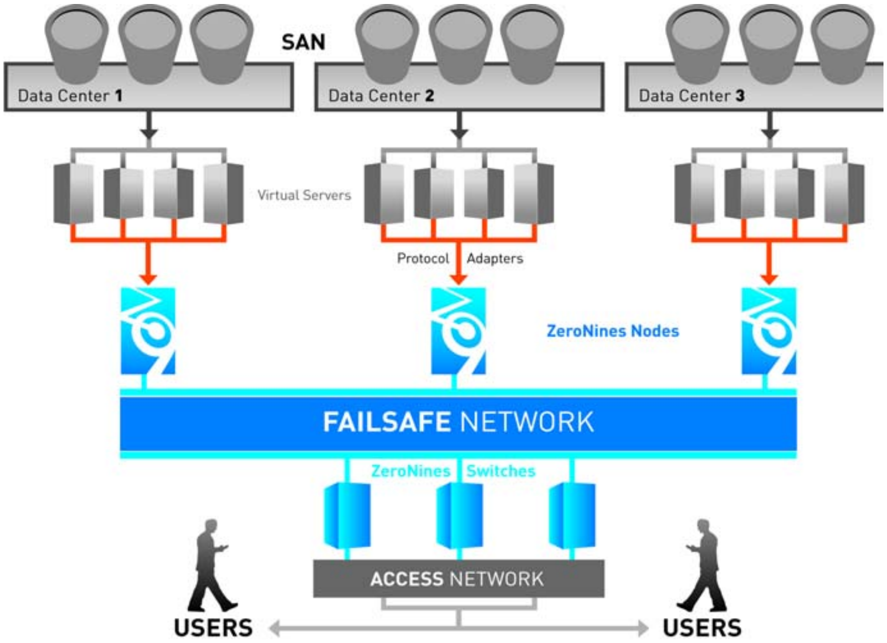

Design principles of a FailSafe solution

The design principles of a FailSafe solution are:

- A one-to-many (1:m) session type is supported
- Server hierarchy is eliminated
- Server sites are diverse
- Heterogeneous product sets are accommodated
- Load balancing is a side effect.

A one-to-many (1:m) session type is supported

A FailSafe configuration maintains application sessions that are one-to-many (1:m) in nature. Each session from a client (service requestor) is maintained with multiple application servers (service responders). Duplicate replies from servers are eliminated during return to the client, ensuring integrity of the application image.

The application need not be session-oriented from the application's point of view. ZeroNines FailSafe supports sessionless and session-oriented applications.

Server hierarchy is eliminated

Each application server image in a FailSafe configuration is always logically primary. In contrast with human relationships, server hierarchy does not exist in a ZeroNines FailSafe configuration. There are no secondary servers—not even the concept of "first among equals." Server primacy is perfectly shared without loss of effectiveness. At least two servers process every client request. Because there are no secondary servers, logical failover at the application layer does not occur, nor does it need to occur. Processing by one site might cease within the FailSafe configuration for typical reasons such as scheduled maintenance or physical trauma, but the other sites in that configuration continue processing in a zero-loss manner that is transparent to the application.

Server sites are diverse

ZeroNines uses the "site diversity" concept to indicate a number of server sites that share no physical exposures, such as infrastructure failure, natural disaster, fire or explosion. When server sites are diverse, dispersed by hundreds or thousands of miles and not dependent on the same infrastructure, FailSafe application availability is feasible.

Example Sites n New York and Singapore are diverse. They share neither natural disasters nor essential infrastructure such as electricity, water, or local exchange carriers. In this example, site diversity is two: two sites with no shared exposure.

Application availability is augmented as diverse sites are added to a configuration: five nines, seven nines or, with larger numbers of servers, effectively zero nines—100% application uptime to client requests, even with unscheduled server maintenance.

The combination of shared server primacy and site diversity obviates application-wide recovery because application-wide failure does not occur.

Heterogeneous product sets are accommodated

Heterogeneity as a design principle produces more robust systems by minimizing system-wide effects of:

- attacks that are specific to a particular operating system
- vulnerabilities to model-specific defects of vendor hardware or software.

Example Every IT professional knows of situations in which Linux servers kept running when NT servers were under attack. Any operating system can be attacked. That said, we have never heard of a successful *all-OS* attack in a commercial setting.

ZeroNines FailSafe capability can be achieved with or without heterogeneous product sets. You can mix and match old and new hardware and operating systems, even from different vendors, without compromising FailSafe integrity. ZeroNines' protocols prevent race conditions and operate asynchronously across thousands of miles.

Removing matched-speed and matched-capacity constraints eases the burden of prototype projects and enables maintenance and upgrade of production servers and networks. You don't have to do everything at once to develop a prototype, deploy, or to maintain production.

The benefits of heterogeneity can be considered in the context of increased complexity. Some IT organizations prefer to standardize on one server operating system to achieve economies of scope and scale in consolidated infrastructure. Other organizations have long ceased attempting such an approach in favor of accommodating top-down decisions driven by user requirments. Being application- and platform-agnostic, ZeroNines' architecture does not constrain the choice of server operating system, hardware or network protocols, enabling heterogeneity as a design strategy for those who choose it without excluding those who do not.

Load balancing is a side effect

The combination of shared server primacy and heterogeneity produces, as a side effect, a survival-of-the-fittest load balancing to support your application layer. FailSafe consolidated servers effectively compete to return results to requesting clients. A server that is closer to the requesting client or that temporarily has less workload might return a result more quickly than a faster processor that is more geographically distant or temporarily under heavier workload.

Designers remain free to match speeds and capacities of servers or networks for proprietary application-layer load balancing algorithms without disrupting FailSafe capability.

An infrastructure before and after

To understand how a ZeroNines FailSafe configuration differs in a general sense from typical application access, consider the following exhibits.

Figure 2-3 depicts a typical application access topology, before FailSafe. An access network links users' application clients to a datacenter's internal network via firewall, router and secure gateway. The server complex responds to application requests. In this example, database service was separately defined for ease of reconfiguration or performance.

Figure 2-3
Typical consolidated application access (not failsafe)

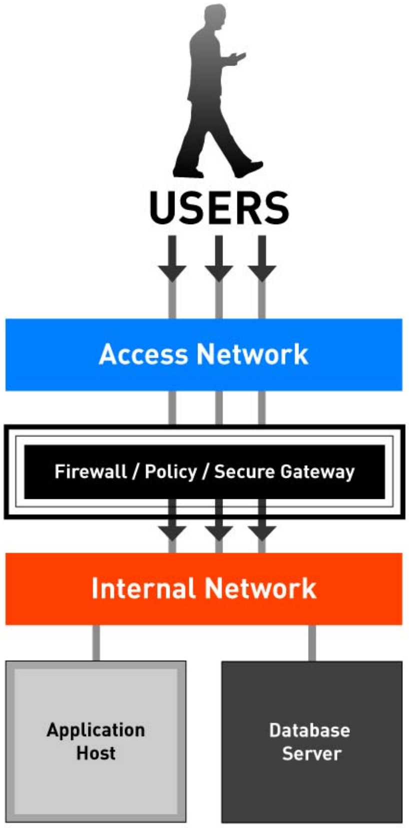

In a ZeroNines FailSafe topology as shown in Figure 2-4, two (or more) ZeroNines FailSafe switches are present between the application user network and the application servers' network. Each FailSafe switch may have one or more state-accurate shadowing switches that continue service to the application clients in case a switch discontinues service for any reason, such as scheduled maintenance. FailSafe Switches may be clustered for load balancing as desired.

Figure 2-4
FailSafe consolidated application access

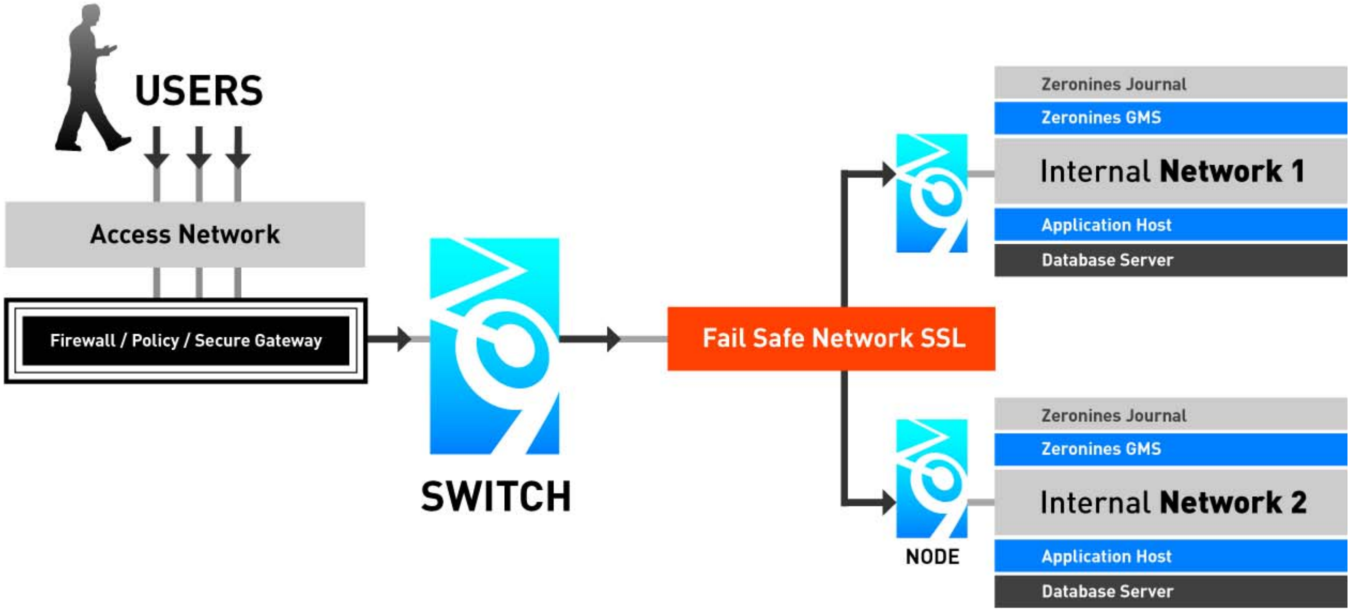

The mere fact that a FailSafe configuration contains fewer single points of failure from a hardware perspective does not fully explain why continuous application availability is assured. Simply buying more servers and configuring them for traditional DR failover is insufficient to enable 100% uptime. Failover is insufficient for continous availability. A FailSafe architecture requires the FailSafe design principles to be implemented.

In a ZeroNines FailSafe configuration, each application server is associated with a ZeroNines FailSafe *node*, a listener function. When a client requests application service, at least two FailSafe switches pass the request to at least two FailSafe node listeners, each of which completely and independently processes the request using the respective servers associated with those listeners. The results generated by the servers are returned by the respective listeners to the switches, which cooperatively return *one* copy of the result to the requesting client. Thus a 1:m session is implemented. Duplication of data is prevented, and integrity of results is ensured, by the ZeroNines protcols and formats that are completely transparent to the application. Listener functions may be implemented as hardware integrated with the consolidated server or one or more software modules running on the associated server.

A ZeroNines configuration can utilize gateway, unicast or multicast protocols, depending upon your requirements. This network protocol flexibility is captured in our Transaction MultiSynch marque.

Relating nodes to nines

ZeroNines has developed configuration guidelines for estimating the number of servers and other elements necessary to achieve desired application availability. We have tested these guidelines in our own business with our own mission-critical application.

Your FailSafe configuration must reflect the imperatives of your organization's Business Impact Analysis, business plan and regulatory requirements. ZeroNines believes that clients appreciate sizing approximations as a starting point for proof-of-concept and prototyping projects. Consultative services are available for the sizing of a FailSafe production configuration.

Table 2-1
Minimal site diversity for desired availability during prototype tests

Availability in prototype (%)	Sites diversity required
99.999	2
99.99999	3
100	> 3

Scheduled maintenance ignored. Minima shown are adequate for prototyping projects. ZeroNines offers services for designing production configurations.

Augmenting the minima shown by adding incrementally diverse sites supports greater availability, such as during routine maintenance, upgrades, or additional trauma that causes simultaneous service interruptions at two or more sites.

Case study: MyFailSafe.com

Design

ZeroNines Technology, Inc., invented MultiSynch technology and has been using it for years in our own business for our own operational continuity. We rely on it.

For us, email is a mission-critical business application, so we commenced a MultiSynch implementation with the MyFailSafe.com email service (Figure 2-5).

- We standardized on one operating system for all three server nodes, but CPU, RAM and disk are neither speed- nor capacity-matched.
- Each server node is scheduled for 15 minutes of downtime per month for log resets, staggered to ensure that no two nodes are ever scheduled for simultaneous maintenance.
- Telecommunication links are described in Table 2-2 on page 2-9.

Figure 2-5
MyFailSafe.com topology, on continuously since 2Q2004

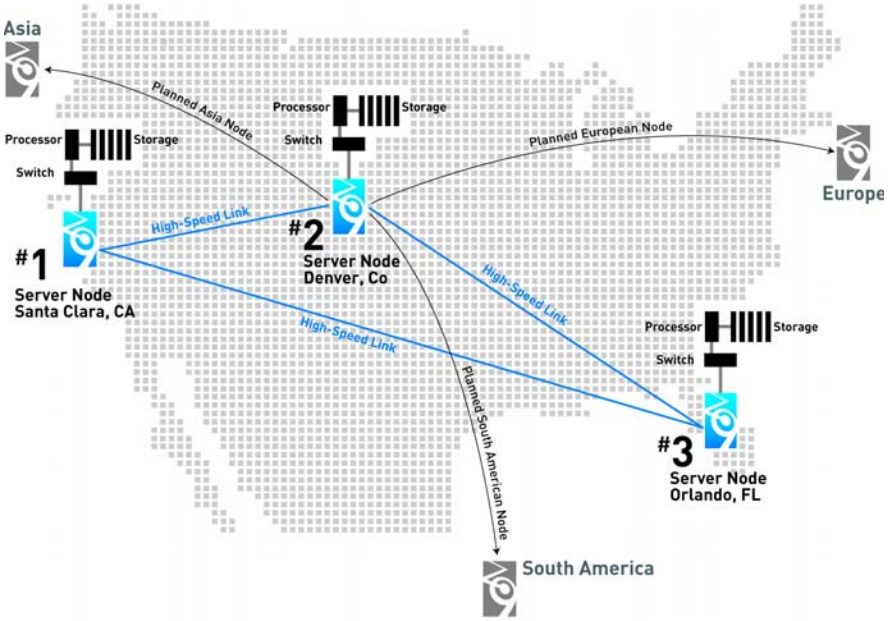

Table 2-2
MyFailSafe.com telecommunication links

City	Carrier	Link characteristics
Santa Clara, California	MCI	1MB, burstable
Denver, Colorado	Level 3	1MB, burstable
Orlando, Florida	Time Warner Telecom	1MB–10MB

Results

Since activation on July 15, 2004, MyFailSafe.com has furnished continuous service to email clients. There has never been an interruption of service to email clients for any cause: scheduled or unscheduled maintenance, server ugrades, virus attack, distributed denial of service attack, or natural disasters. Never, for any cause.