ZERODOWN
SOFTWARE™

Cloud Computing Observations and the Value of CloudNines™

ZERODOWN SOFTWARE: WHITEPAPER

www.zerodownsoftware.com

Inquiries regarding ZERODOWN Software should be directed as follows:

→ For financial matters:

Mr. Sean Myers
Co-Founder & Director

ZeroDown™ Software
5445 DTC Parkway
Penthouse Four
Greenwood Village, CO 80111
1.720.244.2120
Sean@ZeroNines.com

→ For all other matters:

Mr. Alan Gin
President and CEO

ZeroDown™ Software
5445 DTC Parkway
Penthouse Four
Greenwood Village, CO 80111
1.303.814.8121
Alan.Gin@ZeroNines.com

Contents

Opportunities

CHAPTER 1
Opportunities

SECTION 1
Opportunities

Introduction: **Opportunities**

To orient executives to the cloud computing concept and to raise awareness of its potential and peril, we recommend that technical professionals describe the concept of virtualization and then proceed to the cloud. Cloud is a metaphor for the Internet; some firms run their own private clouds.

Current cloud providers integrate the commodity technology and products of the IT product vendors into computing and storage utility services. This is happening now due to the integration of inexpensive computing products with digital networking and ingenious virtualization features, at an adequate scale and low price. This enables a breadth and depth of resource elasticity that is at the heart of leading cloud business models.

The utility paradigm, already well established for energy and telecommunications services, is growing around IT after an incubation period that spans decades. Businesses, their workers and regulators face today with respect to IT most of the same issues that their forerunners faced with respect to electricity a century ago. Standardization, scale economies and labor specialization are potent drivers. Value chains are developing to enable all of this and consultancies to facilitate it. Some of the brands were forged a century ago while electricity was morphing into utilities; others are new. Leading cloud providers have deployed very different services, so today's cloud customers have choices. Regulation of these new public IT utility services has not substantially started.

Today's cloud computing implementations convey the promise of utility computing, but also the peril of legacy architectures and technologies that were never designed for utility-grade service availability. Leading cloud providers have committed glaring service outages in recent years, and continuing problems are making headlines through the time of this writing.

The security, management and availability concerns of enterprises must be addressed for the IT utility model and cloud providers to flourish. We believe the security and management challenges are already on the path to resolution. For availability we offer our Always Available architecture and technology to enterprises and cloud providers, and to cloud-users the CloudNines™ variation.

But problems and solutions are for subsequent chapters. In this chapter, we look at the promise of cloud computing.

"Business continuity may be the killer app for cloud computing. We need reference architectures and models so clouds support business uninterrupted by future disasters. ZERODOWN Software furnishes such an architecture and technology, to cloud providers and cloud users, in Always Available™ and CloudNines™."

Reuven Cohen, Founder, Enomaly

What is "cloud computing"?

To orient executives to the cloud computing concept and to raise awareness of its potential and peril, we recommend that technical professionals describe the concept of virtualization and then proceed to the cloud. Although ZERODOWN Software is concerned with the way that virtualization and cloud computing are implemented today, at a conceptual level they are essential.

In describing virtualization to executives, we frequently use a telephone messaging example.

What is virtualization?

Telephone messaging

Telephone answering service began with human beings: family members or co-workers. A single-purpose box that recorded voices in analog form on reels of magnetic tape became available as the supply of labor for answering service decreased while demand continued to grow. As computer chips advanced, the machines became digital boxes in which messages were stored in binary form in the device's memory. Once all the functions had become fully digitized, renderable in software code, they no longer had to be built as task-specific machines or to reside on customer-specific premises. The physical machine turned into a "virtual machine"— into pure software running in the phone company's network. Where you once had to pay a human being or buy an answering machine, you now can subscribe to an answering service. That's the essence of virtualization.[1] Telephone voice messaging can be viewed as hardware-as-a-service (HaaS). Once the executive understands a virtualization example, it is relatively easy to introduce the software-as-a-service (SaaS) concept and examples such as salesforce.com and SuccessFactors. We furnish more information about these two firms and the XaaS value chain model under "Who is doing this?" on page 22.

1 Adapted from Carr, Nicholas. The Big Switch: Rewiring the World from Edison to Google. New York: W.W. Norton. 2009. 75. Carr's coverage is especially useful for nontechnical executives who need to think through the stakeholder management implications of cloud computing.

What is cloud computing?

Once virtualization is understood, it is reasonable to introduce cloud computing.

In defining cloud computing we are as content as anyone with Wikipedia's approach. Wikipedia defines cloud computing as a business information management style of computing in which typically real-time scalable resources are provided "as a service" over the Internet to users who need not have knowledge of, expertise in, or control over the technology infrastructure ("in the cloud") that supports them.

Cloud is a metaphor

"Cloud" is simply a metaphor for the Internet. It is a reference to diagrams that depict as clouds the wide-area networks provided by telecommunications companies. The phrase "cloud computing" first appeared in mainstream U.S. media in the Houston Chronicle of June 5, 2002. The New York Times had published "cloud of computers" a bit earlier, on April 9, 2001, when referring to Microsoft's positioning of .net technology from 1998.[2] References to "private" clouds mean cloud concepts are implemented by an organization intramurally, not for outside customers. The basics of cloud computing are the same, public and private.

Elasticity is a virtue of virtualization

"Elastic" is frequently used as a synonym for "scalable" in media descriptions of cloud computing. Elasticity is a key distinguishing characteristic of the cloud model. Resource availability expands and contracts easily under changing operational conditions. Virtualization and the massive scale of utilities are the key enablers of elasticity. In the traditional hosted IT model, the IT resources available to a particular customer are fixed, both exclusive and essentially constrained.

When supported with an analogous pricing model, elasticity is not only a distinguishing characteristic. It is probably the most important characteristic.

> [C]ompanies with large batch-oriented tasks can get results as quickly as their programs can scale, since using 1,000 servers for one hour costs no more than using one server for 1,000 hours. This elasticity of resources, without paying a premium for large scale, is unprecedented in the history of IT.[3]

2. Nexis.

3 "Above the Clouds" Technical Report UCB/EECS-2009-28, University of California (Berkeley). http://www.eecs.berkeley.edu/Pubs/TechRpts/2009/EECS-2009-28.html. Herein we refer to this effort as the Berkeley Lab.

Marketing is the key driver

With respect to definitions, our dialogue with senior IT architects is quite telling. Here is a sample of what we hear at organizations that require maximum application availability.

We interviewed the Director of High-Performance Grid Computing at a major bank, an accomplished and articulate technologist who is running a large, private cloud. Among other questions we asked, "From your perspective, what is cloud computing and what makes it different from other computing models?" His reply was, "A complete answer would require a 3-day debate and 6-month Ph.D. dissertation."

If technologists are debating what "it" is, and "it" is in the business media and blog buzz, as cloud computing has been, then "it" is driven more by marketing change than by technology change. Technology is relevant, but technology change is not the key driver of the cloud computing concept today.[4]

It's really a utility model

Cloud computing is a new bottle for the old and fine wine of utility computing, with some cloud providers utilizing grid computing concepts in the mix. For purposes of our analysis, utility computing is the most important of these concepts. It is the concept that computing and its associated data storage are available to users in the same way as electricity or telecommunciations services—from "plants," such as power plants, central offices or hubs.

Essence of the utility model

The essence of the utility model is that, at least at the corporate level, the service user is not the service producer. The user does not own all, most, or even any of the tangible assets necessary for service production. The user may own only a small part of the assets for transmission of the service to its point of usage.

Evolution of the utility model

It wasn't always this way in the case of electricity, and it isn't (yet) this way in computing. Factories once powered their machinery with kinetic energy passed through millworks from factory-owned waterwheels in a nearby stream. Thomas Edison's direct-current system replaced on-site waterwheels and millworks with on-site power plants and cable, enabling dramatic improvements in the dramatic improvements in the reliability, precision and flexibility of manufacturing processes.

4 If we were cynical, we would note pension managers' views of hedge funds as "a compensation scheme masquerading as an asset class." Parallel cynicism would suggest cloud computing is a marketing plan disguised as an information processing architecture. We believe, however, that business sobriety, not cynicism, is the best mind-set from which to evaluate cloud computing.

Consistent with the B2B market mind-set of its time, Edison's model did not aggregate energy demand across scores of customers per plant. In tracing a concise history of commercial electricity production and usage in the U.S., **Nicholas Carr notes:**

"In the early years… the presumption would be that a manufacturer electrifying his machinery would use his own power plant." That presumption is evident in the statistics. As the new [20th] century began, a survey by the Census Bureau found that there were already 50,000 private electric plants in operation, far outstripping the 3,600 central stations.[5]

The geographic reach of Edison's system was limited by its direct current technology. The adoption of Samuel Insull's alternating current system enabled electricity production to occur far away from the user. Such a power plant could serve many customers in a large region, with much greater efficiency as scale economies developed. As specializations increased and concentrated in such plants, reliability and efficiency grew even more, attracting more customers and specialists in a positive feedback loop.

There could be a computing plant in your future

Carr argues that computer processing demand will consolidate into "computing plants" that furnish service across a wide geographical area. We believe this is credible if the security and availability concerns of prospective customers are addressed.

Cloud development today is as private as electricity was when the previous century began. We believe all of the Fortune 500 and most of the Global 2000 firms have private cloud prototypes. We also believe that the firms with the most cloud experience are considering public cloud utilization for some of their processing.

With respect to availability, every disaster recovery architecture that we have seen has the fatal flaw that ZERODOWN Software' architecture fixes. Furthermore, some computing clouds today utilize a computing design principle that, when incorrectly implemented, drives exponentially higher catastrophic risk than the alternatives that it is replacing. ZERODOWN Software furnishes a patented architecture and technology to correct the flaw and harden the implementations, for both cloud customers and cloud providers (Chapter 4).

5 Carr, 37.

Why now?

Our stakeholders have asked us why cloud computing is emerging "now." It is a good question—from technologists and executives.

Our IT colleagues claim cloud computing contains nothing new conceptually or technologically. They point to our comments under "Marketing is the key driver" on page 1-3 and say "See, nothing new!" We concur that computing clouds currently contain only codified concepts. We disagree that clouds contain nothing new technologically.

Invention vs. innovation

So why now?
What is new through the lens of the business decision maker, for that is the lens that counts, is integration.[6]

To understand why integration is key, we first distinguish between invention and innovation:

- A new idea has been invented when it is proven to work in the laboratory.
- That idea becomes an innovation when it can be replicated reliably on a meaningful scale at practical cost.

When a concept moves from invention to innovation, diverse component technologies integrate. Emerging from isolated developments in separate fields or firms, these components gradually form an ensemble of technologies that are essential to each others' success. Until integration, the concept, though possible in the laboratory, does not achieve its potential in practice.

DC-3 exemplifies the value of integration

To see the value of integration, consider for context another highstakes, technologically sophisticated endeavor, commercial aviation. In December 1903 the Wright brothers' aircraft proved that powered flight was possible, but it would take more than thirty years before commercial aviation could serve the general public, when the Douglas DC-3, introduced in 1935, ushered in the era of efficient commercial air travel.

6 The remainder of this is based on Senge, Peter. The Fifth Discipline: The Art and Practice of the Learning Organization. New York: Doubleday. 1990. 5, 6, 271, 342, 363.

The DC-3 was the first plane that supported itself aerodynamically and economically. During those intervening thirty years myriad experiments with commercial flight had failed. The early planes were not reliable and cost effective on an appropriate scale.

The DC-3, for the first time, brought together five essential component technologies that formed a successful ensemble. They were:

- The variable-pitch propeller
- Retractable landing gear
- Monocoque, a light-weight molded body construction
- Radial air-cooled engine, and
- Wing flaps.

To succeed, the DC-3 needed all five; four were not enough. One year earlier, Boeing had introduced the Boeing 247 with all of them except wing flaps. Lacking flaps, Boeing's engineers found that the plane was unstable on take-off and landing and had to downsize the engine. That meant a lower payload for Boeing's potential customers, crimping commercial viability of the 247. For the DC-3, designing the engine specification required understanding the effects of the variable-pitch propellers, the flaps, the retractable landing gear and the stress characteristics of the new monocoque body. The wing and body design depended on the engine's thrust. Integrating the component technologies was more essential to the success of the DC-3 than the task of designing any single component.

So our colleagues who say cloud computing contains "nothing new" are accurate when they say it contains no new inventions—but that is irrelevant and leads to career complacency. In a very real sense, technological history belongs to the integrators. Even though in its infancy and challenged by lack of certifications and management controls, cloud computing is a real integration of technologies that are now available at adequate scale and price.

Integrated technologies of cloud computing

Cloud computing comprises component technologies and a utility business model. The technologies and the skill sets for integration are now available at adequate scale and price. This is why cloud computing is happening "now." The key component technologies of cloud computing are high-capacity digital networking and the virtualization of servers and storage. The integration of these technologies on relatively inexpensive and nearly ubiquitous hardware is new.

Digital networking

High-capacity digital network links were in place a quarter-century ago as the long-haul lines of telecommunications providers. What is new is widespread availability to the "last mile," the final connection between the utility and the customer premises. The bandwidth problem was essentially solved in the first year of the 21st century as even consumers ordered "broadband" connections.

> What the fiber-optic Internet does for computing is exactly what the alternating-current network did for electricity: it makes the location of the equipment unimportant to the user. But it does more than that... By providing a universal medium for data transmission and translation, the Net is spurring the creation of centralized computing plants that can serve thousands or millions of customers simultaneously.[7]

Wide-area bandwidth prices have been dropping, though more slowly than CPU and disk storage prices, according to the Berkeley Lab. Table 1-1 on page 1-8 shows their comparison of 2003 and 2008 expenses. The numbers in this space constantly change, but we think the trends are relatively clear.

7 Carr, 60.

Table 1-1
Expense / performance improvements, 2003–2008*

		WAN bandwidth / month	CPU hours (all cores)	Disk
2003	Item	1 Mbps link	2 Ghz cpu, 2 GB RAM	200 Gb 50 Mb/s transfer rate
	Expense $	100	2,000	200
	Units per $1	1 GB	8 CPU hours	1 GB
2008	Item	100Mbps link	2 Ghz, 2 sockets, 4 cores per socket 4 GB DRAM	1 TB 115 MB/s sustained transfer
	Expense $	3,600	1,000	100
	Units per $1	2.7 GB	128 CPU hours	10 GB
	Better by	2.7x	16x	10x

*Berkeley Lab

"Virtualization," below, traces the history of virtualization by system size as we originally described in the Always Available Server Consolidation brief in 2007. Stakeholders who have already read that brief and who do not require a refresher can, without loss of continuity or meaning, skip to the paragraph that begins "After chasing..." on page 14.

Virtualization

The second key component technology of cloud computing is virtualization, of both general purpose servers and of digital data storage.

Virtualization has a lengthy and distinguished history. It has long been supported on the most expensive, highest-throughput computing systems that have delivered the greatest available uptime. Among the commercial pioneers, IBM has supported the Virtual Machine (VM) capability, hosting multiple operating systems such as OS390 on its own processor chips. The Unix operating system, in both uni- and multi-processor configurations from a variety of vendors, has supported virtualization for applications since the 1980s. Until recently, virtualization support was not available on the smallest systems.

Both the VM/OS390 and Unix implementations have showed scalability and throughput superior to small systems running Microsoft Windows. From their industrial strength design requirements, the larger systems have been more expensive to purchase and easier for large enterprises to manage.

The smaller systems, by contrast, have driven low acquisition expense and the managerial migraines of "server farm" and "storage farm" proliferation. Because of the relatively weak ability of the smaller systems to switch between tasks, the application deployment rule of thumb became "one application, one server." Simultaneously, the need to support surges in computing or storage has required peak capacities far above average capacity utilization. These two factors, 1:1 and surge headroom, have yielded low average system utilization and the perception of waste of electricity and real estate expenditures. Capital expenses have been reasonably low and dropping. Operating expenses have been high and rising. Frustration has persisted.

Consistent with the disruption of computer industry structure that accelerated in the 1990s, virtualization is now available on the least expensive class of general purpose processors, whether they run open-source or proprietary operating systems.

After chasing IBM's chipmakers for almost 40 years, Intel and its competitors now produce chips for small systems powerful enough to support virtual machines, at sufficient scale and price required for widespread adoption. Complimentary technology from VMWare®, Citrix® (Xen®) or Microsoft® (Hyper-V™) enables the required low-level ensemble.[8]

Business rationale of cloud computing

Though hyped by IT vendor marketers under the cloud metaphor, technology integration does not fully explain the allure of cloud computing. Amazon® and Google™ are frequently cited as cloud computing exemplars, yet neither is a traditional IT vendor. In addition to the power of technological integration, executives have business reasons to consider cloud computing that are similar to those reasons of factory owners who considered Samuel Insull's alternating current utility a lifetime ago. Those reasons form a superset of the typical outsourcing business case. Industry structure and professional specialization are additional reasons why the utility model gains strength in the eyes of executives.

The outsourcing rationale distills questions like the following, from the prospective customer of the utility:

> This stuff, whether electricity or word processing software, does not make my business unique or even help it get ahead, since everyone else has it, too. It is the ante for the competitive table. Can I get the same or better benefit, provided at least as reliably, for the same or less money considering expense and opportunity cost, from a utility than from in-house capabilities?

Standardization and scale economies

Though not necessarily so a decade ago, that is now a reasonable question for general purpose computing and storage. There is no longer any substantial difference in the general purpose computers and software available to the different types of customers of the IT vendors—or to cloud providers, legacy IT outsourcers or disaster recovery service providers. Most computing purchases merely maintain competitive parity rather than drive decisive differentiation.

8 ZeroNines Always Available™ technology is Hyper-V certified.

Most of the software and almost all of the hardware that companies use today are essentially the same as the hardware and software that their competitors use. Computers, storage systems, networking gear, and most widely used applications have all become commodities from the standpoint of the businesses that buy them. They don't distinguish one company from the next. The same goes for the employees who staff IT departments. Most perform routine maintenance chores—exactly the same tasks that their counterparts in other companies carry out.[9]

Lack of differentiation is not exclusively an end-customer phenomenon, either. The existence of new cloud service providers, whether Amazon, Google or others, does not require the existence of new IT product providers. Though purchase terms differ, cloud builders purchase essentially the same server, storage and networking gear from the same vendors as end customers do.

Most substantive differences are in price. According to Microsoft research summarized by the Berkeley Lab, very large data centers command 80% discounts from the prices typically offered to medium-sized data centers.

Table 1-2 compares data center expense categories.[10]

Table 1-2
Data center size expense differences

	$ expense by size of data center		
Technology category	**Medium**	**Very Large**	**Premium**
Network (per Mbit / sec / month)	95	13	7.1x
Storage (per GByte / month)	2.2	0.4	5.7x
Administration (servers / administrator)	≅ 140	> 1,000	7.1x

The research rates very large centers as containing tens of thousands of computers and medium centers as containing hundreds or thousands. This is not strict categorization, but the pricing difference is sufficiently large to be clear. Figures here assume 1,000 servers in the medium center and 50,000 in the very large center. Microsoft, Berkeley Lab; ZeroNines analysis.

9 Carr, 57.

10 Berkeley lab.

The IT product vendors rely on the same contract or in-house manufacturing capabilities and the same in-house or independent software suppliers. Broad and deep standardization enables integration that is essential to the innovation (not invention) and massive scale economies necessary for, and structured by, a utility model.

Even in a somewhat fragmented industry, independent software vendors have lusted after the utility computing model for years because a server-centric model—network-centric, eventually plantcentric—drives standardization that enables higher quality at lower price. The customer does not need to perform complex software installation on its own "appliances," but instead merely needs to plug those appliances into the network as an extension of the clientserver model.[11] The customer gains the same or better software interaction benefits more reliably because layers of error-prone procedures never arrive on the customer's premises.

Specialization and quality of scope

The utility firm that performs complex configuration procedures eventually builds a team of individual contributors who have mastered the technology and the process. Their error rate drops to a lower level, at a faster pace, than their self-provisioning prospective customers could expect to achieve in their relative isolation.

The utility furnishes technical specialists with a comparatively endless supply of professionally challenging technical problems to solve and a commensurate budget of tools and techniques for solving them—or at least a budget proportionately greater than similar specialists would have in typical user organizations.

Teaming effects are germane because solving IT service problems is at the heart of the IT utility's value proposition. The specialist is surrounded by others with a similar focus and passion for the subject matter at hand. More case examples are available more frequently to technicians inside the utility than to those outside.[12] With reasonable management, the presence of more colleagues means more numerous, tighter and faster feedback loops. Error correction and insight thereby accelerate. Quality of customer service, professional satisfaction and competitiveness of the individual and the utility company increase together.

11 A client is a service requestor; a server is a service provider. A Web browser is a client; a Web site runs on a server. Although we like Carr's contextual work on cloud computing, we believe he casts the client–server model narrowly.

12 Web-based peer dialogue forums remove some of the tilt from the professional playing field, but do not level it.

Comparing in-house vs. cloud uptime for the average firm
Application uptime will be higher with a utility model for the vast majority of businesses. Intuition suggests that uptime is the business of a computing utility, even if that utility is unregulated.

A quick calculation points in the same direction:

- IDC research indicates targeted annual uptime of less than 99% for 2007, the latest year for which figures are available, in a survey conducted for its client SunGard®, a privately held disaster recovery service provider. The average number of employees per firm in the survey is 3,298.

- For perspective, according to the U.S. census, approximately 80% of U.S. business establishments have fewer than 500 employees and, we infer, disproportionately less IT support than larger firm.

- We conclude that average uptime in the IDC survey is a likely upper bound on the true average uptime by firm across U.S. business.

- Amazon's cloud computing service level agreement, to take one example, implies 99.95% availability, less than 5 hours of downtime per year—dramatically greater than average, though still unacceptable to clients of our Always Available architecture and technology.[13]

Opportunity cost and hybrid models
Every investment decision drives an opportunity cost. Because economy of scale and quality of scope from skill specialization are natural consequences of a properly governed utility model, customer opportunity cost with utility computing is likely to be lower than the self-provisioned alternative for commodity applications. The utility invests in skills that are appropriate to its core business while its customers invest in skills appropriate to their own. Cloud customers are thus likely to pursue hybrid models, retaining proprietary applications in-house and accessing most others through off-premises models such as cloud computing.

I really do think that the winning model here is not going to be everything in the cloud or... everything in the data center. I think the hybrid model is the big winner.[14]

Having mentioned comparative opportunity cost, we recognize that "now" is still early in the game for the cloud providers and their customers. Amazon's EC2 beta outages received extensive industry media coverage; Microsoft's pre-production Azure was down 22 hours during the weekend of March 14, 2009. But there is no more stark reminder than Google's Gmail outages of 2008 (Table 1-3). Recognize, however, that the uptime that we posit for the average business in the U.S. is lower than implied by the published production SLA's of public cloud providers.

13 IDC, 2006: "Optimizing Business Performance Requires Optimizing Information Availability Investments." Amazon SLA: http://aws.amazon.com/ec2-sla/. fBusiness census: http://www.census.gov/compendia/statab/tables/09s0736.xls.

14 Joe Tucci, CEO, EMC Corporation, quoted in "Microsoft and EMC renew their vows"4 February 2008, cnet news.

Table 1-3
Google Gmail outages

Date of Gmail outage	Reported duration (hours)	Comments
16 July 2008	2.5	4:30 a.m. Eastern time, affecting business everywhere but the Americas
6 August 2008	15	Struck early-afternoon Eastern time, wiping out the rest of the day for the Americas and the start of business in Asia.
11 August 2008	2	5 p.m. to 7 p.m. Eastern time
15 August 2008	> 24	Every time zone Monday
16 October 2008	30	More than one business day worldwide. "In companies that were affected, Apps administrators told of very tense situations, in some cases involving having to deal with extremely upset CEOs and other high ranking executives who got locked out of their e-mail."
24 February 2009	2.5–4	Business morning in Europe, business day in Gulf states

*Google, *PC World*; ZeroNines analysis.

Asset management & financial benefits

Apart from skill specialization benefits, the question of ownership of tangible assets is another distinguishing dimension of cloud computing. Ownership of such assets has rewards and risks. In a customer's transition of some applications to the cloud, many if not most of those risks and rewards transfer to the cloud provider.

Asset ownership in legacy outsourcing

In legacy IT outsourcing, tangible assets such as servers, storage and network routers can be owned by the client while operated and maintained by the service provider. That business model is about IT skill specialization on the provider side and, occasionally with disaster recovery service providers, liability deflection by the client's officers.

In many legacy outsourcing contracts the client's financial capital is trapped in the tangible assets that the outsourcer manages on the client's behalf. Upon contract conclusion the client owns tangible assets, the IT gear, that are fully depreciated and technologically a generation or two behind the times. The client's people have learned little or nothing about the management of those assets along the way, though they have been able to focus on other issues such as the client's core business.

Asset ownership in cloud computing

In the cloud computing model, as in a utility model, asset ownership by the service provider minimizes the client's capital expenditures for the provisioning of service. The client utilizes operating expenditures instead, preserving capital for investment in differentiating projects.

Though a vivid selling point at any time, capex → opex conversion is even more attractive as corporate balance sheets compress in The Not-So-Great De-Leveraging. Prior to the financial crisis that began in 2007 and spread outside the financial services industry in 2008, capital had been inexpensive for several years. We don't see it as accidental that the utility computing model—a shift in asset ownership risk—gained greater attention during the first global monetary panic in a century.[15]

How virtualization drives asset management benefits

How does virtualization drive asset management benefits?

The most commonly cited financial effects of virtualizations are constrained growth of capital expense and operating expense. In this view, the traditionally calculated financial present value of the firm is higher with consolidated infrastructure as enabled by virtualization. This business case for virtualization, unadjusted for tail risk, has proved almost irresistible for the large enterprises that have considered it.

Figure 1-1 graphically depicts the asset management and financial benefits of the virtualized server and storage functions.

The remainder of this section traces virtualization history by system size as in our Always Available Server Consolidation brief from 2007. Stakeholders who have already read that brief and who do not require a refresher can, without loss of continuity or meaning, skip to "Who is doing this?" on page 22.

15 Prior to the week of September 12, 2008, the highest ever spread of the 90-day commercial paper rate above Federal Funds rate was 1.48 points, a record from 1971. Starting in September 2008 the record rose eightfold to 12.4 points, and the spread stood near 6 points as we wrote this in 1Q2009. In the autumn of 2008, as money market funds "broke the buck" (reduced below $1 their net asset value per dollar of client investment), a prospective investor approached ZeroNines as a place for "safe keeping" (investor's phrase) a portion of his wealth.

Figure 1-1
Benefits logic of virtualized infrastructure

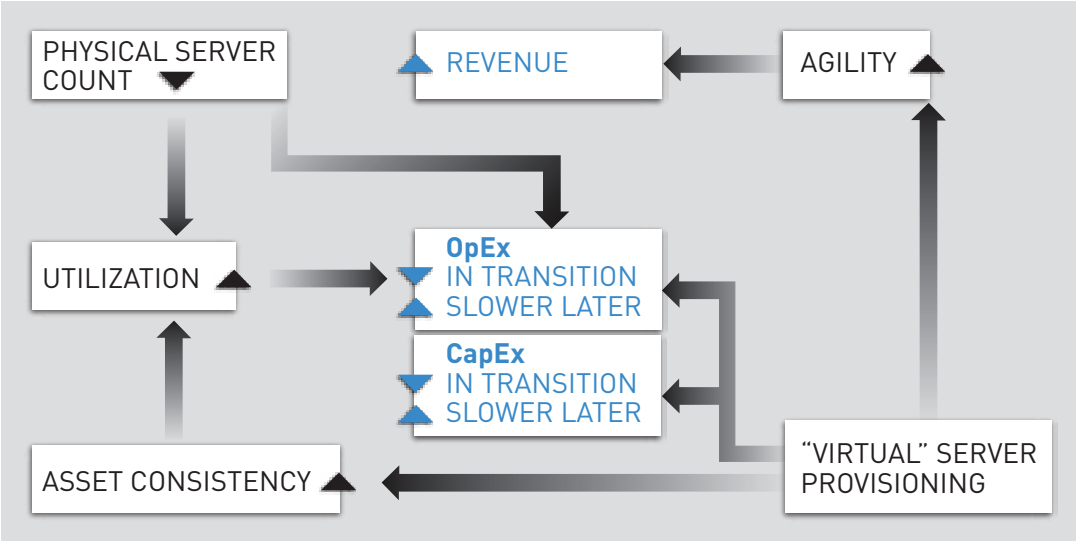

The pillars of IT virtualization business value are:

- Reduced quantity of physical assets
- Increase in consistency of logical assets
- Virtual asset provisioning.

The assets are the tangible capital assets, such as servers and storage array units, and relatively intangible "human capital" assets such as procedures for provisioning and maintenance. The latter can be made more tangible and repeatable through knowledge management services and, later, process automation software.

Details follow. Readers who do not require details can, without loss of continuity or meaning, go to "Who is doing this?" on page 22.[16]

Quantity reductions in physical assets
Decreases in physical asset quantity reduce operating expense through lower real estate and utility expenses. New multi-core processor chips and higher data-density disks cost less per application load per watt. Greater power at higher utilization means less silicon and metal are required to perform a given amount of work. Less hardware is required, so less space is needed to house it.[17]

16 Many of our stakeholders will recall coverage similar to the following from our Always Available™ Server Consolidation brief.

17 ZeroNines has also seen research showing data center wattage requirements can be reduced even more with clever cabinet design, replacement of traditional rack fans with more reliable blowers, deactivation of unneeded chips from system boards, and fluid dynamics studies."Keeping Your Cool in the Data Center while Consolidating and Virtualizing your IT Infrastructure." Appro International, Inc., 2006.

Asset consistency
Increases in asset consistency drive higher utilization of the assets that remain. There is less variety and more efficiency. The operating expense efficiency increases again as asset consistency and asset quantity reductions interact.

Virtual server provisioning enables more benefits
Virtual asset provisioning is a dramatically important benefit of consolidation. Instead of days or weeks to prepare a new physical server for a business application, per with the "one application, one server" rule, the IT function can prepare a new "virtual" server in a matter of minutes. The variable cost of enabling a new application is slashed dramatically even as the consistency and quality of server provisioning increases. Similar benefits are enabled with virtual storage provisioning.

Provisioning expense growth slows
Physical provisioning expenses grow slower later due to the ability of the new infrastructure to scale by means of the virtualization features of the processor chips, operating system and disk array controller software. Virtual provisioning permanently reduces unit expense while constraining capital expense growth.

Agility raises revenue growth curves
Faster provisioning means greater business agility and uptime. For many business executives, this translates directly to higher enterprise revenue. Entire revenue growth curves shift upward simply through the power of rapid provisioning.

Flexibility helps the back office
Virtual provisioning enables a great deal of behind-the-scenes flexibility in the IT function. Server and storage unit images can be easily provisioned and decommissioned for testing, maintenance and other purposes that are practically invisible to senior management, yet essential to the success of better business models.

Though failover is the flawed core of disaster recovery, one benefit of this flexibility is failover of virtual machines within the data center using the "movement" feature of some virtual machine management software. Within the confines of the disaster recovery paradigm, this is a benefit. We have not seen virtualization software that replicates transactions or otherwise supports failover between data centers.

Who is doing this?

Who is doing cloud computing today? What are customers doing with it?

To address these questions, we develop an industry value chain model in the form of an "XaaS" framework, where "x" varies by layer in the model. We are not the only ones doing so. Every XaaS framework is subject to substantial revision as chaotic market signaling intensifies from myriad pure-play and conglomerate firms while consultative organizations jockey for position. Some of the models that we have seen from the consultancies are abstruse.

The healthiest way to approach exercises like this is to recognize that it is early in the game for everyone except the big IT product vendors and their largest legacy customers—and personnel turnover has left even them with few of the seasoned hands who midwifed the service bureaus decades ago.

XaaS value chain concept introduction

The implicit flow through our relatively simple XaaS value chain, starting upstream from the end customer and moving downstream, is:

1. **IT Product Vendors**
2. **HaaS** (hardware as a service)
3. **PaaS** (platform as a service)
4. **SaaS** (software as a service)
5. **End users.**

We begin our value chain analysis with the product vendors as HaaS providers. Some vertical integration strategies start at HaaS.

Product vendors as HaaS providers

Established IT product houses such as Hewlett-Packard, IBM® and Sun™ have announced strategies focused mainly on enabling others to provide cloud services. These are essentially defensive or conservative pick-and-shovel strategies, reasonable first moves for these enormous firms that have much to lose as well as to gain from the adoption of cloud computing.[18]

HaaS providers are typically the legacy IT product vendors or long-established and large customers of those vendors. Disaster recovery service providers on that traditional model are HaaS providers. The newer HaaS providers focus most of their marketing on would-be PaaS providers, whether those PaaS players intend to offer public or private (internal corporate) clouds. Today's "cloud" offerings from HaaS firms usually are marketing overlays on their substantial legacy capabilities. Some of these companies will direct new investment to consultative or professional services for the build-out, movement and optimization of large data centers—which will, of course, be rebadged as cloud computing centers, cloud banks, or with other mildly amusing monikers. We hope the cloud banks stay solvent.

PaaS is the main cloud layer today

As large as the HaaS firms are, it is the PaaS players who get the most attention as cloud providers today. It is also the locus of most "real customer" usage examples in the business and trade media.

The main audience for the marketing literature from PaaS providers is software development and deployment organizations, whether independent software vendors, the development departments of large enterprises who might need to populate their private clouds, or SaaS players.

Platform-as-a-service (PaaS) is an intermediate layer of functions between HaaS and SaaS and the principal cloud offering of Amazon, Google and Microsoft. Their cloud offerings are not identical, and the differences are highly relevant to prospective customers.

The Berkeley Lab team argues that cloud offers are distinguishable by the degree of abstraction presented to the programmer and the degree of management of the resources. As technical treatments go, this is a plausible starting point. Choices along these two dimensions drive trade-offs in system design. Stated coarsely, there is a PaaS service spectrum with flexibility on one end and programmer convenience on the other. We adopt Berkeley's analytical lens to describe PaaS offerings here and to summarize comparisons in Table 1-4 on page 25.[19] The firms are discussed in alphabetical order.

18 In a gold rush, those who sell picks and shovels have the highest probability of turning a steady profit.

19 We recommend the Berkeley Lab paper to readers with technical backgrounds. The remainder of this PaaS section is based on it: http://www.eecs.berkeley.edu/Pubs/TechRpts/2009/EECS-2009-28.html.

Amazon's Elastic Computing Cloud (EC2)
Amazon's EC2 platform virtualizes with a server metaphor, enabling explicit management of images that logically start from the operating system kernel and proceed upward to user space. A programmer or operations technician provisions a virtual server approximately equivalent to a 1GHz x86 machine in two to five minutes and controls it with a remote desktop program at a price of $0.10/hour.[20] More powerful resources can be provisioned. Storage services are also available on a pay-as-you-go basis.

The virtual server resources to be managed are low-level, such as CPU cycles and IP connectivity. The low-level approach enables developers a great deal of control. It also ties state management to the application in question, frustrating automatic scalability and failover.[21] Higher-level services are available with higher latency and nonstandard API's.

Google App Engine™
On the flexible end of the spectrum are platforms that enable particular application domains, such as:

- Force.com™, the development platform from salesforce.com.
- Google's App Engine.

Consistent with Google's origins, App Engine enables traditional Web applications that separate stateless computation based on a request-reply protocol from a stateful storage tier. Scalability and failover are relatively automatic.

Flexibility ends there, however. "Particular application domains" entails constraints, notably:

- App Engine limits the amount of CPU time that an application can use per request.
- Force.com is designed for usage of the salesforce.com database, period.

Microsoft Azure™
Between the fluid scalability of Google and the programmer's precise control in Amazon is Microsoft's Azure service. Azure leverages Microsoft's .NET library and is compiled to a language independent managed environment. In this sense Azure supports general purpose computing in which the developer can choose a language but not the operating system, and is required to specify application parameters to enable scalability and failover services.

20 Pricing as of 1Q 2009.

21 ZeroNines recognizes that failover is the default design approach for application resilience, so we understand that it is naturally considered in analyses today. We analyze weaknesses of the failover architecture later in this brief. Customers of our Always Available™ architecture believe failover is fatal.

Table 1-4
Virtualization models across leading cloud providers*

	PaaS Provider		
Model	**Amazon**	**Google**	**Microsoft**
Computation	• x86 Instruction Set Architecture via Xen VM • Elasticity enables scalability, but developer builds machinery or a third party (e.g. RightScale) furnishes it.	• Application structure and framework based on 3-tier Web model. • Programmer writes "handlers" in Python, with persistent state storage in MegaStore facility outside Python code • Automatic scalability and failover.	• Common Language Runtime VM, a common intermediate form in a managed environment • VM's provisioned from declarations (e.g. "roles") • Automatic load balancing
Storage	• Multiple models available, from block store to key/blob, which drive scaling and consistency traits • Scaling from none to fully automatic • Consistency assurance varies widely • API's vary from standard to proprietary	• MegaStore / BigTable	• Restricted version of SQL Server capabilities • Azure storage services
Networking	• Declaration of IP topology without internal placement detail • Node communication may be restricted with Security Groups • Independent network failure abstracted with availability zones • Persistently routable network names enabled with elastic IP addresses	• Fixed by 3-tier structure • Automatic scaling is invisible to programmer	• Based on declarations of application "roles" (components).

*Berkeley RAD Lab; ZeroNines analysis

SaaS

Relying on the PaaS players or on traditional hosting are the Software as a Service (SaaS) firms. SaaS providers typically focus on end customers and sometimes sell outside the IT functions of their targeted prospects. Salesforce.com has targeted marketing and sales executives; SuccessFactors focused on senior human resources executives until 2008, when it seemed to begin a CxO strategy.

These examples represent sales of service for what their customers call "real work," something at the core of their operations or business processes. Salesforce.com furnishes browser-based sales process automation. SuccessFactors furnishes browser-based talent management processes such as goal management, performance appraisal and 360/multi-rater feedback. The services of each firm are configured to some degree to the customer's requirements.

Data confidentiality is misunderstood

Salesforce.com and SuccessFactors strike at the heart of the confidential-data myth, the notion that prospective customer companies will not adopt cloud computing because of concerns about data confidentiality. A brief check on the histories of these companies is clear counter-evidence. Salesforce.com stores sales process data that every commercial organization views as extremely sensitive. SuccessFactors stores talent management data that every organization, commercial or not, views as extremely sensitive. Both firms have been thriving. The fact that they have thrived suggests that data security is simply another design requirement of, not a fundamental barrier to, adoption of cloud computing, at least by customer firms that are free of pertinent regulatory proscription.

Data confidentiality matters. It is mostly likely to be driven by governmental demands, at least in democracies, either in the form of:

- Industry regulation, where in many cases there are well established patterns and market understandings, such as payments processing and the banking industry; or in the form of
- Administrative procedures established in law or managerial systems, such as evidentiary chain-of-custody in law enforcement and chain-of-provenance in intelligence analysis.

Professor Benn Konsynski of Emory University notes:

"Exposures associated with data content must be addressed. The user must trust the cloud. The cloud provider needs to maintain high transparency. Tight logging and trails are required, especially if the customer faces issues of data leakage and requirements for eDiscovery."[22]

22 ZeroNines interview, 18 February 2009.

Usage examples

In concluding our tour of the XaaS value chain with end users, we believe cloud customer usage categories are unlikely to remain stable as the latest manifestation of utility computing unfolds. So far, many customer case examples with publicly available quantification revolve around predictable fluctuations in resource needs. Some examples are summarized in Table 1-5 on page 28. Some of the principles in play are as follows:

- One-time fluctuations in a company's processing load might, to some meaningful degree, be predictable, as the Washington Post case shows.
- Routine end-of-period processing, such as annual tax-oriented scenario planning, can be considered.
- Project-oriented development, such as backtesting of algorithmic trading concepts by quantitative funds, tends to be CPU- and disk-intensive but not a 24 x 7 requirement by many.
- According to Intel®, software applications are easier to design for scalability for cloud infrastructure. The firm has seen case examples of 10% to 60% reductions in software maintenance and licensing expense in cloud alternatives.[23]
- Applications that benefit from multi-threading or other parallel processing concepts are likely to realize great gains on an "infinitely scalable" infrastructure. Privacy concerns aside for a moment, paycheck processing and Federal income tax return preparation are heavily "seasonal" and utterly parallel from service provider and regulatory perspectives.
- Bursty demand on Web sites is another category of examples in which the site owner might have a sound idea that traffic will increase after, say, a television advertisement during the Superbowl, but cannot quantify it precisely or pay the scale premium that would be required in the absence of a "cloud bursting" alternative.[24]

ZERODOWN Software also believes that the "infinitely scalable" promise of cloud computing will be embraced for contingent demands, such as the failover model for disaster recovery—but only if the fatal flaw of failover is fixed. As we have noted, cloud computing per se does not solve failover. Within the confines of the failover approach, however, the economics of disaster recovery, for the service provider and the customer, probably improve with a cloud approach.

23 ZeroNines interview with Jake Smith, Advanced Server Technologies, Intel Corp, 17 February 2009.

24 IBM cites numerous categorical examples here: http://tinyurl.com/aab3j7.

Table 1-5
Quantitative case examples

Firm	Essentials
Washington Post	
	• Major media company • IT department needed to convert 11,000 low quality .pdf pages of Hillary Clinton's travel records from the National Archives into a quickly searchable file set for editorial staff. • In-house gear would have required 30 minutes per page conversion time, too long for deadlines. Launched 200 server instances at a cloud provider, cutting processing to 60 seconds per page and delivering searchable files within 9 hours.
FlyMiwork ‡	
	• Early-stage air charter reservations firm. • Seat prices are calculated dynamically, requiring nimble software. • Excluding facility and personnel expenses, data center expansion to support new markets was priced at $250,000. • Switched to a cloud provider and expects $28,000 in analogous first-year expense.
TC 3 Health ‡	
	• Checks insurance claims for duplication or fraud. • A client asked for one-time scanning of 20 million old records. • Data center expansion would have required $1 million for this one-time job. • After switching to a cloud provider, has spent less than $1 million and is now more responsive to future customer requests.

† http://tinyurl.com/dzzre2. Searchable archive: http://tinyurl.co m/3a7mf6.

‡ "Technology: When the Forecast Calls for Clouds," Michael Fitzgerald. Inc.

Challenges

SECTION 2
Challenges

INTRODUCTION

Challenges of cloud computing

Cloud computing conveys both promise and peril. We explored the promise of cloud computing in Chapter 1.

We now turn to the challenges that cloud computing conveys to its providers and users. Challenges can be classified as operational or existential.

Operational challenges

In broad strokes, the operational challenges of cloud computing today fall into these categories:

- Multitenant implications
- Management problems.

The challenges posed by multi-tenancy are the performance and security implications of virtualization, which is the chief enabler of multitenancy.

An obvious challenge of multitenancy is lack of performance assurance to the user. Every virtualization-based IT business model, whether the traditional disaster recovery service provider model or the cloud model, relies for financial leverage on the oversubscription of assets. The provider bets that not all subscribed customers demand access to all expected resources at a given moment. In many or even most cloud cases this is a reasonable bet for the provider and the customer. It cannot be a reasonable bet for a cloud user who requires an ironclad service level agreement, nor do we think it reasonable for business continuity, a topic that we explore in Chapter 2.

Unlike the security implications of multitenancy, we believe the solution to the performance assurance problem is mostly in the cloud user's hands through the choice of hybrid models. Customers of the XaaS chain will naturally choose to keep in-house those applications that require performance assurance, while moving other applications partly or completely into clouds if their business continuity requirements are met. Such movement is beginning. It is as sure as the transformation from plant-specific DC electricity generation to utility-based AC generation in the prior century. ZERODOWN Software enables business continuity of such hybrid models with our CloudNines™ offering, which we introduce in Chapter 4.

Unlike the performance assurance issue, solutions to the security challenges of multitenancy are more evenly shared between cloud user and provider, or fall predominantly in the provider's problem set. The precise center of gravity depends on the application domain and user preferences.

Clients are reasonably concerned about data confidentiality, whether their data storage is in- or out-sourced. Assurance that data in one virtual machine is protected from encroachments by another is required, whether encroachment would be accidental, as in a system failure, or deliberate, as in an insidious attack. A more complex security scenario develops when a virtualization enabler, such as a hypervisor, fails with the virtual machine still intact. At that point it may be reasonable to ask, "Who controls that virtual machine?"

The latter scenario spans the boundary of security and management concerns. Beyond the security challenges, cloud implementations pose management challenges. The systems management provisioning from cloud providers today is beginning to emerge, but is neither standardized nor robust. It is beginning.

How do we feel about these operational challenges? ZERODOWN Software monitors provider progress against these security and management challenges at multiple points in the XaaS value chain. We monitor providers' market and industry signals for problem awareness, understanding, and commitment to solutions. In some cases we contact providers directly to ensure that we understand their intent and capabilities. We are also following developments at the Cloud Computing Interoperability Forum, where we are a sponsor, and the Unified Cloud Interface project.

On the basis of our experience and continued monitoring, we are satisfied that the XaaS value chain participants will successfully address the operational challenges of security and management.

ZERODOWN Software' main concern for prospective cloud users and providers is not the operational challenges. We are much more concerned about the existential challenges of cloud computing. Those are the challenges that our architecture and technology address.

Existential challenges

When we say "existential" we mean risks or exposures that, if inadequately addressed, threaten the survival of the cloud provider or user business. Existential challenges are about tail risk.

The existential challenges of cloud computing are not new because, as we suggested under "Standardization and scale economies," cloud component products are not new. We have seen this before.

The existential challenges of cloud computing, to cloud providers and therefore cloud users, are the:

- Disaster of the disaster recovery paradigm
- Catastrophic risk increase from poorly architected usage of virtualization, notably in server consolidation.

It is these existential challenges that are the subject of the next two chapters.

INTRODUCTION

The Disaster of Disaster Recovery

The first existential challenge of cloud computing is the disaster recovery paradigm. Service uptime must be greater from a utility provider to its customer than do-it-yourself alternatives or the utility model fails.

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

The problem with the cutover archetype is that it requires an event that halts the business. The disaster recovery architecture, which uses the synonym "failover," is based on the cutover archetype and suffers from similar risks that are amplified by disaster trauma.[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, a 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." See Senge, 21.

2 Dorian Naveh, Director, Product Marketing.

3 Conversation with ZeroNines, Benjamin Taylor, Chairman Emeritus, Disaster Recovery Institute

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 assumptions, organizations surrendered in advance and accepted the weaknesses of the DR paradigm, much as stock market investors accepted the buy-and-hold mantra. But stakeholder expectations have risen and continue to rise, not only because people can be impatient, but because they pursue growth, improvement and excellence.

ZERODOWN Software' 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 many organizations go down because of the limitations of widely used DR implementations. ZERODOWN Software has developed the patented Always Available method and architecture to enable real vendor and platform-agnostic business continuity. Our technology is fully compatible with leading server virtualization products.

To paraphrase Sam Nunn, former US Senator and Chairman of the Nuclear Threat Initiative: if an application outage damages our cloud, 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 chapter. We first examine the value of 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.

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

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? Business continuity is 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 2-1
Indirect vs. direct losses, financial services firm crises (McKinsey)

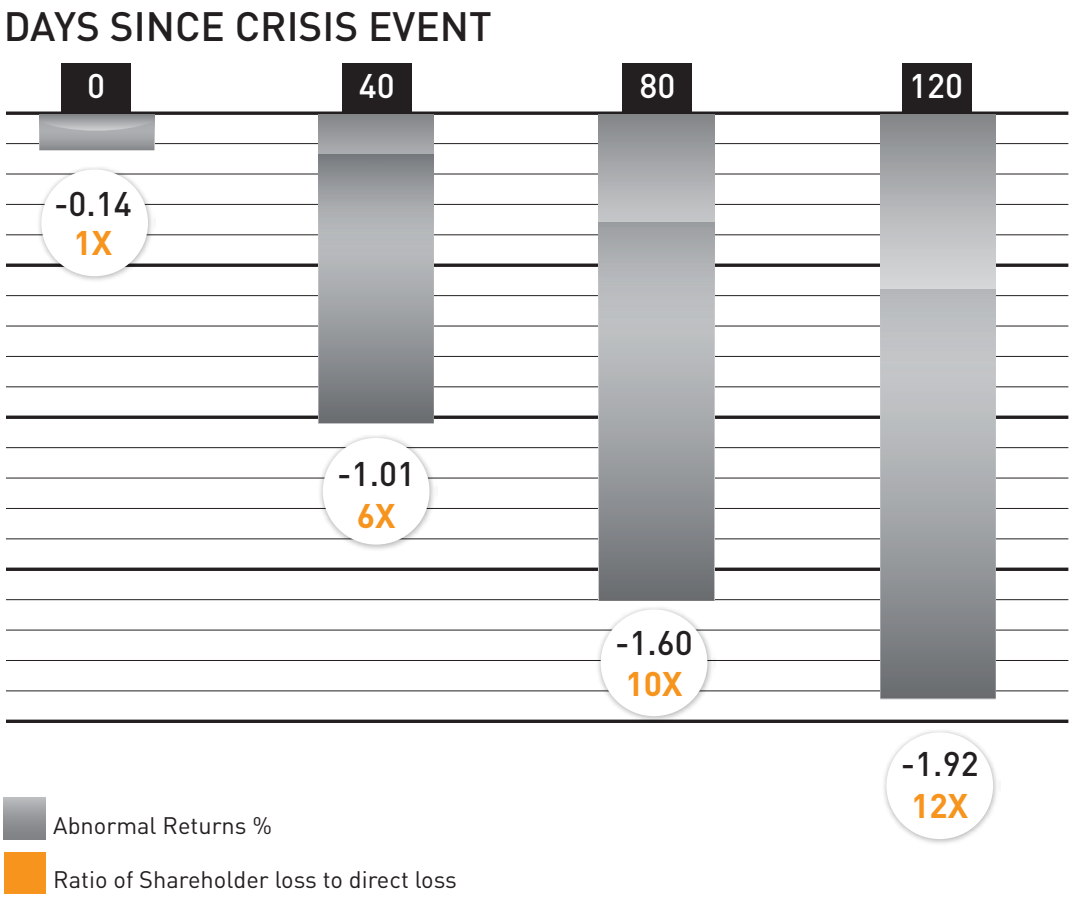

Average direct loss is equal to -0.16% of shareholder wealth, so the 0-day indirect impact of -0.14% rounds to 1x the direct impact. Indirect loss metastasizes to just under 2% of shareholder wealth over 120 working days.

1 The study assessed 350 events since 1990 from Fitch Risk Management's OpVar Loss database. Events were classified with guidance from the Bank for International Settlements."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 truck 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.[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 plausible.

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.

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. And 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 will face the 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 2-2 depicts IDC research indicating a 53% reduction in commercial expectations of planned + unplanned downtime through CYE2007.

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

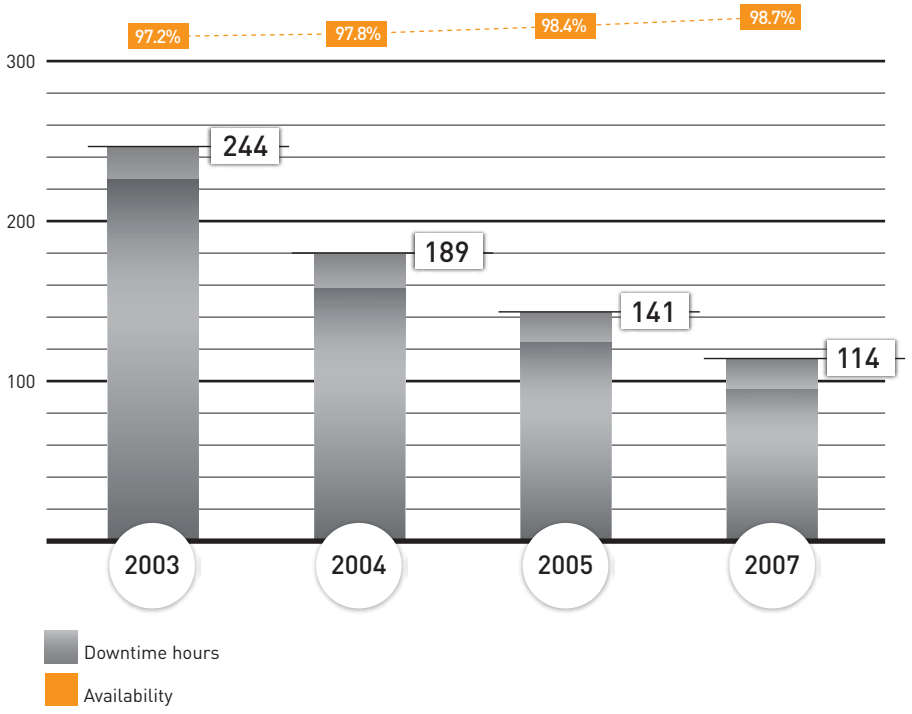

7 The study omits 2006 data. "Optimizing Business Performance Requires Optimizing Information Availability Investments." IDC, 2006.

The threats

Table 2-1
Threats summary
(ZeroNines)

Given the value of business continuity—of disaster avoidance—what threats must be recognized? We summarize the breadth of the threat universe in Table 2-1.

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 retardant 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 ZERODOWN Software Always Available architecture and technology has remained available to its application clients' network 100% since implementation. There has never been a case of an Always Available application client failing to reach its Always Available 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 the Orlando node of the ZeroNines Always Available infrastructure. 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 failure of failover

ZERODOWN Software believes that existing disaster recovery designs are weak. These weaknesses aren't the fault of IT departments, but flaws propagated by vendor designs that have been present for years.

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 2-3):

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

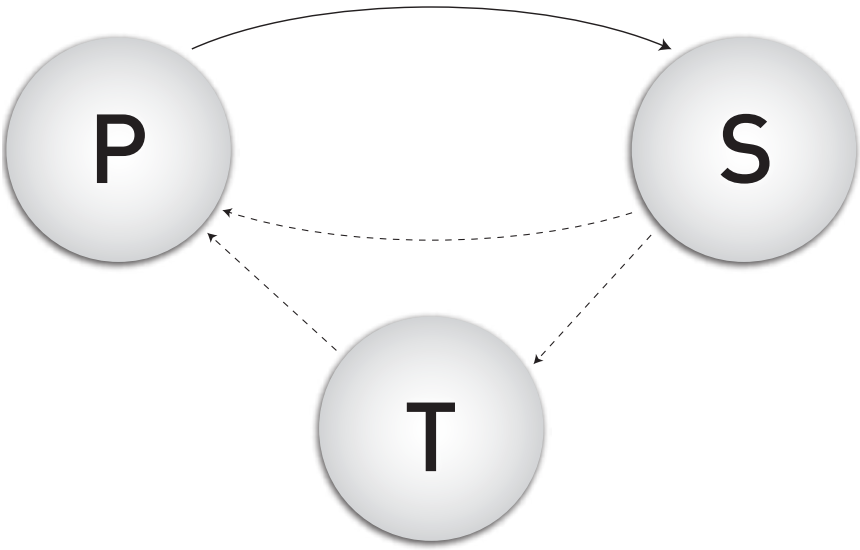

Figure 2-3
At least two cutovers per disaster

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 2-4 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 2-4
Why downtime is inevitable with disaster recovery architecture

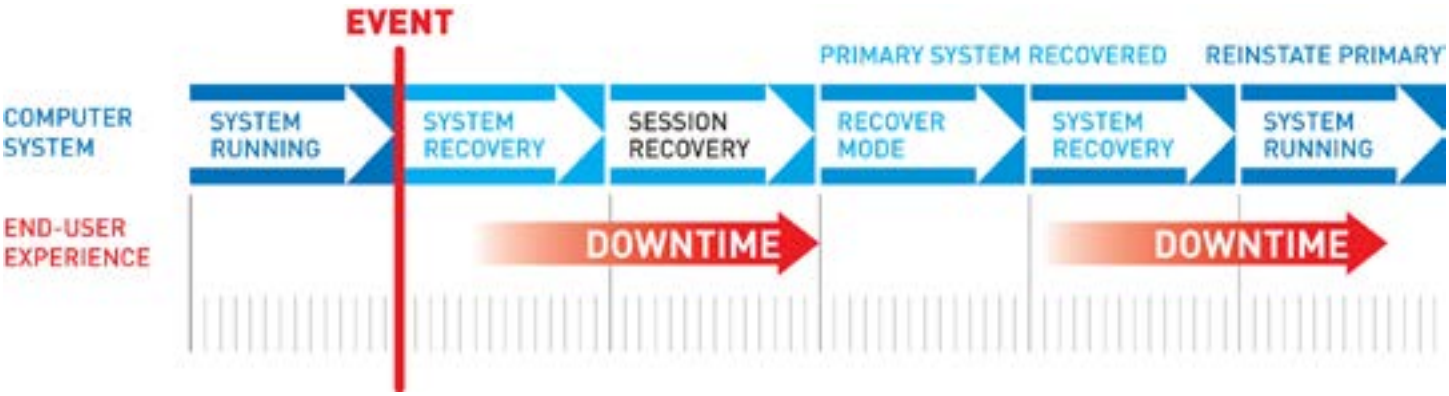

What's recoverable from recovery?

Before 9/11—indeed, before Hurricane 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 (page 40)
- Remote vaulting recovery (page 43)
- Failover and clustering recovery (page 45).

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 2-5 and described as follows.

Figure 2-5
Tape-based recovery architecture

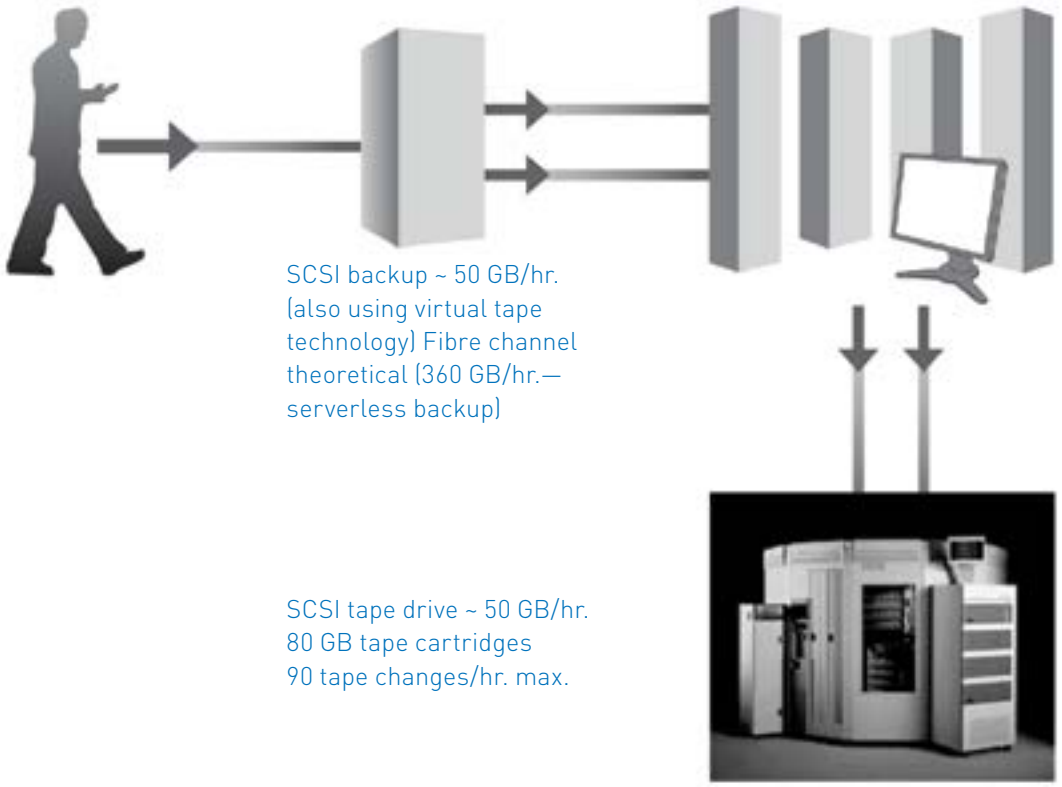

1: Periodically, backup copies of essential business data are produced at the primary site and transported to an offsite storage facility. For 90% of Global 1000 firms that have used failover services,[8] each backup copy utilizes myriad magnetic tape cartridges, each about the size of a paperback book.

2: The primary site fails.

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 seconary 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 38, 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 2-2
Tape recovery attempt time flow

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

8 GartnerGroup.

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 2-3 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 2-3
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 and the threat event is potentially lost. This appears to be the central flaw.
- Tape inventory management must itself 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 the service level agreement queuing of DRSP's, 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, 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 Earthquake In January 1994.

Remote vaulting recovery

Attempting to address the weaknesses of tape-based recovery, vendors now support remote vaulting with split-mirror imaging (Figure 2-6). Vaulting has the advantage of reducing data transportation risk to practically zero by utilizing highly reliable telecommunications networks.

Figure 2-6
Remove vaulting with split-mirror imaging

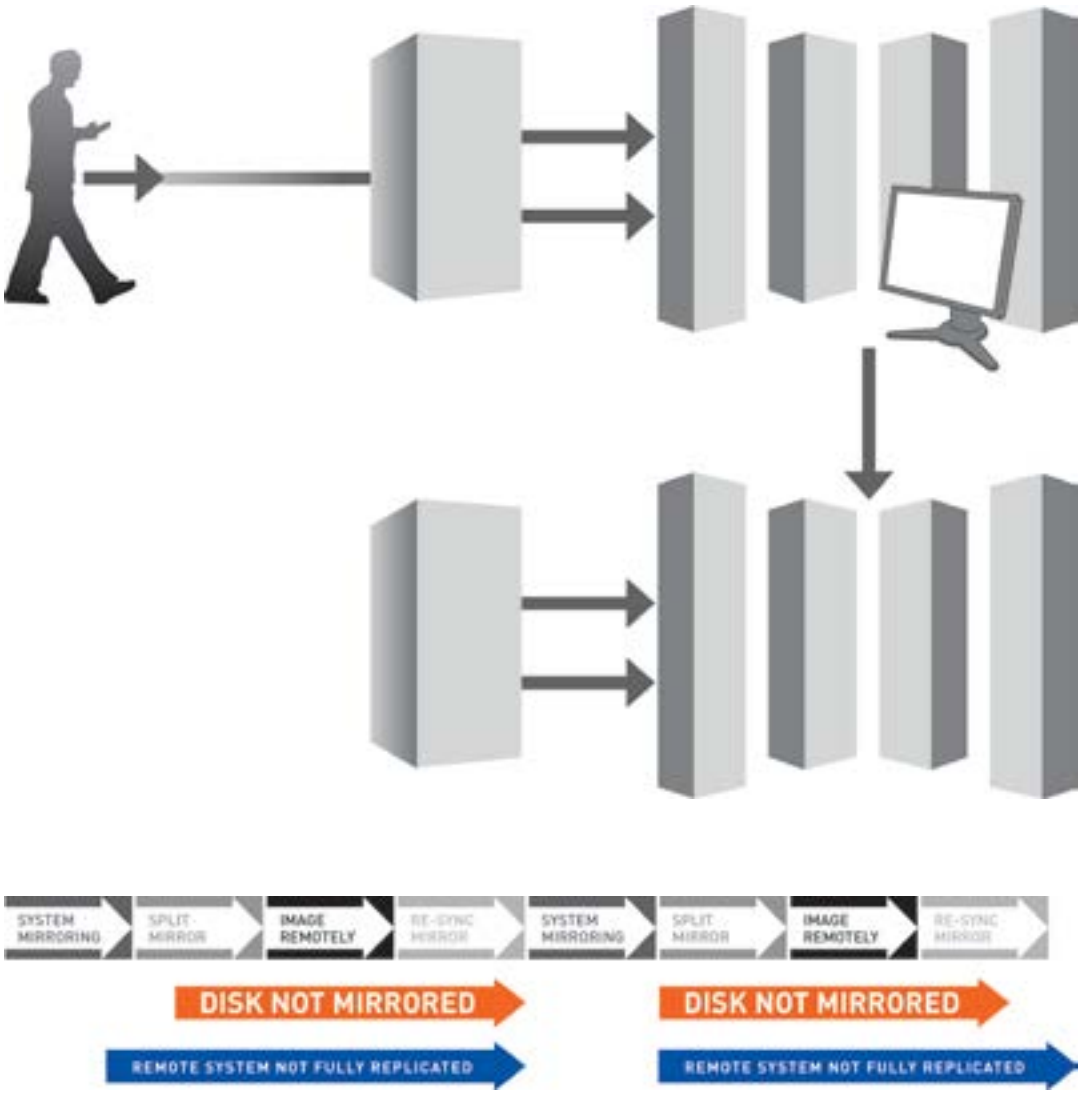

Data written to DASD on the primary system is mirrored locally.

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 reestablished and re-synchronized.

A representative timeflow of a remote vaulting-based recovery attempt is as follows:

Table 2-4
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 are 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 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 explicitly 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 2-7
Server-based failover and clustering

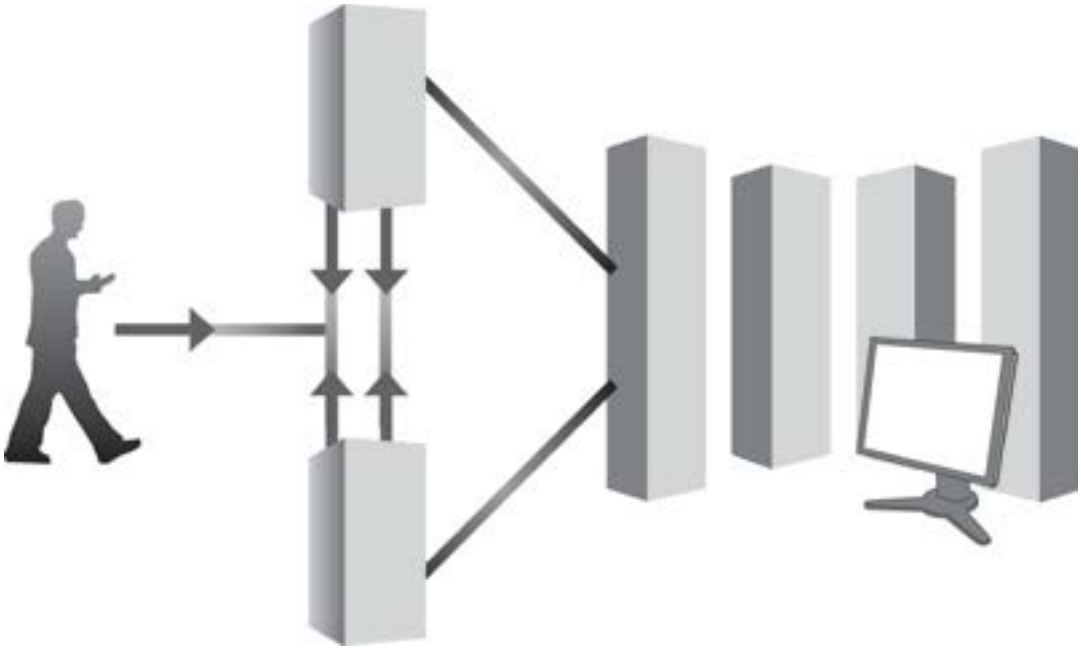

The key failover method is as follows: The secondary node monitors the primary node through a "heartbeat" connection. When the primary fails, the secondary takes over processing. Application sessions are thus maintained. Usually the primary and secondary share disk space, and the distance between servers 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. Disaster recovery architectures have fundamental design exposures that cannot be worked around. IT organizations cannot circumvent the weaknesses with clever and diligent implementation of either the DR architecture or virtualization. Disaster recovery designs are indadequate to support continuous application availability. The two-hour rule and the dispersal rule are not jointly satisfied by any alternate commercial disaster recovery technology from any other leading service provider or vendor.

Table 2-5
Problem summary

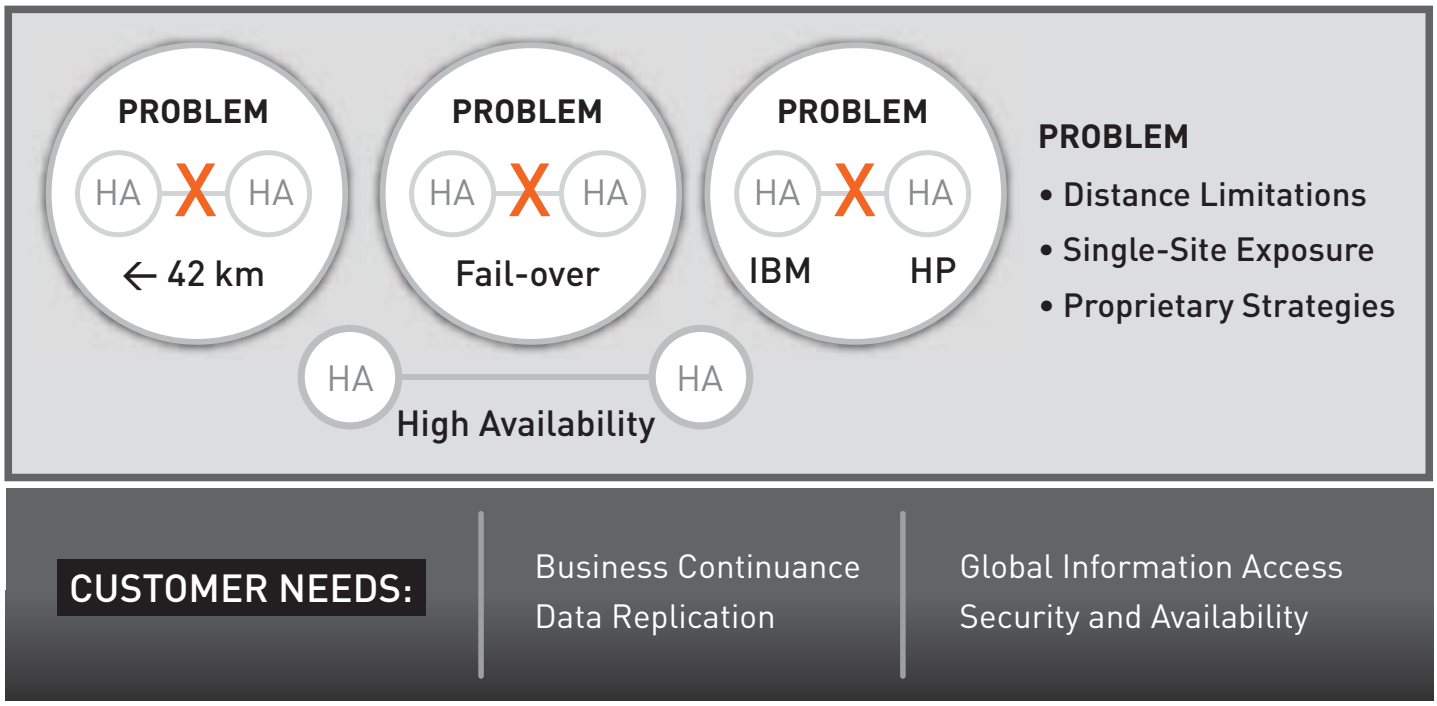

Sungard, a disaster recovery service market share leader that was taken private in August 2005, issued a press release in response to the draft Interagency guidance along these same lines. **Sungard wrote:**

> [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....[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.

Challenges

CHAPTER 3
The Catastrophe of Consolidation

INTRODUCTION

The Catastrophe of Consolidation

The second existential challenge of cloud computing is poorly architected usage of virtualization. As an IT design principle, virtualization is indispensable. So is the avoidance of application catastrophe.

Recent executive attention has focused on server consolidation, the coalescing of numerous services on fewer computers running at higher average utilization. Diminished expense growth and augmented agility attract this investment.

ZERODOWN Software believes executives have received an incomplete picture of the risk-adjusted present value of virtualization, especially the tail risks of server consolidation. Enterprises exponentially increase their probability of application catastrophe by performing consolidation without additional infrastructure improvements. Consolidating customers also face site and cutover risks. Leading server virtualization software was not designed to address these risks. Given the role of virtualization in the cloud model, cloud providers who use legacy disaster recovery architectures face all of these risks. The customers of those providers face the consequences.

In this chapter we address this concern. We first review server consolidation context to make concrete the virtualization coverage from Chapter 1, where we have reviewed the most cited benefits of consolidation. We then extend the traditional business case with a quantitative analysis of catastrophic risk. We describe our solution in Chapter 4.

Consolidation context

Server consolidation implements the virtualization concept that we examined in Chapter 1.

Virtualization support in chips and operating systems consolidates server functions and simplifies server provisioning. The consolidated infrastructure uses fewer physical servers running at higher utilization. The consolidated infrastructure is easier to manage and more responsive to changes in demand, both quantitative (surge) changes and qualitative changes in business requirements such as needs for new products and services. Analogous effects are also available from storage virtualization.

Figure 3-1 depicts a 6:1 server consolidation ratio. The application processing load of six physical servers is consolidated to one physical server of identical capacity. Prior utilization of 10% per physical server coalesces to 60% utilization on the target server. This is the conceptually simplest form of server consolidation.

Figure 3-1
6:1 server consolidation

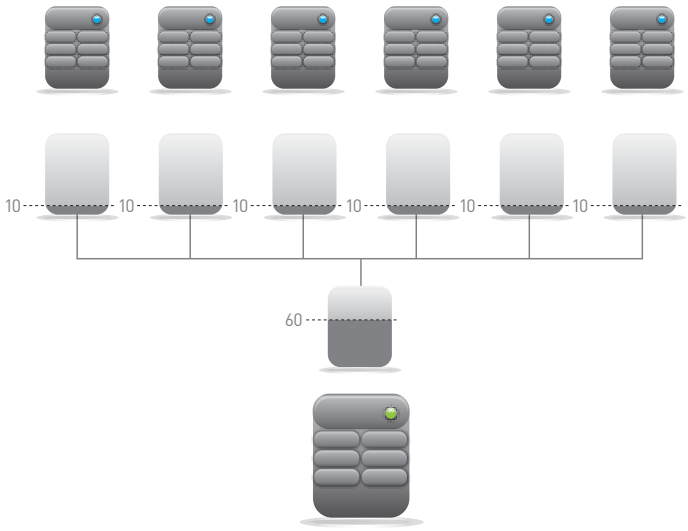

The more common variation of consolidation deploys newer, faster, more efficient multi-core processors in the target server to accommodate higher consolidation ratios. Intel's IT organization has reported an 8:1 consolidation with a 66% speed increase and 86% electricity decrease, suggesting direct annual operating expense savings over $6,000 per consolidation (ignoring space, networking and power backup).[1]

Intel documents cases of higher ratios:

- Mechanics Bank reports a 12:1 consolidation ratio and 5-year expense savings of $1.5–1.7 million.[2]
- AtlantiCare, a health services provider, reports a 16:1 ratio.[3]

Research suggests that simple server consolidation can save many enterprises 60% of expenses driven by server hardware. This analysis is incomplete because it ignores changes in risk, particularly tail risk. We summarized the simplistic business case under "How virtualization drives asset management benefits" in Chapter 1. We next analyze the tail risk of server consolidation.

1 "Server Consolidation Using Quad-Core Processors." Intel, 2006.

2 http://tinyurl.com/yqt62k

3 http://tinyurl.com/yva2fp

Risks of consolidation

ZERODOWN Software believes executives have received an incomplete picture of the risk-adjusted present value of virtualization, especially the tail risks of server consolidation. Enterprises exponentially increase their probability of application catastrophe by performing consolidation without additional infrastructure improvements. Consolidating customers also face site and cutover risks. Leading server virtualization software was not designed to address these risks. Given the role of virtualization in the cloud model, cloud providers that use legacy disaster recovery architectures face all of these risks. The customers of those providers face the consequences.

These risks are discussed below. Details follow. Readers who do not require details can, without loss of continuity or meaning, go to "Improvement strategies" on page 53.

(In the explanation that follows, preconsolidation is abbreviated as "precon," and postconsolidation is abbreviated as "postcon.")

Catastrophic risk

We define an application catastrophe as the concurrent absence of all mission-essential application service. Application catastrophes are worse than the sum of their parts. When customer- or supplier-facing applications are down and the email or incident management system is also down, the company has difficulty communicating with important external parties and its own ability to solve the problem decreases because needed internal coordination is more difficult. "Going dark" is dark indeed.

That is exactly what happened to Google's team in the February 2009 Gmail outage mentioned in Table 1-3. "Google itself depends on the service and press spokespeople for the firm were unable to e-mail journalists with statements regarding the problem."[4]

Catastrophic risk is driven by the eggs-in-one-basket problem. If all your application eggs are in one server basket that fails, you skip at least one meal.

The quantitative analysis of catastrophic risk rests on a joint probability analysis. Figure 3-2 on page 51 illustrates the increase in this risk.

The horizontal axis is the precon annual probability of N hours of downtime due to a server failure. This is the basic rate that assumes failures occur independently, that the failure of one server does not change the probability of another server failure. ZeroNine's analysis of downtime studies suggests that most firms' probabilities fall somewhere within the range of these baseline risks. Organizations who wish to identify their own position on the horizontal axis consult their downtime record, simulation, or Business Impact Assessment, and then select known probability percentage for a known N hours of downtime.

The three mildly curved lines that slope downward represent annual precon catastrophic risk of varying server quantities: 6, 10 and 14 servers. These represent probabilities of simultaneous server failure calculated from the basic failure rates. They overstate precon risk, so the chart understates postcon risk growth depicted in the orange lines. Organizations can plot their position by using the probability percentage and then selecting the downward sloping line that most closely corresponds to the actual or planned consolidation ratio.

4 "Google Users Struck By Gmail Outage For Over Two Hours" Rahul Chatterjee, eBrandz, 25 February 2009.

Figure 3-2
Exponential growth of catastrophic risk with consolidation[5]

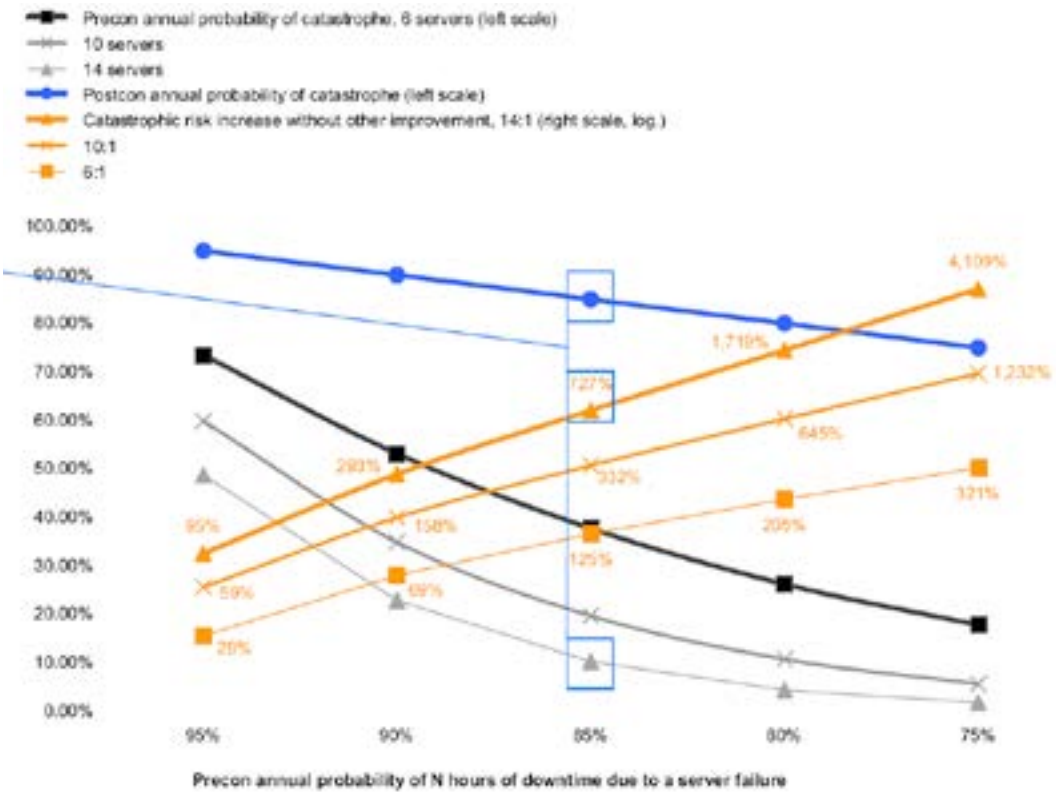

Each orange line depicts the increase in catastrophic risk from its corresponding dark-pigment line, precon catastrophic risk, to the blue line, postcon catastrophic risk. Actual (lower) precon risks imply higher growth rates in postcon catastrophic risk because postcon risk, blue line, is constant. Postcon risk grows exponentially with consolidation ratio and precon reliability: the higher the consolidation ratio or precon reliability, the faster the postcon risk growth.

The thick blue line represents the annual postcon catastrophic probability. This is the same for all cases because, absent other nonconsolidation improvements, the probability is the same as the horizontal axis value at each data point. It is the simple probability of failure of the unimproved postcon server environment.

The orange lines indicate the increase in postcon catastrophic risk at varying consolidation ratios and precon server failure probabilities, assuming no other infrastructure improvements. Catastrophic risk always increases. The best-case change, at the left end of the thinnest orange line, is a 29% increase. These risks worsen exponentially with consolidation ratio and precon reliability (the right-hand scale is logarithmic).

Site risk

We now turn to two other categories of risk: site and cutover risks.

Site risk is the risk from losing application service from a data center due to site-wide failure. These failures are almost always caused by a local physical disaster, such as a fire, flood, earthquake or an electrical surge that overwhelms protective circuits.

In the "should not happen" category, regional electrical outages driven by natural disasters, such as those caused by hurricanes or even grid management errors, can drive intermittent failures over days that collectively exhaust battery backup systems of telecommunications providers. Recall our Hurricane Charley example on page 37.

5 It is mere coincidence that "95%" is the same value at the left end of the orange 14:1 postcon risk growth line and on the horizontal axis below it. The orange line value represents the increase in catasrophic risk, a growth rate. The horizontal axis represents the annual basic server failure probability, assumed constant.

Note Every application service protected by the ZERODOWN Software Always Available architecture and technology has remained available to its application clients' network 100% since implementation. There has never been a case of an Always Available application client failing to reach its Always Available application service across an operational network.

Although site risk is not driven by server consolidation, the temptation to centralize everything is quite strong and is best resisted.

Leading server virtualization software was not designed to address site risk. These products do not perform transaction mirroring within a site, much less across the distances necessary to ensure business continuity in case of a disaster.

ZERODOWN Software recommends that organizations capture the risk management benefits of multi-site application support. ZERODOWN Software uses the "site diversity" concept to indicate a number of server sites that share no 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, Always Available application availability is feasible.

Example Sites in 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.

Cutover risk

Cutover risk is change risk, the risk that something will go wrong during a change that causes unplanned downtime.

The prevailing practice of server consolidation requires business to halt. Consolidation requires business application service to cease while technicians cut over to the postcon server. The most common euphemism for this is "quiescing" the application. Quiescing the application means stopping the business. By definition, it is a high-risk event.[6]

Quiescing an application and achieving cutover usually requires changes to the application and to network addressing. The execution time for these changes depends largely on the volume of work and human error rates. Aside from utilization of our Always Available architecture, we have never seen this time measured in seconds for a server consolidation.

Because of the high risk of the event and variable time required, cutovers are scheduled to occur when the business impact will be minimized, away from processing peaks at least. The pressure to "make it happen" on schedule drives personnel to extra effort in the weeks and days prior to the event. In practice, therefore, cutovers are attempted outside of normal working hours when personnel are sleep-deprived from a combination of long and shifting work hours. Sleep deprivation impairs reaction time and judgment similarly to alcohol intoxication.

The theoretical flaw of the cutover archetype is that it requires an event that halts the business. The disaster recovery architecture, which uses the synonym "failover," is based on the cutover archetype and suffers from similar risks that are amplified by disaster trauma.

6 Most IT organizations manage this risk as a "configuration management" or "change control" problem. The most spectacular career-crunching failures of the IT function that do not involve fraud reside in this category of operations. The other euphemism, most often used by business executives, is "migrating," and it also misleads. Migration is a smooth and natural process, not an event.

Improvement strategies

Organizations that consider server consolidation face the choice of higher catastrophic risk or implementation of additional improvements to the baseline case.

Improvement strategies to consider and the impact on each consolidation risk are summarized in Table 3-1.

Table 3-1
Improvement strategies to address consolidation risk

	Impact by category of risk		
Strategy to consider	**Catastrophic**	**Site**	**Cutover**
Hardware improvement	Reduces	None	None
Application distribution	Reduces	Reduces, if precon sites retained	None
Always Available™ architecture	Virtually eliminates	Virtually eliminates	Virtually eliminates

Target hardware improvement

Absent an architectural change such as we deliver in Always Available, reduction of catastrophic risk can be pursued by reduction of the failure rate of each component of the postcon system. ZERODOWN Software recommends that the best budget-feasible server hardware available be utilized
for consolidation. Thanks to improvements from Intel, AMD and other chipmakers, the latest processors really are the greatest as well: more reliable, effective and efficient. We see customers doing this build-down already, and we applaud it.

Note Customers of ZERODOWN Software do not need to discard their fully depreciated server hardware, and we suggest that it be considered for Always Available implementation. Our Always Available technology does not require matching server and network specifications. The existence of computing viruses as a corporate threat is a strong case for retaining diversity of vendors and operating systems. ZERODOWN Software Always Available is vendor- and platform-agnostic.

ZERODOWN Software recommends that customers seeking server hardware improvements consult their hardware vendors carefully about the achievements and limits of statistical process control. We would not be surprised to hear that SPC has driven server manufacturing to the statistical space where quality differences cannot be explained systematically with high confidence. Eventually randomness rules.

We suggest there is a natural limit to the additional reliability that management can accomplish within a centralized architecture. To the extent that brand-name vendors rely on fewer contract manufacturing firms, there are fewer opportunities for vendor differentiation on reliability. Buying better hardware can only take you so far, and you might be closer to that statistical wall than you recognize.

Required increase in server quality

How much better do servers need to be so postcon catastrophic risk is no worse than precon? The required increases in percentage match the increases in catastrophic risk from the base case (see Figure 3-3). Vendors increase server reliability in the normal course of their product design activities, but the required increases that we are discussing are not trivial. Higher consolidation ratios require stupendous server quality improvements. Vendors might achieve these increases over time, but probably not within one or a few depreciation cycles. Because the wall is tall, we think it imprudent to bet an improvement campaign solely on a server-improvement strategy.

Figure 3-3
Server quality increases required to equalize catastrophic risk

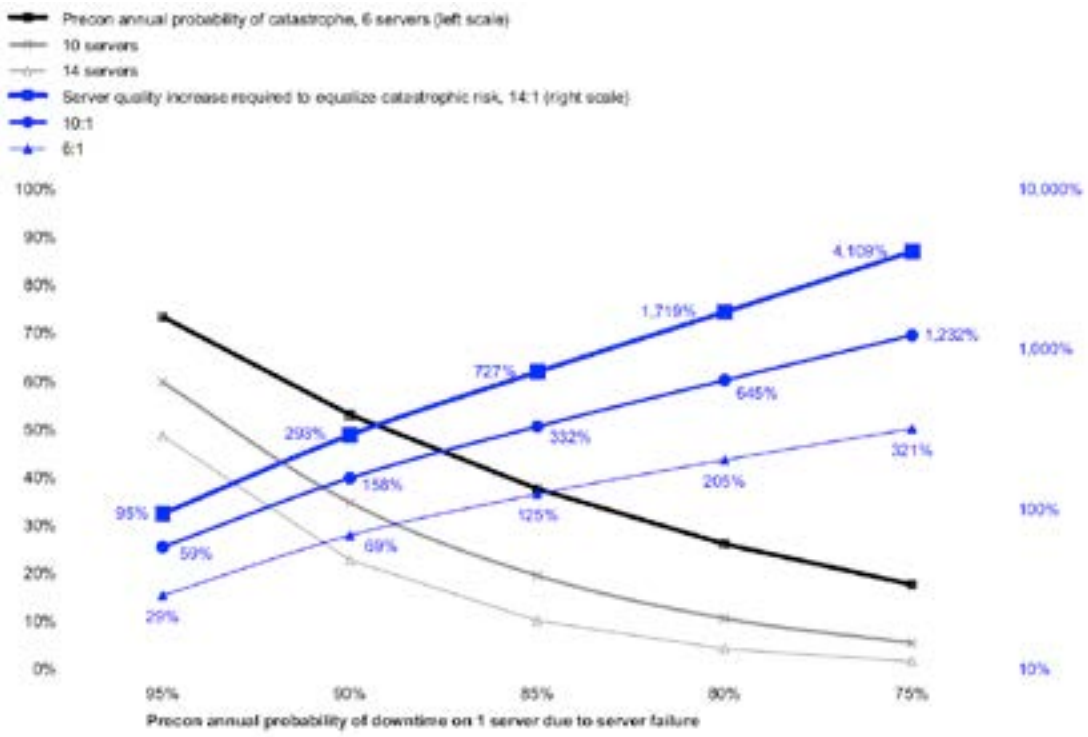

Application distribution

Organizations that perform consolidation are unwise to take the concept to its logical conclusion and centralize all applications on one server. Although this is possible and even seems desirable, the likelihood that prudent availability levels can credibly be sustained in this manner is remote.[7] ZERODOWN Software recommends that application instances be distributed across multiple physical servers at diverse sites, even if our Always Available architecture and technology are not in use.

The best consolidation is not the greatest consolidation. Purists argue that "the effort is about consolidation, so consolidate!" We believe that the effort is about securing the risk adjusted present value of shareholder wealth. Consolidation is best managed as a means to this end.

By operating at the transaction level, ZERODOWN Software Always Available technology supports virtually 100% application uptime from diverse hardware, operating systems and networks, unconstrained by distance and unaffected by latency. We are aware of no alternative that can do this.

7 We have seen exactly one analysis showing that a single-server approach produced a higher risk-adjusted value than a multi-server approach. The superior product was discontinued by its manufacturer because it was too expensive to build for the pricing model used at the time.

Always Available™ as a Solution

CHAPTER 4
Always Available™ as a Solution

SECTION 3
Solutions

INTRODUCTION

Always Available™ as a Solution

The risk-adjusted present value of cloud providers and cloud users credibly increases with ZERODOWN Software' Always Available business continuity architecture and technology.

- Cloud providers utilize our solution internally, driving benefits for themselves and their customers.

- Customers implementing CloudNines™ are those who require availabilty greater than assured by their cloud providers, need a gateway function to link disparate clouds, or wish to implement an Always Available infrastructure without commissioning additional data centers of their own.

ZERODOWN Software Always Available architecture and technology virtually eliminate application downtime of a cloud service. The CloudNines implementation of Always Available virtually eliminates downtime of application software that uses such a service. By linking servers and operating
systems in a platform- and vendor-agnostic manner—with our patented transaction-level technology that is fully compatible with leading virtualization products—we support the unprecedented combination of virtually 100% availability and operational flexibility for cloud providers and cloud users.

The always-on nature of our technology means you always know that your business continuity is always working, rather than waiting for a "recovery" test or actual disaster to learn that something does or does not work.

In this chapter we describe how Always Available enables these benefits. We do the following:

- Discuss the requirements of Always Available architecture

- Introduce topologies for cloud providers and customers and show before/after views

- Explain the key design principles and their beneficial side effects

- Convey a case study of true 100% email system uptime since July 15, 2004 utilizing our Always Available architecture.

Always Available™ requirements

Given the disaster of disaster recovery and the catastrophic risk of consolidation, what requirements must cloud continuity solutions address? We suggest the following:

- Mitigate regional disasters
- Virtually no loss of in-flight transactions
- Do not require speed- or capacity-matched hardware
- Leveraging current assets, even if fully depreciated
- Platform- and vendor-agnostic (hardware, operating system, network)
- No prolonged application customization
- Simple, elegant and cost-effective.

Our Always Available solution fulfills these requirements, as summarized in Figure 4-1.

Figure 4-1
Always Available™ solution summary

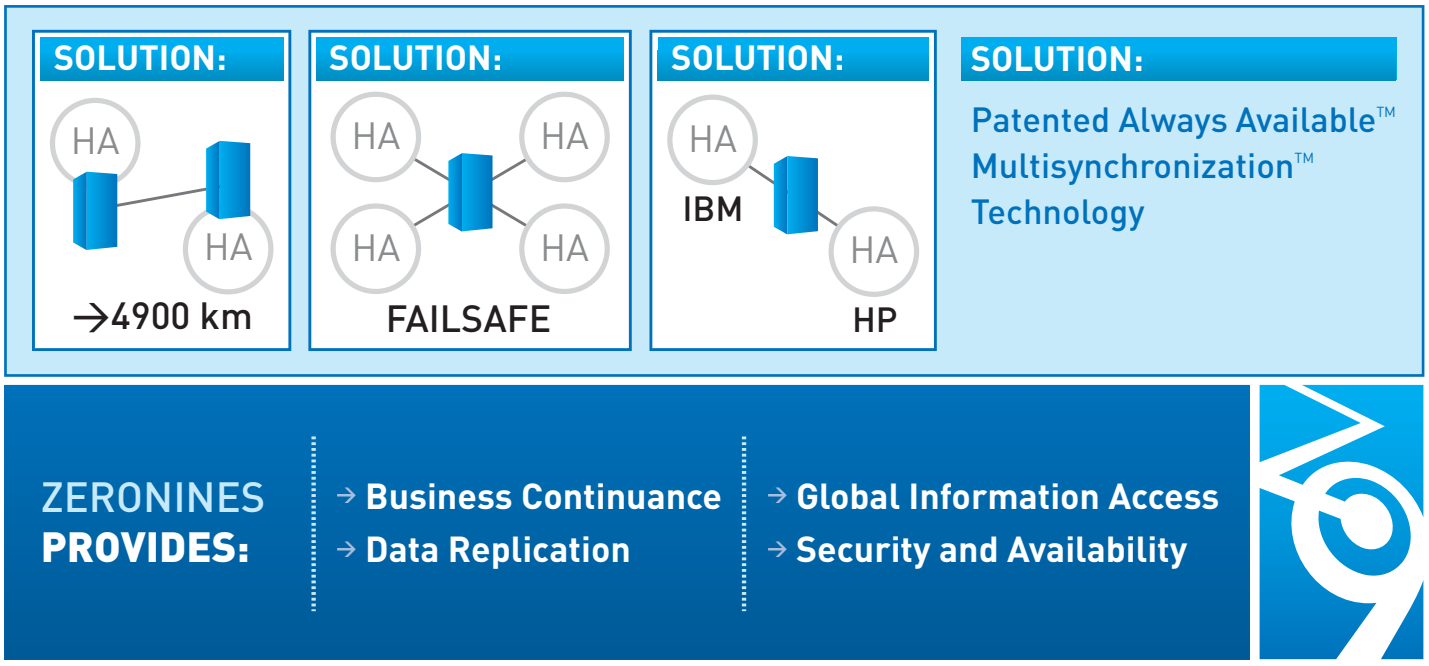

ZERODOWN Software' Always Available architecture disaster-proofs an application or cloud service without a wholesale application rewrite. The protected application or service communicates to the infrastructure through one or more Always Available protocol interfaces (adapters).

Examples include the following:[1]

- Cloud services
- Databases
- Transaction monitors
- Email systems
- Storage configurations
- Other business application software.

Availability from a ZERODOWN Software Always Available 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 4-2
Always Available™ solution topology (cloud provider)

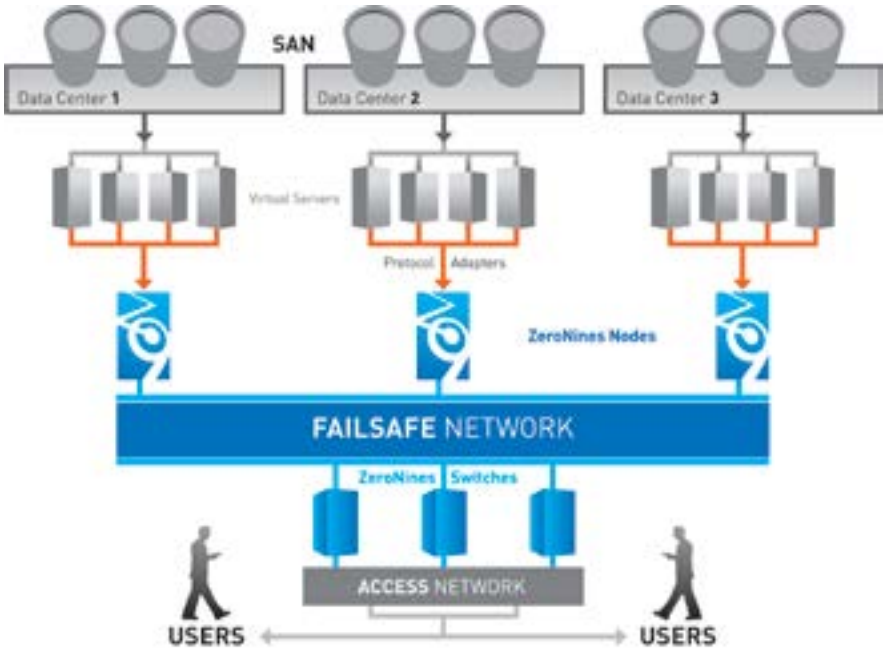

Cloud providers utilize our Enterprise solution. Customers implementing CloudNines are those who require availability greater than assured by their cloud providers, need a gateway function to link disparate clouds, or wish to implement an Always Available infrastructure without commissioning additional data centers of their own.

These alternatives are depicted in Figure 4-2 and Figure 4-3.

1 Through the remainder of this brief, we use "application" to include both discrete business applications and cloud services.

Figure 4-3
CloudNines™ topology
(cloud user)

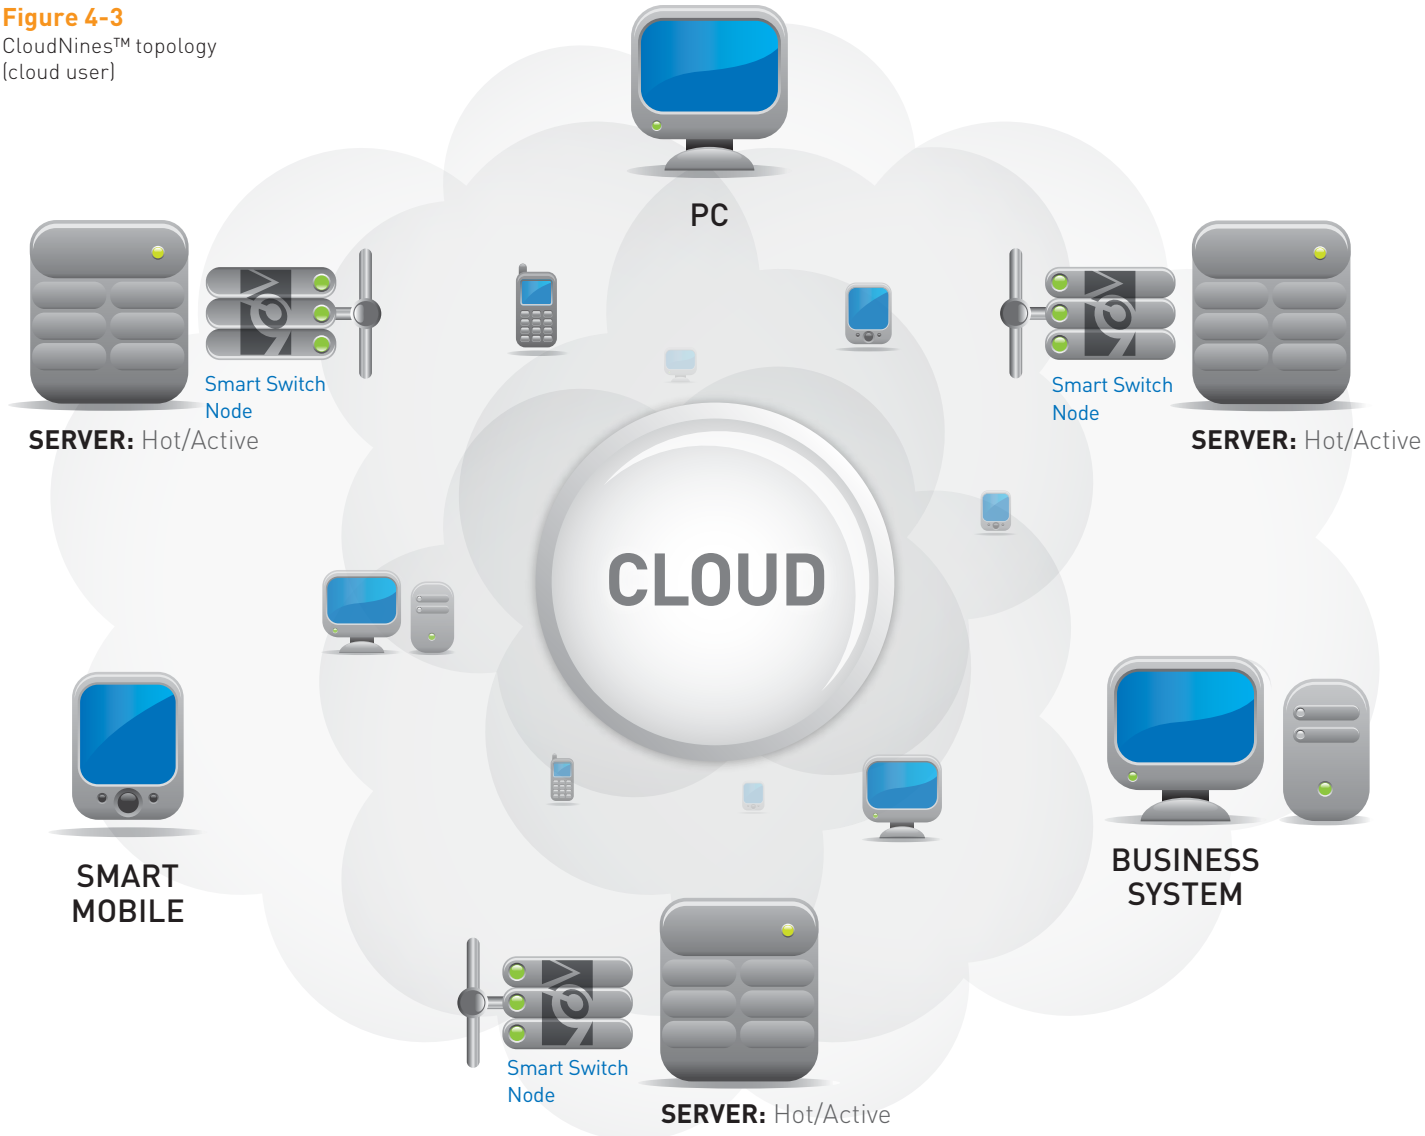

Design principles of an Always Available™ solution

The design principles of an Always Available 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**
- **Every transaction is journaled**
- **Load balancing is a side effect.**

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

An Always Available 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) or, for a CloudNines user, maintained with one or more clouds that themselves contain at least one Always Available node each. 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. ZERODOWN Software Always Available technology supports sessionless and session-oriented applications.

Server hierarchy is eliminated

Each application server image in an Always Available configuration is always logically primary. Server hierarchy does not exist in an Always Available 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 and, in a CloudNines configuration, those servers are in separate clouds. 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 or cloud might cease within the Always Available configuration for typical reasons such as scheduled maintenance or physical trauma, but the other sites or clouds in that configuration continue processing in a zero-loss manner that is transparent to the application.

Server sites are diverse

ZERODOWN Software uses the "site diversity" concept to indicate a number of server sites that share no 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, Always Available application availability is feasible.

Example Sites in 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—virtually 100% application uptime to client requests, even with unscheduled server maintenance.

With parallel reasoning, CloudNines™ users select diverse clouds. 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.

ZERODOWN Software Always Available 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 Always Available integrity. ZERODOWN Software' patented MultiSynch protocols prevent race conditions and operate asynchronously across thousands of miles. CloudNines™ scales these concepts to the level of disparate clouds.

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 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 quality of scope and economy of scale in a consolidated infrastructure. Other organizations have long ceased attempting such an approach in favor of accommodating top-down decisions driven by user requirements. Being application- and platform-agnostic, ZERODOWN Software' architecture does not constrain the choice of operating system, hardware, network protocols or cloud provider. Our technology enables heterogeneity as a design strategy for those who choose it without excluding those who do not.

Every transaction is journaled

An Always Available configuration journals every transaction. This core function enables Transaction MultiSynch. In addition to supporting application uptime of virtually 100%, journaling also enables new nodes to synchronize to a working configuration either as members of a new diverse site or as nodes returning after hardware or software refresh. Beyond this operational benefit, this all-transactions record is indispensible for audit, e-discovery and other business uses. Because our Always Available engine works below the application layer, a transacation record cannot be deleted merely through the application.

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 the application. Always Available servers, whether cloud based or internal, 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 Always Available capability.

An infrastructure before and after

To understand how a ZERODOWN Software Always Available configuration differs in a general sense from typical application access, consider the following exhibits. The principles in these exhibits pertain both to cloud providers, who implement this internally, and the cloud users who utilize our CloudNines™ solution.

Figure 4-4 depicts a typical application access topology, before Always Available. 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 4-4
Typical consolidated application access (not Always Available™)

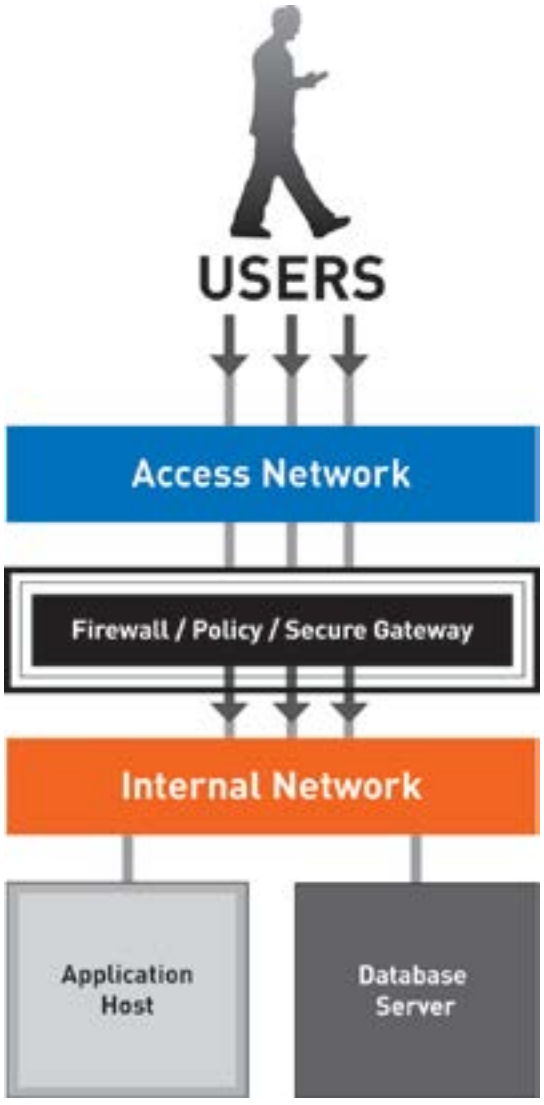

In a ZERODOWN Software Always Available topology as shown in Figure 4-5, two (or more) ZERODOWN Software Always Available switches are present between the application user network and the application servers' network. Each Always Available switch may have one or more state-accurate shadowing switches that continue service to the application clients in case a switch discontinues service for any reason, such as scheduled maintenance. Always Available Switches may be clustered for load balancing as desired.

Figure 4-5
Always Available™ consolidated application access

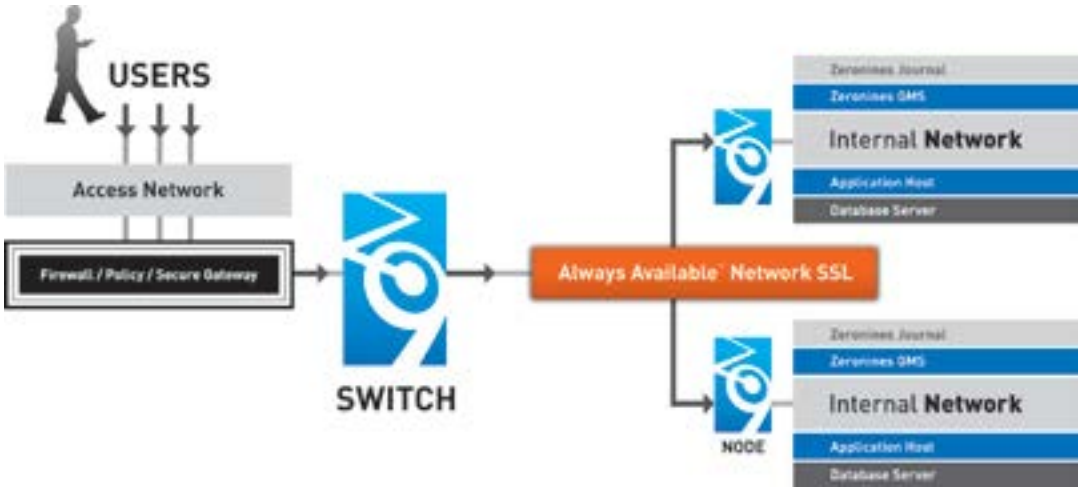

The mere fact that an Always Available configuration contains fewer single points of failure from a hardware perspective does not fully explain why continuous application availability is assured. Failover is insufficient for continuous availability. Simply buying more servers and configuring them for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

In a ZERODOWN Software Always Available configuration, each application server is associated with a ZERODOWN Software Always Available node, a listener function. When a client requests application service, at least two Always Available switches pass the request to at least two Always Available 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 ZERODOWN Software protocols 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 ZERODOWN Software 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

ZERODOWN Software 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 Always Available configuration, whether internal to your cloud or a CloudNines™ implementation, must reflect the imperatives of your Business Impact Analysis, business plan and regulatory requirements. ZERODOWN Software 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 an Always Available production configuration.

Table 4-1
Minimal site diversity for desired availability during prototype tests

Availability in prototype (%)	Residual probability	Site diversity required
99.999	1e10 $^{-5}$	2
99.99999	1e10 $^{-7}$	3
≈ 100	1e10 $^{-27}$	→ 3

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

Adding more sites than the minimum required will increase availability, providing greater protection during routine maintenance, upgrades, or additional trauma that causes simultaneous service interruptions at two or more sites.

See Figure 4-6 for perspective on the reduction in tail risk as site diversity increases.

Figure 4-6
Decrease in residual probability by site diversity (inverted scale)

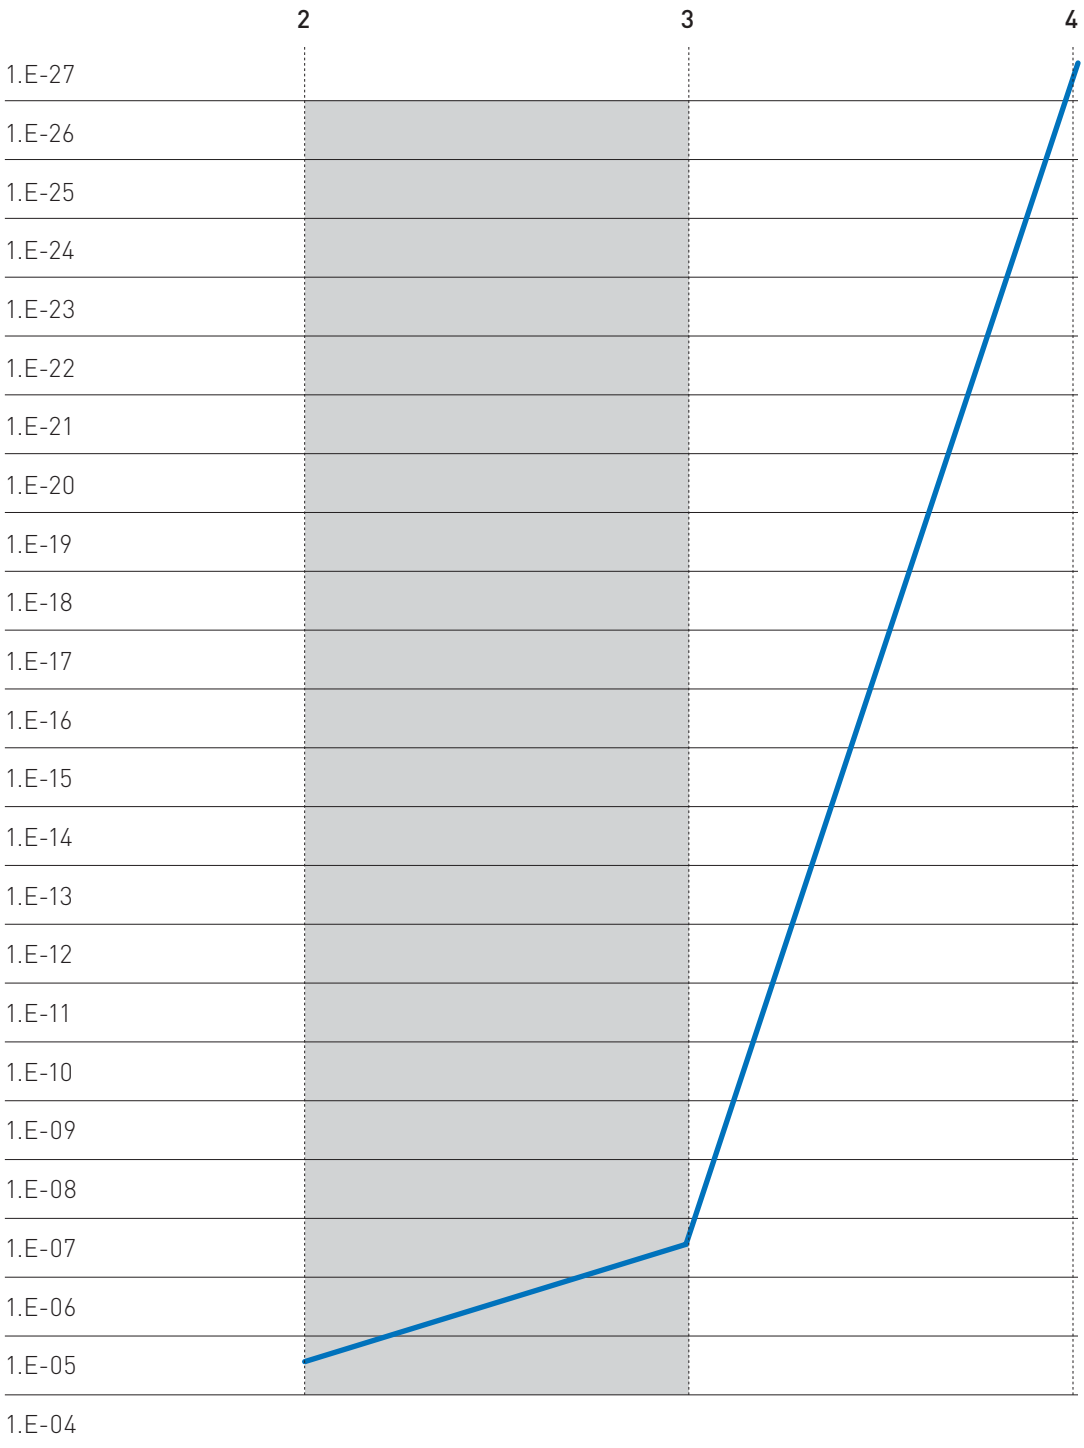

Case study: **MyFailSafe.com**™

Figure 4-7
MyFailSafe.com topology, on continuously since 2Q 2004

ZERODOWN Software 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 4-7).

- 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 in our original hosted implementation are described in Table 4-2 on page 69.

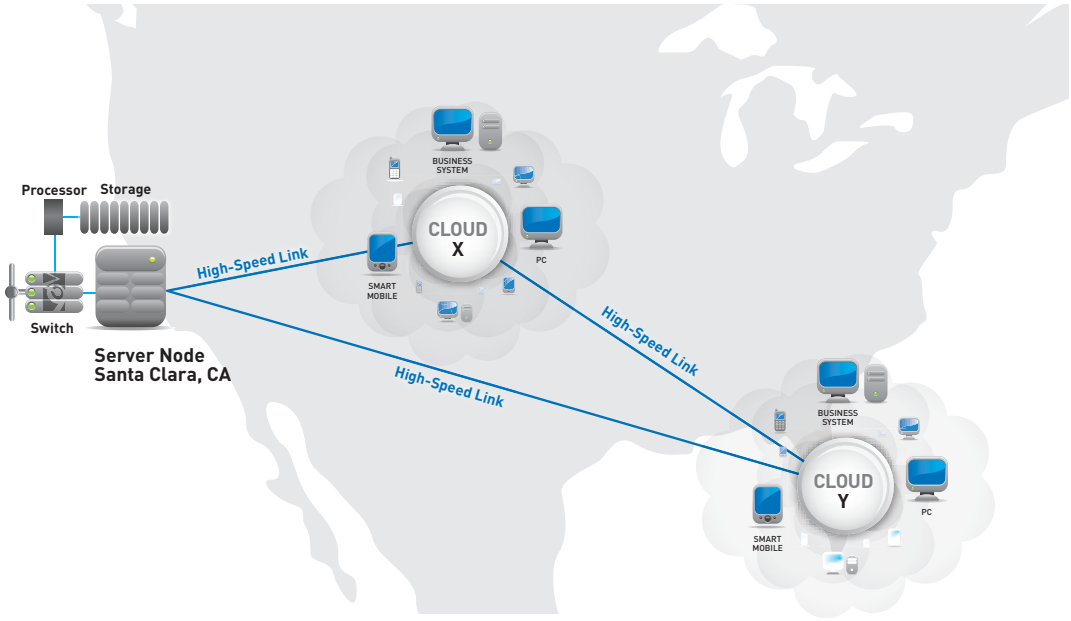

Since activation on July 15, 2004, MyFailSafe.com has furnished continuous service to email clients. There has never been an interruption of service to email clients for any cause: scheduled or unscheduled maintenance, server upgrades, virus attack, distributed denial of service attack, or natural disasters.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

Table 4-2
MyFailSafe.com telecommunication links

City	Link characteristics	Decommissioned for CloudNines implementation
Denver, Colorado	1MB, burstable	2009
Orlando, Florida	1MB–10MB	2007
Santa Clara, California	1MB, burstable	N/A

ZERODOWN Software supports cloud computing as an evolutionary technological change for enterprises with our CloudNines™ offer. Beginning in 2007, ZERODOWN Software began adopting CloudNines™ for the MyFailSafe.com service. For some of the nodes in the MyFailSafe network we have been switching from exclusive-access colocation to public cloud providers. The Orlando, Florida, node was the first to be moved into a public cloud. With CloudNines™, we are now running our own private cloud and utilizing public clouds for additional nodes. Continuous MyFailSafe.com service has been maintained during the transition.

ZERODOWN
SOFTWARE™

www.zerodownsoftware.com

For more information: **info@zerodownsoftware.com**

ZERODOWN SOFTWARE: **WHITEPAPER**

Corporate HQ

ZERODOWN™ Software
5445 DTC Parkway
Penthouse Four
Greenwood Village, CO 80111