CLOUD COMPUTING
BEST PRACTICES
CLOUDBESTPRACTICES.NET

Always-On Digital Government

Cloud Computing as a Catalyst for Legacy Transformation

By

Neil McEvoy
Founder and President, Cloud Best Practices Network

Alan Gin
Co-Founder & CEO, ZeroNines Technology, Inc.

DigitalGovernment.info

August 2012

Contents

Executive Overview

Two years ago the 'Cloud First' program ignited the trend of Cloud adoption in the U.S. Government, part of a 25-point plan for transforming its enterprise IT footprint through data-centre closures and legacy migrations to the Cloud.

This has indeed acted as a successful catalyst and considerable progress has already been reported. Additionally, momentum has continued to grow and evolve, with new programs building on this first step.

In particular, recent developments have included:

- Best Practices for Acquiring IT as a Service: A white paper from CIO.gov that stipulates a best practice framework for buying services from Cloud Providers, covering a range of topics such as contracting, records management, compliance, legal, security, and so forth, so that agencies can be suitably diligent in procuring their first Cloud projects.[1]
- Digital Government strategy: A program focused on stimulating economic growth through IT innovation, driven by the public sector. This was recently announced by the new Whitehouse CIO Steven VanRoekel.[2]

This paper identifies key recommendations from across these practices, with a special focus on the backbone activity of "legacy transformation" – the process of migrating existing applications and data-centres to new Cloud scenarios.

In the Whitehouse Cloud Best Practices guide they highlight:

> *"The US Federal Government spends approximately $80 billion dollars [sic] on Information Technology (IT) annually. However, a significant portion of this spending goes towards maintaining aging and duplicative infrastructure. Instead of highly efficient IT assets enabling agencies to deliver mission services, much of this spending is characterized by low asset utilization, long lead times to acquire new services, and fragmented demand. To compound this problem, Federal agencies are being asked to do more with less while maintaining a high level of service to the American public.*
>
> *Cloud computing presents the Federal Government with an opportunity to transform its IT portfolio by giving agencies the ability to purchase a broad range of IT services in a utility-based model. This allows agencies to refocus their efforts on IT operational expenditures and only pay for IT services consumed instead of buying IT with a focus on capacity. Procuring IT services in a cloud computing model can help the Federal Government to increase operational efficiencies, resource utilization, and innovation across its IT portfolio, delivering a higher return on our investments to the American taxpayer."*[3]

The aging nature of these systems presents an immediacy of need that the Cloud can address, but there are also broader and longer-term considerations, especially in how to achieve new online Open Government models.

Managing Government Records – Open Government Best Practices

The need for these improvements and technologies is documented in other headline initiatives.

In particular in the program announced last year 'Managing Government Records', President Obama describes how the IM (Information Management) function is critical to the successful implementation of Open

Government[4] – transparent and participative modes of public sector operation. He considers this a top priority area, so many private- and public-sector organizations are likely to put significant effort towards creating records management technology and best practices. Much of this will no doubt carry over to the Cloud, not only becoming new Cloud offerings in their own right but opening the door for development of related technologies.

This is repeated in the CIO.gov "Best Practices" guide, where from page 31 it focuses on these records management aspects and how to implement them in Cloud scenarios

> *"Many Federal agencies have older record schedules in place which fail to account for modern electronic records and may contain outdated references to superseded software platforms and applications. For these Federal agencies, a transition to cloud-based systems holds the potential to provide an agency`s records officer(s) with a chance to start fresh, identifying records and potentially updating schedules or creating them anew.'*[5]

Open Standards

Because these are universal issues that governments everywhere struggle with, equally common global standards will be key.

In this document and our ongoing Digital Government series, we will focus on key enabling technology trends and open standards from bodies like the Organization for the Advancement of Structured Information Standards (OASIS)[6], who pioneer a range of projects that are key to Cloud adoption, such as:

- **SAML, OpenID and OAuth –** Identity Management standards that provide the foundation for what NSTIC call the emerging 'Identity Ecosystem'[7].
- **Records Management –** The open standards needed to enable secure sharing and access between on-premise and Cloud-based systems for compliance needs, like FOIA (Freedom of Information Act) requests.
- **BPM and Integration –** Standards for "orchestration", the automation of buying and configuring virtual Cloud services, and also a platform that can enable better integration and workflow between legacy applications.

The Need for Legacy Modernization

The headline theme for this paper is a need for 'Legacy Modernization', a reference to upgrading older computer systems because their age presents a risk.

This situation isn`t limited to only the US Government. Indeed, it`s common to pretty much any and all large organizations that have been accumulating technology platforms since they started the earliest data processing centres.

For example the fundamental nature of this challenge is conveyed in the report Aging Information Technology Systems[8] by the Auditor General for the Government of Canada. In it, they introduce:

> *“Aging information technology (IT) systems refers not only to a system’s age in years but also to issues that affect its sustainability over the long term, such as the availability of software and hardware support and of people with the necessary knowledge and skills to service these systems. The term also relates to a system’s ability to adequately support changing business needs or emerging technologies, such as 24/7 online availability.”*[9]

The report then goes on to document their analysis of the current (as of 2010) state of the main IT systems used by the Government of Canada.

It points out that the five largest spenders of IT budget all have significant aging IT systems risks—the Canada Revenue Agency, Public Works and Government Services Canada, Human Resources and Skills Development Canada, the Royal Canadian Mounted Police, and Citizenship and Immigration Canada.

The analysis describes a number of significant levels of risk across this estate of applications, such as:

> *“The federal government relies heavily on IT systems to deliver programs and services to Canadians. Even though these systems are functioning, many of them consist of legacy applications that are supported by old infrastructure and <u>are at risk of breaking down</u>. A breakdown would have wide and severe consequences—at worst, <u>the government could no longer conduct its business and deliver services to Canadians</u>. Even applications that meet current business needs can be difficult and expensive to operate and may not be flexible enough to respond quickly to changes.” [emphasis added]*[10]

As examples of the risks in each individual area, they identify:

For PWGSC:

> *“PWGSC stated in its 2008 corporate risk profile that some outdated IT systems such as the Pay and Pension systems were close to imminent collapse, and compensation specialists were leaving as a result. The Department has initiated new projects to modernize both the Pay and Pension systems. We did not audit these systems.”*[11]

> *“Due to growing demand for departmental services because of the current economic downturn, and existing technologies that are reaching the end of their useful life, the Department recognizes that there is a high risk that its IT infrastructure will not be able to support the delivery of its core programs, such as Employment Insurance (EI). HRSDC also identified the lack of sustainable funding for renewing IT infrastructure as a significant corporate risk.”*[12]

Citizenship and Immigration Canada:

> *“The Field Operations Support System is a 29-year-old system critical to the National Immigration Program. It is considered high risk because the programming language is no longer being taught, and staff familiar with it are retiring. It is also very difficult, if not impossible, to integrate this application with newer systems.”*[13]

RCMP:

> *"One of those aging IT risks involves radio systems that use older technology unable to support current security and privacy requirements. According to the RCMP, this increases the risk to police and public safety and could lead to injury or death."*[14]

Digital Government: Framework for Cloud-Powered Modernization

Individually each of these departmental risks is concerning, but perhaps the most worrying aspect of the report is captured in this point, which sets the scene for the principle value of Cloud Computing:

> *"Although the Chief Information Officer Branch of the Treasury Board of Canada Secretariat is aware that the aging of IT systems is an issue, it has not formally identified it as an area of importance for the government. Nor has it assessed the issue from a government-wide perspective or worked with departments and agencies to develop government-wide solutions."*[15]

The core ideas of this paper are that these enterprise-wide strategies are obviously needed, and that a broad Cloud Computing strategy meets this need at that level. 'Moving to the Cloud' is inherently a process of Modernization, with a specific benefit being increasing resilience of these at-risk systems.

More so, a Cloud-enabled 'Digital Government' program fully encapsulates this within an overall framework for transformation. It incorporates other emerging technologies that streamline the online experience for citizens, as well as core infrastructure upgrades.

Data-Centre Modernization

Examples of the need for these infrastructure upgrades are conveyed through observations such as:

> *"For example, the heating ventilation and air conditioning system in the Montreal data centre is over 16 years old and the vendors no longer exist or make parts. As a result, the Department's Innovation, Information and Technology Branch (IITB) has spent $152,000 for repairs and maintenance contracts to maintain cooling capacity in the past year."*[16]

And for the Canada Revenue Agency a similar risk was identified:

> *"This data centre will not be able to support the Agency's long-term service needs because it is located in a 40-year-old complex that was not built to accommodate a data centre. Its age, location, and other factors pose a significant risk."*[17]

The core reason for adopting the Cloud is to migrate away from antiquated systems and move to modern equipment and data-centres to nullify such risks. A key point to note is that older systems carry much higher maintenance costs, and that across large enterprise estates there are likely many small and large instances like these.

Online Service Delivery

The overall compelling business need is that in today`s Internet economy users expect 24x7, multi-device customer service.

This always-on mode dictates that this infrastructure need is 24x7, and also that the services themselves need to be e-enabled and accessible via the web, smart phone, et al. This dual challenge is difficult for many agencies to meet.

The need for resilience is clear: These systems are often the only source of income for many citizens so delays and failures have real hard-hitting consequences. However this also does require that the systems are online and available 24x7.

There is often a direct correlation between the age of a system and the amount of downtime it experiences; the older the system, the greater the downtime. In some cases the people with the skills required to modify the platform are all long-retired and so it is simply a 'black box' – it works but no one knows how.

So moving to the Cloud isn`t just a "lift and shift" exercise where you are moving computing from one location to another, but it remains logically the same. It also dictates the need for business and technical transformation, such as increased adoption of social media and other tools to better improve communication with citizens and facilitate the flow of feedback and new ideas – i.e. all the inputs required to enhance and modernize services.

Lowering the Cost To Serve

The potential for Digital Government to achieve this is clearly conveyed through the SOCITM 'Better Served' report[18], which describes a 'Cost To Serve' ratio that explains how much each different CRM channel costs:

- Face to face : £7.40
- Telephone: £2.90
- Web: 32p

A Digital Government strategy is therefore about two key points:

- Transforming these expensive, paper-based processes that incur so much unnecessary cost
- Improving customer services as people now expect online delivery.

Using the Cloud to Prevent Data Disasters

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

The Need for Preventing Downtime

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

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

Using the Cloud to Create Expendable Data-Centres^{SM}

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

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

Benefits of Always Available technology include:

- Improves reliability of applications and data stored in clouds and traditional data-centres, enabling application and data uptime in excess of 99.999% (five nines).
- Protects against monetary losses caused by outages.
- Reduces customer attrition and brand damage caused by service disruptions.
- Increases the power of your cloud by synchronizing multiple private, public or hybrid clouds at multiple vendors.
- Facilitates migration to the cloud.

- Helps meet government regulations and industry standards.
- Extends the useful life of current technology assets.

An Example from Calgary

As we saw in Calgary on July 11, 2012, a single disastrous event can lead to widespread loss of business and municipal services. On that day, a transformer explosion knocked out data services at medical facilities, took an IBM data-centre offline, and crippled some city government systems for two days or more including the 911 emergency system[19].

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

It's All about Disaster *Prevention*, not Recovery

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

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

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

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

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

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

Won't Failover and Backup Prevent a Data Disaster?

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

Since 1989, we have asked countless Fortune 500 and Global 2000 clients what their level of confidence in their DR plans and recovery strategies is. Without exception, the response is “None, but this is all we have.”

Yet we continue to see these household name companies invest billions in failover- and backup-based DR strategies knowing they are unlikely to work and that the outcome will probably be disastrous.

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

Failover

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

Tape or Optical Backup

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

Other DR systems are basically new variations on failover and backup. Although some of these modernized techniques can significantly reduce the actual outage to a few seconds, many of the risks of data corruption, failed failover, and failed restoration still exist. Again, just look at the news for examples. The Amazon EC2 Cloud went down twice in June 2012, and we can assume they used the most effective DR methods known to them[20].

Cloud BCP - DR in the Cloud

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

Cloud outages occur with frightening regularity so customers need good business continuity planning (BCP) and they need to know exactly what kind of DR support to expect. The cloud provider may offer a variation on the failover paradigm to try to maintain continuity. They may offer multiple availability zones as Amazon does[21]. They may urge their customers to provide their own DR system. And as some companies have found out recently, it is a mistake to assume that your cloud provider will make backups for you:

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

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

This was said in June 2012 by Cameron Peron, VP Marketing at Newvem, a cloud optimization consultancy that specializes in the Amazon cloud[22].

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

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

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

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

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

Here's What Happens

1. Server calls are originated by networked clients: computers, email, PDAs, phones, business software, databases, etc.

2. Each transaction passes through the ZeroNines Smart Switch Nodes, which route it simultaneously to all clouds and other data-centres in a one-to-many (1:m) session.

3. Processing and storage take place equally in all locations. All clouds and servers are hot, and all are active. There are no primary or secondary clouds or servers.

4. All transactions and data

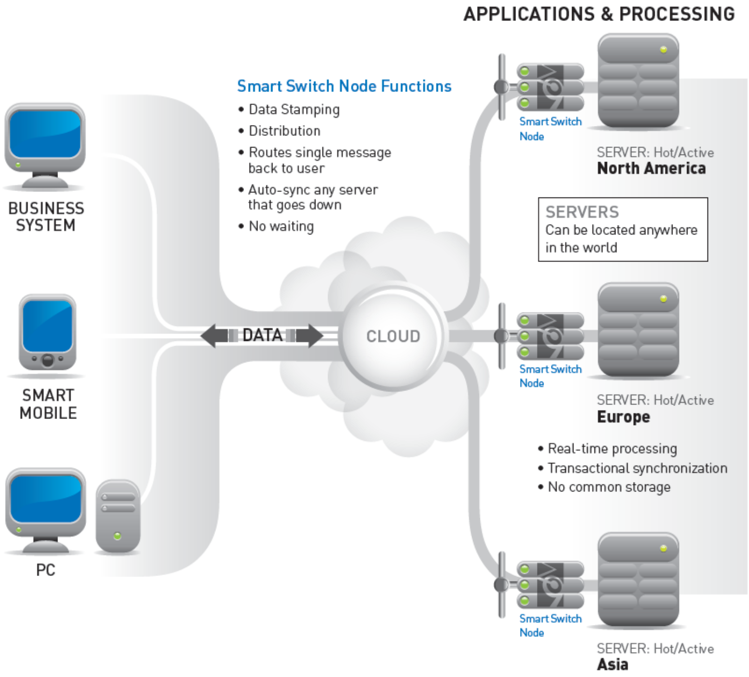

exchanges are recorded, verified, and subjected to security measures.

5. Each location sends its responses back through the Smart Switch Nodes.

6. The Smart Switch Nodes cooperatively eliminate duplicate responses and return a single response to the client that originated it.

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

Interoperability

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

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

Facilitating Cloud Migration

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

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

Case Study: Web Portal Startup ZenVault

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

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

Savings and Security Justify the Effort

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

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

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

We believe that the leading best practice for continuity in the cloud will be to dispense entirely with failover. Instead, government IT departments, government contractors, and businesses will adopt a technique that uses the cloud itself and its affordable computing capacity to virtually eliminate outages. The price of downtime events is highly unpredictable and extremely high, starting at millions of dollars per hour for even a moderately sized business. Governmental agencies are concerned about budgetary constraints, lost lives, and the appearance of stability. Exchanging these extreme risks for the known costs of cloud fees and disaster prevention makes a highly attractive cost-benefit equation. Many businesses are already discarding old and inefficient hosting models, buying extra cloud capacity, and reinvesting the savings into reliability. Their strategy is to compete based on reliability. Government entities, though not concerned about competition, will likewise be eager to eliminate the unknown but undoubtedly bitter cost of future outages.

About the Authors

Neil McEvoy

Founder and President - Cloud Best Practices Network

Neil McEvoy is a Cloud Computing entrepreneur who has been pioneering new innovations in this industry for twenty years. At age 28 he launched his first company, one of Europe's first ASPs (Application Service Providers), a joint venture with Microsoft to bring hosted CMS systems to small businesses and funded by the elite of the UK Internet entrepreneur market. Since then Neil has repeatedly brought new Cloud products and managed services to market across a spectrum of different industries and product segments, both in Europe and now more in North America. Most recently Neil has founded and launched the Cloud Best Practices Network in Toronto, with plans to expand throughout the nation, the USA, Europe and Asia.

http://CloudBestPractices.net

contact: neil.mcevoy@l5consulting.net

Alan Gin

Co-Founder and CEO, ZeroNines Technology, Inc.

Alan Gin founded ZeroNines Inc., a predecessor of ZeroNines Technology Inc., in May 2000 to implement and deploy the Internet operating system of the new millennium. Mr. Gin led a virtual team as the Chief Architect, developing the ZeroNines® Architecture. Prior to founding ZeroNines, Mr. Gin held numerous executive and management positions at Hitachi Data Systems, Storage Tek, AT&T Global Information Solutions, Wang Laboratories and Coopers & Lybrand. Mr. Gin attended Hawaii Pacific College from 1974 to 1978 in pursuit of a Bachelor of Science in Business Administration. He has lectured, co-authored manuals and reports and conducted classes on designing and auditing networks, client-server strategies and implementing financial systems. He is listed in Who's Who in the Computer Industry for his experience in designing and implementing mission-critical financial applications on controlled global area networks.

http://www.zeronines.com

contact: alan@zeronines.com

About the Cloud Best Practices Network

The Cloud Best Practices Network is an industry forum owned and operated by L5 Consulting in Toronto, Canada. The objective of the forum is to build a knowledge base of Cloud Best Practices that underpin a consulting and solutions program that assists customers to migrate successfully and profitably to Cloud Computing.

http://CloudBestPractices.net

Notes

[1] "Creating Effective Cloud Computing Contracts for The Federal Government: Best Practices for Acquiring IT as a Service" A joint publication of the Chief Acquisition Officers Council in coordination with the Federal Cloud Compliance Council, February 24, 2012, http://www.cio.gov/cloudbestpractices.pdf

2 VanRoekel, Steven, Federal Chief Information Officer, "Hitting the Ground Running with the Digital Strategy", http://www.whitehouse.gov/blog/2012/06/21/hitting-ground-running-digital-strategy June 21, 2012

[3] "Creating Effective Cloud Computing Contracts…" p. 1

[4] Obama, Barak, President, "Presidential Memorandum -- Managing Government Records" http://www.whitehouse.gov/the-press-office/2011/11/28/presidential-memorandum-managing-government-records, November 28, 2011. Section 1, Paragraph 2

[5] "Creating Effective Cloud Computing Contracts…" p. 31

[6] Website for the Organization for the Advancement of Structured Information Standards, as accessed August 9, 2012, https://www.oasis-open.org

[7] Website for the National Strategy for Trusted Identities in Cyberspace (NSTIC), as accessed August 9, 2102, http://www.nist.gov/nstic/identity-ecosystem.html

[8] Report of the Auditor General of Canada to the House of Commons, Spring 2010, Chapter 1, "Aging Information Technology Systems" http://www.oag-bvg.gc.ca/internet/docs/parl_oag_201004_01_e.pdf

[9] "Aging Information Technology Systems", p. 1

[10] "Aging Information Technology Systems", p. 2

[11] "Aging Information Technology Systems", p. 10

[12] "Aging Information Technology Systems", p. 11

[13] "Aging Information Technology Systems", p. 12

[14] "Aging Information Technology Systems", p. 13

[15] "Aging Information Technology Systems", p. 2

[16] "Aging Information Technology Systems", p. 11

[17] "Aging Information Technology Systems", p. 9

[18] Sargent, Vicky, "Adopting A Digital Approach To Engaging Citizens In Local Councils", The Information Daily (formerly eGovmonitor), http://www.egovmonitor.com/node/40858, February 23, 2011

[19] Yevgeniy Sverdlik, "Transformer explosion knocks out hospital, IBM data centers in Calgary" DatacenterDynamics, http://www.datacenterdynamics.com/focus/archive/2012/07/transformer-explosion-knocks-out-hospital-ibm-data-centers-calgary, July 13, 2012

[20] Rich Miller, "More Problems for Amazon EC2 Cloud" Data Center Knowledge, http://www.datacenterknowledge.com/archives/2012/06/29/another-outage-amazon-cloud/, June 29, 2012

[21] Cade Metz, "Amazon Floats New Cloud Over Pacific Northwest" Wired.com, http://www.wired.com/wiredenterprise/2011/11/amazon-in-orego/, November 9, 2011

[22] John Koetsier, "Using cloud services? 40% of you aren't ready for the next outage" VentureBeat, http://venturebeat.com/2012/06/21/using-cloud-services-40-of-you-arent-ready-for-the-next-outage/, June 21, 2012