



I D C A N A L Y S T C O N N E C T I O N



David Senf
Research Director, Infrastructure Solutions

How Does Cloud Fit Into Your Business Continuity Planning?

Sponsored by Allstream and Geminare

November 2011

For better or worse, a lot more people today are familiar with service level agreements (SLAs) than ever before. With evening and weekend work common among many employees, any planned or unplanned IT interruptions get noticed. Cloud adoption grows quickly in Canada to meet demands from business and IT departments under pressure to improve service with the same resources. On the topic of cloud and continuity, IDC analyst David Senf addresses key questions for Canadian organizations.

Q. Why is disruption and downtime garnering more attention today?

- A. Business dependence on both information and technology is higher today than in the past. Improvements in productivity, agility and revenue are inseparable from technologies, whether at our fingertips or in the server room. Most processes within a company rely on IT availability. Downtime of any length is intolerable for employees, partners and customers. It is not so surprising then that IT disruption is the top threat that businesses are concerned about. In fact, IDC research shows that failure of IT systems trumps all other business continuity threats, exceeding "acts of God" and power outages.

There is an unfortunate irony, however, that more business reliance on IT leads to more technology deployment — and in turn more points of potential failure. The likelihood of downtime further increases as the number of devices in the IT environment expands — and there has been no shortage of growth lately. Over the past couple of quarters in Canada, IDC saw an 11% increase in server shipments, a 54% increase in storage capacity, and a 13% jump in virtualization. The number of network switches arriving in Canadian IT shops grew by over 30% last year alone. These large growth numbers highlight the challenges IT faces in keeping up with the pace of change and avoiding downtime.

There are two fundamental aspects to downtime, and both are strongly related to change. One is mean time between failures (MTBF). This boils down to how often a system goes down. The other key metric is mean time to repair (MTTR), which tracks how long it takes to then get a system back up and running. In both cases, IDC finds organizations that have IT management processes (e.g., ITIL) in place score much better. However, in many firms, rolling out processes is easier said than done. Because most of IT's time and resources are typically dedicated to maintenance, there isn't a lot of time available to train and certify staff, let alone to develop and implement new processes, such as business continuity or change management.



Process is only one variable in the equation for reducing downtime; people and technology factor prominently too. And all three legs of this well-known stool require planning. IDC sees good news on the horizon for IT system uptime, given the focus that firms are placing on planning. 44% of large businesses have a business continuity plan in place, while an additional 33% have departmental/business unit plans in place. Further, nearly 20% are launching plans in the next 12 months. There is much less planning in place in the SMB market segment. However, nearly one-quarter of small business are launching plans in the next 12 months.

Q. What does the future hold for cloud and recovery?

A. A shift in IT thinking builds momentum for replication and recovery in Canadian firms, and virtualization plus cloud is behind this shift. The breadth and scope of IT changes, mentioned above, equates to more potential points of failure, and thus inspires action. But the low cost of redundancy and rapid provisioning of virtualization has reshaped how IT views recovery. Essentially, virtualization makes recovery very easy, and draws a picture that clearly outlines the value of elasticity, self-service, and shared services, which by no mistake are the core tenets of cloud. Further, virtualization redefines location limitations. Moving — or recovering — an application to and from a cloud is going to become commonplace. Referred to as hybrid cloud, both public cloud and on-premises systems provide a balance for a range of use cases — a key one being replication for planned and unplanned downtime.

Public cloud activity is ramping up fast in Canada, with a growth rate of more than five times that of the traditional IT market. Public cloud spend is close to \$500 million today in Canada and IDC forecasts it to approach \$1.2 billion by 2015. Our surveys show 41% of organizations use SaaS (e.g., CRM, ERP, Email). IaaS (e.g., storage and server) usage comes in at 32% of firms, compared to 11% for PaaS (e.g., development platform).

Demand is strong for cloud services such as backup, archiving, replication, and recovery — services geared towards minimizing downtime. Close to 40% of Canadian firms are likely or highly likely to subscribe to cloud-based recovery services; an additional 30% are considering but undecided. Overall, storage services hold an average of 11% of Canadian organizations' data, according to IDC cloud survey research. These organizations further believe that in five years their cloud storage will climb to 24% of their total data.

Q. Why is the network a critical piece of this puzzle?

A. The new number-one buying criterion for networking equipment is "uptime." Price and ease of integration round out the top three, but both have been bumped lower by the concentration on uptime. Whether it is communication across the WAN to a branch office or to a cloud service, network availability is critical.

Recovery time objective (RTO) is an important concept for setting the priorities that dictate which applications need to come up first when a small disruption or a large disaster happens. RTO requirements specify how quickly users need to have access to the data and the functionality of everything from email to finance to sales applications. Without the network, information doesn't move between technologies — there is no IT. Performance and uptime of business applications are often determined by latency, loss, and other factors at the network layer. Meeting recovery time objectives is determined by a range of factors, many of which start and end with the network.

Recovery is about more than bringing an HR application online. Traffic types at the network are changing. Virtualization, video and voice occupy more capacity. If a remote office system goes down, recovery can take the form of spinning up entire workloads across the WAN from

the virtual machine on up. Activity that used to take place locally on the desktop is now transacted over the network as well. In fact, the entire desktop experience is shifting to the datacentre and to the cloud. From boot storms to general QoS constraints, end users complain about any latency, not to mention outright system failure.

New bottlenecks form while others are reduced. WAN optimization, deduplication, and other optimization techniques can mitigate some performance issues, but the fact remains that a lot more mobile computing, virtual server, video, voice and storage activity is running along corporate and public pipes.

Public cloud infrastructures create new considerations for network administrators and architects, from availability to performance to security. IDC finds many organizations don't cast their network planning far enough into the future. The strain on QoS builds as activity on the network shifts towards increased cloud.

Q. With cloud in the picture, what steps should IT consider to reduce disruption and downtime?

A. Every firm suffers from planned and unplanned downtime — the goal is to mitigate user impact. In an ideal world, IT will find the time to work through the best practices that IDC outlines below. Time constraints being what they are, we recommend testing the applicability of cloud replication services as an option, particularly to avoid capital investment.

- **Plan.** IT departments that are more advanced in their use of change, configuration, availability, capacity and a range of other management processes have less downtime and score better on MTBF and MTTR. Planning should involve all three legs of the stool: process enhancement, people (e.g., training) and technology (e.g., tool investment).
- **Architect for future trends.** Security, compliance, peak versus average load/seasonality, acceptable QoS, cost, staff expertise, and sudden business change (e.g., M&A) all determine which type of deployment model is used where and when. Cloud shouldn't be merely a sourcing option, but an extension to IT architecture. Start with hybrid cloud as the goal.
- **Know your applications and data.** With a holistic view of IT in the context of business, consider the priority of applications based on factors such as recovery time objective (RTO). The top three applications that Canadian organizations want to recover first in the event of a disruption are: email, finance, and Web applications. Ensure complete QoS coverage (e.g., P, V, LAN, WAN).
- **Have end-to-end visibility/monitoring.** Each application has requirements, both while running live and in disaster recovery scenarios. Set an acceptable baseline for latency, loss, and MTBF to then monitor for deviation over time. Consider, too, incident prevention rather than incident response.

ABOUT THIS ANALYST

David Senf is IDC Canada's Research Director, Infrastructure Solutions. His team tracks security, networking equipment, servers, storage, virtualization, management tools and cloud. The ecosystem that surrounds these technologies is also examined from ITIL to compliance to business continuity. He works with vendors, the channel, and end-user customers to form a complete picture of market direction.

In the decade prior to joining IDC, he sold, managed and implemented IT consulting initiatives ranging from Web portals to online presence. His accumulated understanding of technology trends from a business strategy and IT "nuts and bolts" perspective enriches his research with both strategic and tactical advice. He has taught courses he developed on SOA and XML at the University of Toronto.

In demand as a speaker on a range of emerging technologies, he is also a frequent commentator and writer for both IT trade and business media.

ABOUT THIS PUBLICATION

This publication was produced by IDC Go-to-Market Services. The opinion, analysis, and research results presented herein are drawn from more detailed research and analysis independently conducted and published by IDC, unless specific vendor sponsorship is noted. IDC Go-to-Market Services makes IDC content available in a wide range of formats for distribution by various companies. A license to distribute IDC content does not imply endorsement of or opinion about the licensee.

COPYRIGHT AND RESTRICTIONS

Any IDC information or reference to IDC that is to be used in advertising, press releases, or promotional materials requires prior written approval from IDC. For permission requests contact the GMS information line at 508-988-7610 or gms@idc.com. Translation and/or localization of this document requires an additional license from IDC.

For more information on IDC visit www.idc.com. For more information on IDC GMS visit www.idc.com/gms.

Global Headquarters: 5 Speen Street Framingham, MA 01701 USA P.508.872.8200 F.508.935.4015 www.idc.com