



Cloud-based continuity of operations

New challenges and opportunities in enterprise-class continuity

Table of contents

Elections and natural disasters: a perfect storm	2
Challenges to continuity of operations.....	2
Key considerations	2
How it works	3
Benefits of cloud continuity	4
Conclusion.....	4
Authors	5

Elections and natural disasters: a perfect storm

Disasters can have a dreadful, direct impact on citizens and the operations of state and local governments. In the last 10 years, elections officials have faced natural disasters that have impacted elections administration, most notably:

- August 2005 - Hurricane Katrina's impact along the Louisiana and Mississippi Gulf Coast delayed county deployments of HAVA-compliant solutions
- October 2012 - Hurricane Sandy's impact to New Jersey, New York, and Connecticut resulted in polling place closures for the November Presidential Election

In today's "instant-on" world, continuity for state elections officials is gaining more attention from the media and public alike. In this business white paper, we present a more efficient approach to re-establishing IT operations by leveraging cloud and virtualization technologies.

Challenges to continuity of operations

Enterprises and governments are challenged by operational risks to their often complex and mission-critical IT environments. While innovations like virtualization and cloud can bring major value to enterprises, these new architectural models can also increase operational risk if proper controls are not put in place. Other long-term challenges include the increased reliance on IT and the growing threat of natural and man-made disasters. Organizations must protect their critical systems and data, while maintaining the agility to quickly and easily support changes in their IT environment, people, and organizational processes.

As governments and IT become more tightly integrated than ever, forward-looking organizations understand the importance of maintaining IT-based business processes. As IT-dependent processes become more critical to the enterprise, downtime and data loss are less and less acceptable. Traditional recovery times—typically 24 to 72 hours from a most-recent tape backup—are inadequate for most crucial enterprise applications.

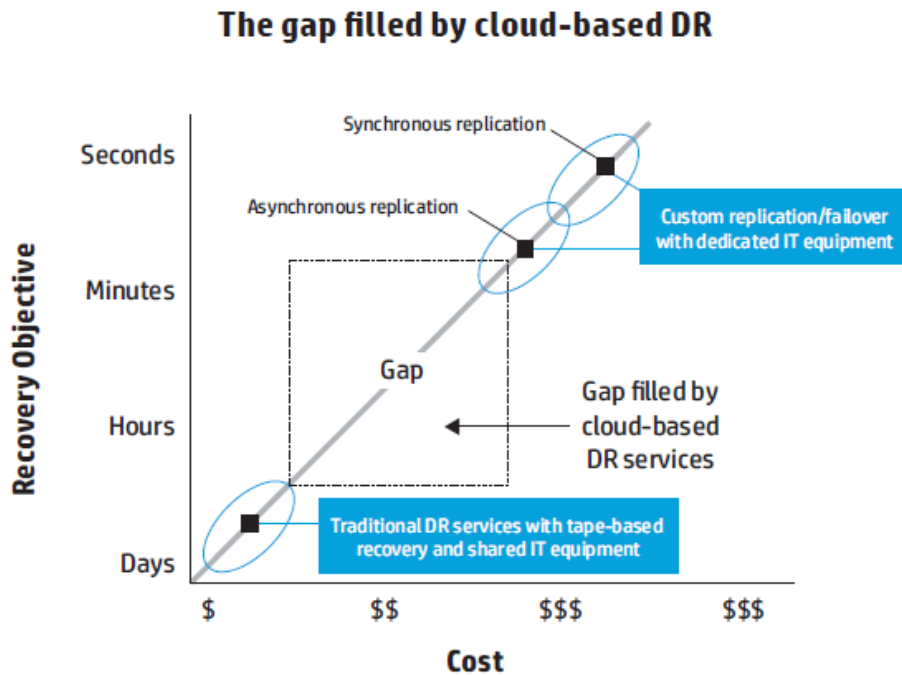
Today, falling costs and increased reliance on IT have yielded cost effective recovery times of less than four hours, with just minutes of potential data loss, for many more key applications. Both lack of flexibility and cost concerns have convinced many organizations to take their disaster recovery (DR) or continuity of operations (COOP) solutions in-house. Yet in-house COOP solutions may present their own unique set of challenges. A number of organizations have evaluated solutions that leverage test and development resources as a failover recovery environment, on the assumption that those less-critical assets were effectively cost-free. In reality, many organizations have learned that locating vital continuity assets in the test and development environment carries less tangible but very real costs.

The good news: a new and robust disaster-recovery-as-a-service (DRaaS) model has now emerged that promises to deliver enterprise-class continuity, cost-efficiency, and flexibility. Now considered by many observers to be one of the hottest new areas for the enterprise cloud, DRaaS has captured the interest of many organizations.

Key considerations

When assessing available continuity models, organizations may wish to consider the limitations of traditional tape-based systems, and the true cost and reliability of internal solutions that rely on test and development resources. While the drawbacks of those traditional models are well understood, cloud-based continuity provides a robust solution with security measures tuned to the needs of medium and large enterprises.

Traditionally, there has been a major cost/recovery objective gap between tape-based recovery solutions based on shared IT equipment and the high-end, dedicated replication/failover solutions. While many organizations wanted to recover more quickly than the one- to three-day times offered with traditional shared IT/tape-based solutions, they couldn't justify the 10 times greater cost of high-end, custom asynchronous, and synchronous replication solutions. This big gap between solutions has existed for many years. Figure 1 illustrates this important cloud recovery gap.

Figure 1. Cloud-based DR filling an important gap in enterprise business recovery.

The alternative to these high-end custom replication/failover solutions for many organizations has been either a dedicated or repurposed in-house recovery environment. Unfortunately, these solutions can be costly, cumbersome, and less than flexible. Many methods require substantial investments in a full set of hardware to mirror mission-critical systems, including significant data storage capabilities—and that mirrored data center must be co-located or situated in another remote environment.

Ongoing maintenance fees add up, as do the cost of routinely replicating production data to the mirror systems and the expense of keeping application software synchronized and updated across both the production and DR locations. Fortunately, organizations can now deploy a dynamic cloud-based continuity solution to protect against mission-threatening disruptions or data loss.

As envisioned, the DRaaS approach allows the resumption of normal organizational processes in one to four hours following a major interruption and replicates data to typically ensure less than 15 minutes of data loss.

The cloud-oriented replicated environment can also be used for non-disaster purposes, such as for user acceptance testing, quality assurance, and to proactively avoid certain other causes of planned and unplanned downtime. All of this can be done at an aggressive price that is roughly the same as the 24 hour Return-to-Operations (RTO) solutions of several years ago! Furthermore, the cost of a recovery-as-a-service solution is typically 15 to 50 percent less than comparable in-house solutions based on repurposing a test/development environment.

How it works

Cloud-based continuity replicates data as it is changed in the traditional enterprise IT production environment. Asynchronous replication is used through a host/appliance hybrid model. Data is transported in encrypted format over the Internet or a private dedicated circuit to a hosted cloud core. In addition to application data, the operating system is also replicated for physical and virtual servers. In the event of a disaster, or at an organization's request, the environment can be rapidly re-created from cloud allocated virtual servers because all necessary context and state information is recorded in the operating system.

Applications are then restarted automatically with proper sequencing and dependencies. All network connections are rerouted to the cloud environment, and normal production processing resumes. All of this takes place in a short period of time with basically no data loss—except that which was in transit or buffered at the time of the failover.

Return of processing to the normal enterprise production environment, referred to as failback, is an important but often overlooked aspect of replication. The cloud model facilitates not only the failover, but also the failback.

At the most fundamental level, failback is essentially a reversal of the invocation process. The challenge is to correctly replicate data back to the primary production data center. This is a key concern as failback is an area in which do-it-yourself software replication solutions can easily fall prey to data corruption. To accomplish a clean failback, positioning the physical server at the enterprise site to capture replication traffic during the failback process is recommended. Data is encrypted and forwarded over the network to that physical server while the enterprise continues live operations at the recovery center. This helps ensure no downtime during the replication mapping process.

Benefits of cloud continuity

Organizations of various types and sizes, across a wide range of industries, can benefit from the unique characteristics of as-a-service DR solutions. Cloud continuity services:

- **Improve RTOs and RPOs** - enabling end-user Recovery Time Objectives (RTO) of one to four hours and Recovery Point Objectives (RPOs) of zero to 15 minutes—significantly better than many current recovery models.
- **Provide more flexibility** - allowing testing with an average of just two days' notice—much better than the six-to-nine-month delays frequently imposed by traditional DR service providers. Cloud continuity also does not impact internal, repurposed IT environments such as test and development systems.
- **Decrease costs from 15 to 50 percent** - reducing costs significantly when compared to the repurposed test/development model. Cloud continuity leverages a multitenant model, leveraged facilities, and leveraged personnel to achieve a lower total cost of ownership.

Conclusion

Continuity and disaster recovery are established and crucial elements of an enterprise-class IT environment.

Yet, many traditional tape-based and hosted continuity solutions simply no longer provide the Return of processing/Return-to-Operations performance, cost efficiencies, and testing flexibility needed by modern organizations. Not surprisingly, organizations are rethinking both the dominant service provider model and the internal approach to continuity.

The emergence and rapid maturity of as-a-service cloud continuity capabilities may fill the current gap in flexibility and in failover and recovery performance.

By understanding the current challenges and opportunities of enterprise-level continuity, and by addressing the technical and process requirements of the as-a-service approach, organizations can reap the recovery, cost, and agility benefits of cloud-based continuity.

Authors

George Ferguson

George Ferguson is the product marketing manager of Continuity Services for HP Enterprise Services and has held various portfolio, marketing, and business management roles in the areas of security, compliance, and continuity services over the past 14 years. In these roles, he has been instrumental in integrating HP's cross-company hardware, software, and services offerings into integrated solutions to help manage the operational risk of enterprises' IT-supported business processes.

Ferguson has worked in the IT industry for 33 years, of which he has spent the last 24 years at HP. He holds an M.B.A. from the Anderson Graduate School of Management at UCLA and a B.A. in computer science from Brigham Young University.

Michael Boyd

Michael Boyd is the Elections Practice Leader for Hewlett-Packard State and Local Enterprise Services. In this role, he is responsible for charting the strategic direction for elections IT services for HP and leading its elections applications delivery team. He is a seasoned professional who has delivered IT projects successfully for state governments and previously served as Governor's Policy Director and State Elections Director.

Learn more at:

To learn more about HP Continuity Services and Enterprise Cloud Services–Continuity, contact Mike Boyd, Elections Practice Leader, or visit hp.com/go/continuity.

Sign up for updates

hp.com/go/getupdated



Share with colleagues



Rate this document

© Copyright 2013 Hewlett-Packard Development Company, L.P. The information contained herein is subject to change without notice. The only warranties for HP products and services are set forth in the express warranty statements accompanying such products and services. Nothing herein should be construed as constituting an additional warranty. HP shall not be liable for technical or editorial errors or omissions contained herein.

Trademark acknowledgments, if needed.

4AA4-xxxxENW, June 2013

