

HP's Cloud Recovery-as-a-Service (RaaS)

June 2012

The protection of mission-critical services from outages has become a critical priority for many enterprises. Financial service organizations, mobile communication carriers, health care providers, and web hosting services are only the tip of the iceberg. Almost all companies have applications that will wreak havoc with their operations should they go down, often costing them thousands of dollars a minute, not to mention erosion of customer loyalty and other penalties.

Though the cost of downtime is a serious consideration, even more important is the loss of data following an outage. A company can usually survive downtime to some extent, but it may not be able to survive the loss of its corporate data.

Classic tape-recovery strategies today are woefully inadequate for these critical applications. They often provide recovery times measured in days, with potential data loss also measured in days because all data since the last backup is lost. The alternative – data replication to a hot standby site in a remote location – can provide recovery times measured in hours and data loss measured in minutes. However, this solution is often prohibitively expensive.

A common alternative is to turn to a collocation recovery service provider that will provide systems to carry a failed production load following an outage. However, it requires some time to configure the required systems at the time of an outage, the organization loses control of its disaster recovery site, the hardware and software configuration must be kept up-to-date – an error-prone process, testing is difficult and risky, and ongoing costs are still high.

What is needed is a flexible, easy to schedule continuity solution that has predictable cost with no upfront large capital expense. The new cloud-based Recovery-as-a-Service solutions provide these advantages. A synchronized copy of all application and system data is maintained in the recovery cloud via data replication. Virtual machines (VMs) stand ready to be deployed to take over production processing upon the failure of the production environment. The service cost is minimal during normal production when continuity services are not needed. However, the cloud can spring into action rapidly should the production site go down and carry on production processing with minimal downtime and data loss.

Moreover, one of the predominant causes of failover faults is corrected. Companies often find that when faced with an outage, the failover to the backup site fails because its hardware or software versions have not been kept up to date. With cloud continuity services, hardware configuration changes simply mean a change in the virtual machine and storage allocation for the applications. Software upgrades are automatically taken care of by data replication.

Because of the risk of failover faults, many companies are reluctant to test failover. Should an outage happen, their strategy is based on faith and hope. With cloud continuity services, failover is automatic and relatively risk free. A failover test can be scheduled on short notice, or a clone of the production

environment can be brought up in the cloud to test the backup environment. Should an actual outage occur, the success of the failover can be assured.

HP's Cloud-Based Continuity Services

HP has recently announced its cloud-based continuity services, *HP Enterprise Cloud Services – Continuity*.¹ These services extend its classic continuity services that it has been providing to its customers for more than two decades via its disaster recovery centers. HP's *Enterprise Cloud Services – Continuity* provides cloud-based Recovery-as-a Service (RaaS) to enterprises. For a cost estimated to be 20% to 50% of that needed to repurpose a development and test system for disaster recovery, HP's RaaS meets Recovery Time Objectives (RTOs) of one to four hours and Recovery Point Objectives (RPOs) of no more than fifteen minutes of lost data. There are no upfront capital expenses. HP's RaaS cloud services are a cost-effective alternative to traditional service-provider or in-house models.

HP's cloud continuity services are designed to protect applications running under Windows or Linux. Application may be running in a virtual environment under VMware. For continuity protection for other environments, HP provides its Managed Resiliency Services.

The HP RaaS cloud is very flexible. It only takes a few days to initially provision and deploy the backup environment. Changes can be set up in a day. Recovery testing can be scheduled on just two days' notice. The backup environment is a clone of the production environment and can be used for other purposes such as application testing, user acceptance testing, and quality assurance, without affecting production processing.

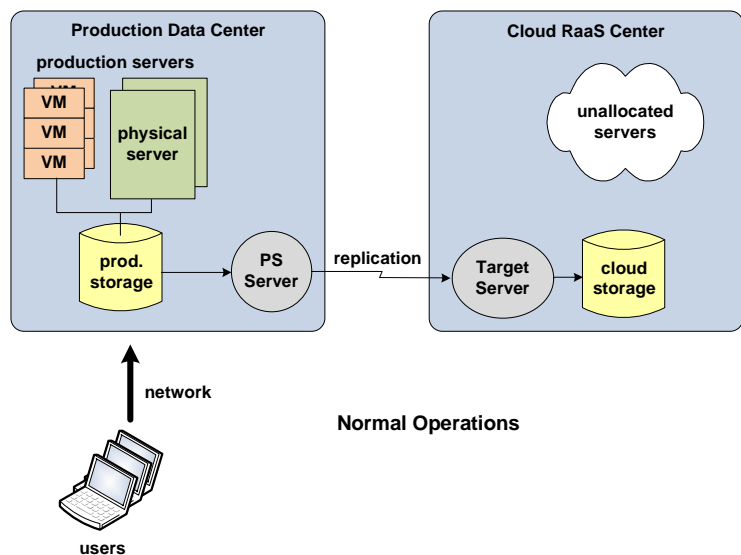
In the event of a business interruption, the production environment is rapidly recreated from cloud-allocated virtual servers. All necessary context and state information for these servers has been replicated and is immediately available to the appropriate servers. Applications are restarted in the proper order, networks are rerouted to the cloud, and normal production processing resumes. Only that data that was in transit in the replication pipeline is lost.

Equally important, fallback is easy and risk-free.

Normal Operations

During normal operations, Windows and Linux servers in the production data center provide application processing for the enterprise. Applications may be running in virtual machines or on their own dedicated servers. The storage systems may be local to the servers or may be storage array networks.

The only modification to this configuration is the replication network that keeps the cloud RaaS storage synchronized with the application data and system software changes being made by the production systems. HP-provided replication software runs on every protected production machine in the data center. Since replication is asynchronous, no transaction is held up; and the replication



¹ George Ferguson, *Cloud-based business continuity*, *HP White Paper*, May 2012.

software imposes only a small load, in the order of 2-3%, on the production servers.

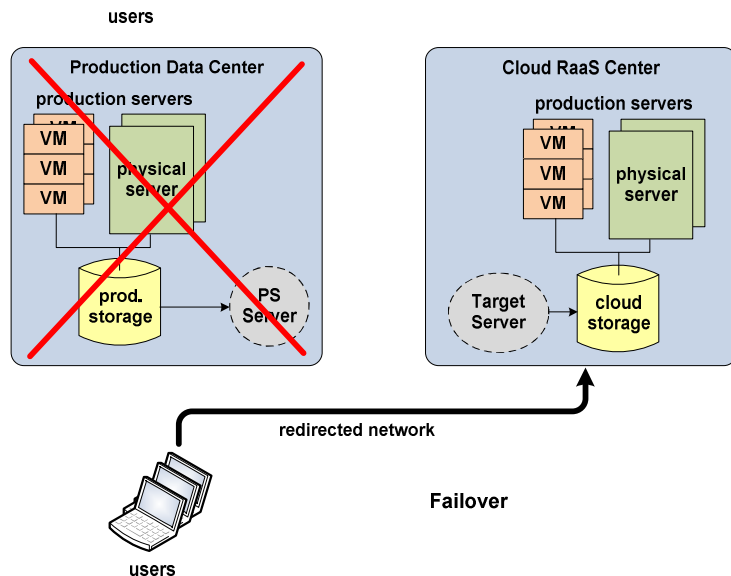
One or more dedicated Windows or Linux Process Servers (PS Servers) are provided at the data center. They receive all changes to the application data and system files that are being replicated. The PS servers are the replication engines. The replication network may be a VPN over the Internet or private dedicated communication links. The PS Servers provide the compute-intensive functions of compression, encryption, WAN acceleration, and bandwidth management for the replication network.

The PS Servers replicate the data that they are receiving to one or more target servers in the RaaS cloud that are dedicated to the production data center. These servers decrypt and decompress the data and store it on HP 3PAR storage arrays in the cloud data center. The 3PAR storage arrays provide high performance, high availability, and strong security.

Therefore, during normal operations, the only additional function when using the HP RaaS services is that all application data and system software changes are being replicated to the cloud. Resident in the cloud is a synchronized copy of all production application data (within the replication latency) and current images of the operating systems for each server and virtual machine in the production center. The RaaS cloud is ready to continue the processing functions of the production data center should the data center experience an outage.

Failover

Then the unthinkable happens! Application processing for part or all of the data center is halted by some serious event. System operators and management scurry to determine the cause and the impact of the outage. Decisions are made whether to try to recover the systems in the data center or to fail over one or more of the production systems to the RaaS cloud.



The decision to failover is not a simple decision, as it might be faster to attempt a local recovery. For this reason, the initiation of the failover sequence is never made automatically. However, once the decision is made, it is simply a matter of calling the cloud operators and instructing them on the extent of the failover to be executed. Failover is quickly initiated via a control console and proceeds completely automatically from that point.

The required virtual servers are allocated from the virtual server pool in the cloud, and their operating system images are built from the replicated data. The pertinent application data is attached to each server, and the servers are booted in the order required by the application. The VMs in the recovery center match the servers in the production center. The user network is redirected to the RaaS center, and production operations are restored. The allocated servers are dedicated to the company as long as they are needed.

A company can elect to have its servers in the RaaS cloud be leveraged from the server pool. Alternatively, if there is a concern about the available capacity in the cloud in the event of a wide-spread event taking out several data centers, the company can elect to reserve servers for their exclusive use. Reserved servers command a higher price. For leveraged servers, HP follows stringent risk-management

controls and has supported its customers' leveraged requirements successfully for over 25 years of classic recovery services.

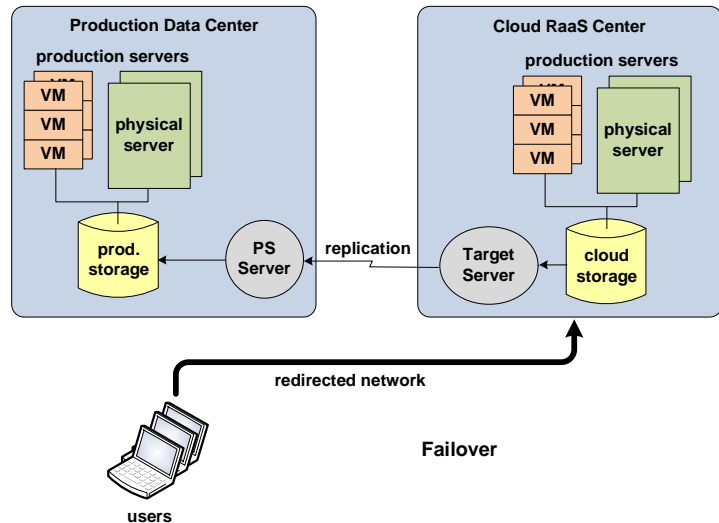
Fallback

When the outage at the production center has been corrected, it is time give the data center control over production processing once again. This process is called *fallback*.

Fallback occurs in two phases – replication and recovery.

Replication

The first step is to resynchronize the production system's application data and system software with the current state that is represented by the RaaS cloud storage. To do this, the replication channel is reversed, and the production system's databases are brought up to date. During this phase, applications continue to run in the recovery center.



Recovery

Once the production databases and system software have been brought up-to-date, production can fall back to the production data center. However, there is no need to hurry this. Fallback can be scheduled for a convenient time for the company. During this wait, production continues at the recovery center, and the production databases and system software are kept current via the replication channel.

When it is time to fallback, production processing is paused for the few minutes that it takes to complete any transactions in process and to drain the latest changes from the replication channel to the production storage devices. The user network is redirected back to the production data center, and processing can then resume at the production data center with no lost data and no lost transactions. Replication to the cloud recovery center is restarted to keep it properly synchronized.

Security

Security is always a concern when a company allows its data to be stored in a cloud. One common concern is where the data will be located. HP recognizes that companies and governments are often restricted as to the location of their data by confidential and regulatory restrictions. HP is building RaaS cloud centers around the world to accommodate these concerns. Recovery cloud centers currently exist or are planned for the United States, the European Union, Australia, and Japan. Other countries will be added over time.

The full range of physical and logical security is provided at all recovery centers. Physical access is controlled by key card or biometric access, with video surveillance and 24x7 onsite security personnel. Networks are protected by firewalls, intrusion detection, VPN monitoring, and VLAN separation. All data transfers are encrypted in flight, and data can be encrypted in place if desired. Leveraged storage is triple-scrubbed before reallocating it. Access to user accounts is based on the customer's standard authentication procedures.

All security events are logged by HP's ArcSight Security Information and Event Management (SIEM) facility and analyzed to provide Security Incident Responses through HP Global Security Operations Centers.

HP Continuity Services

To aid customers in achieving assured continuity of their IT services, HP offers many consulting services, including

- Business Continuity Consulting
- High Availability Consulting
- Workplace Continuity Services
- Data Center Continuity Services

Summary

Cloud-based recovery services are available today to provide hot standby, data replication RPOs and RTOs at tape backup prices. HP's Enterprise Cloud Services – Continuity is an excellent example of such a service. With no upfront capital expenses and little effort, this service can reliably provide RTOs of just a few hours with only minutes of lost data following a primary production data center outage.