HPE Shadowbase Solutions for the Cloud

A Gravic, Inc. White Paper

Executive Summary

Companies are moving more and more IT services and utility applications to public clouds to take advantage of the economy and flexibility of cloud computing. There are no initial expenditures to purchase equipment, datacenter space, and staff. A company must only cover the costs of the computation and storage (and in most cases communications) resources that it uses.

However, there remains a reluctance to move critical applications to the public cloud and for good reason. Notable instances of public cloud failures (service outages) and data loss have been reported, and security is a concern because the organization has lost direct control over its environment.

There is motivation, though, to find a way for critical applications to take advantage of the benefits of cloud computing. A hybrid approach that assigns critical processing to highly available private critical systems such as HPE NonStop servers (among others) and noncritical processing to the public cloud is a concept that is gaining momentum. The HPE Shadowbase data replication engine plays an important role in this approach by providing the 'glue that binds' virtual machines running in the public cloud to highly available private servers handling the critical roles.

There are many HPE Shadowbase-supported architectures that allow critical systems to take advantage of cloud computing by offloading non-critical functions to the cloud infrastructure. These include:

- Using the critical server as the master database while replicating the database (or portions thereof) to the cloud applications and cloud storage infrastructure for local query processing to distributed cloud users.
- Backing up cloud-resident critical data onto a critical system either periodically or on a continuous real-time basis.
- Moving noncritical transactions to applications running in the cloud if the critical system becomes heavily loaded (load sharing).
- Integrating distributed applications and data running both in the cloud and on critical systems (a form of distributed data integration and application integration[1]).
- Integrating multiple heterogeneous clouds for application interoperation, including the capability to run an application in an active/active mode in two clouds for a continuously available business continuity architecture.

In this white paper, we discuss the role that HPE Shadowbase replication solutions can play in public, private, and hybrid-cloud computing for critical applications and how it can lower your IT costs, improve fault tolerance of your applications, and increase flexibility.

Hewlett Packard Enterprise globally sells and supports Shadowbase solutions under the name *HPE Shadowbase*. For more information, please contact your local HPE Shadowbase representative or visit our website.

[1]For more information on HPE Shadowbase solutions for data integration and application integration, please see our companion documents: *HPE Shadowbase Streams for Application Integration* and *HPE Shadowbase Streams for Data Integration*.

Table of Contents

Executive Summary......2

Cloud Computing......4
Architecture......4
OpenStack......5

Integrating the HPE Shadowbase Data Replication Engine with the Cloud......6
The HPE Shadowbase Architecture......6
The HPE Shadowbase Replication Engine in the Cloud......7

HPE Shadowbase in the Cloud Use Cases......8
Critical System as the Master Database......9
Critical System as a Backup Database......9
Cloud Bursting......9
Distributed Cloud Application Integration......9
Cloud Integration for High/Continuous Availability as well as Application Interoperability......9

HPE Shadowbase Data Replication Software – The Ideal Way to Integrate Applications with the Cloud......10

Summary......11

International Partner Information......12

Gravic, Inc. Contact Information......12

Table of Figures

Figure 1 – Virtual Machines in the Cloud......4
Figure 2 – The OpenStack Architecture......5
Figure 3 – The HPE Shadowbase Data Replication Engine......6
Figure 4 – HPE Shadowbase (SB) Replication Engine in the Cloud (External Interfaces)......7
Figure 5 – HPE Shadowbase Cloud Use Cases......8

HPE Shadowbase Solutions for the Cloud

Cloud Computing

One common model for cloud computing comprises a network of computer resources, including CPU, memory, and persistent storage. These resources are assigned to users of the cloud on an as-needed basis. If a particular application running in the cloud is suddenly presented with an increased transaction rate, the cloud infrastructure will assign more resources to it. As the application's load decreases, excess resources are returned to the cloud pool. The cloud business model is pay-as-you-go, meaning a user pays only for the resources actually utilized.[2]

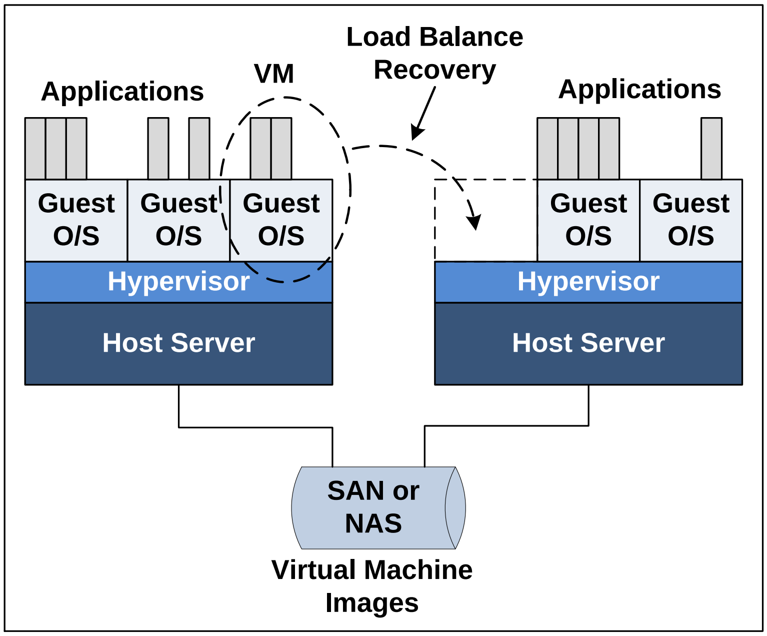

Figure 1 – Virtual Machines in the Cloud

Architecture

The architectural foundation for clouds is virtualized machines (VMs), as shown in Figure 1. A VM is a virtualized server that runs with other VMs on a single physical host server. A VM comprises a guest operating system running one or more applications. Supported guest operating systems typically include Linux, Windows, some UNIX operating systems, as well as HPE's new NonStop virtual machine, vNonStop.

The VMs share the compute, memory, communications, and storage resources of the host server. Typically, multiple interconnected host servers exist within the cloud. Access to and allocation of these resources is adjudicated by a *hypervisor*.

When a user decides to run an application, the cloud will often create a VM running on some provisioned physical server in the cloud environment. The guest operating system to be used for that VM will be booted, and the user application will be started. Thereafter, the user's application runs just as if it had sole possession of its own physical server – it is unaware that it is sharing resources with other VMs. Except for the overall load sharing of the physical resources, the applications and data running within the VM are isolated from all other VMs running on the same physical equipment.

A key point is that the physical representation of a VM is nothing more than a server image stored on SAN or NAS storage accessible by all physical servers in the cloud. The server-image file contains the executables

[2] This model is often referred to as a "capacity on demand" model. In a private compute environment, the excess capacity typically sits idle and unused until loading increases to the point where it is needed, forcing the user to pay for equipment and infrastructure that is under-utilized. In the cloud paradigm, the model shares the excess resources among many cloud customers, thereby spreading the extra costs of the additional resources among many users, resulting in lessening the overall costs for every user.

for the guest operating system used by the VM and the executables for all of the applications that are to be started. This accessibility provides a great deal of flexibility in a cloud. A VM can be moved from one physical server to another simply by terminating it on the old server and rebooting it onto the new server via its server image on the common storage. This flexibility allows the cloud (specifically, the hypervisor) to move VMs around for load balancing and to relocate them in the event of a physical-server failure.

Equally important, if a VM begins to handle a larger load and requires more resources than it can get from its current physical server, it can be moved to another server that has the necessary resources. As that need dwindles, the VM can be relocated to a server whose resources more closely match its needs. The resources used by a VM are monitored by the cloud and are used to bill the user for their actual usage.

There are *public* clouds, *private* clouds, and *hybrid* clouds. A public cloud serves any user who wishes to subscribe to its services. A private cloud is owned and managed by an organization, for its own internal use. An application running in a hybrid cloud spreads its resource requirements across public and private clouds or other internal IT infrastructure (for example, an HPE NonStop server which is not running as a part of any cloud).

OpenStack

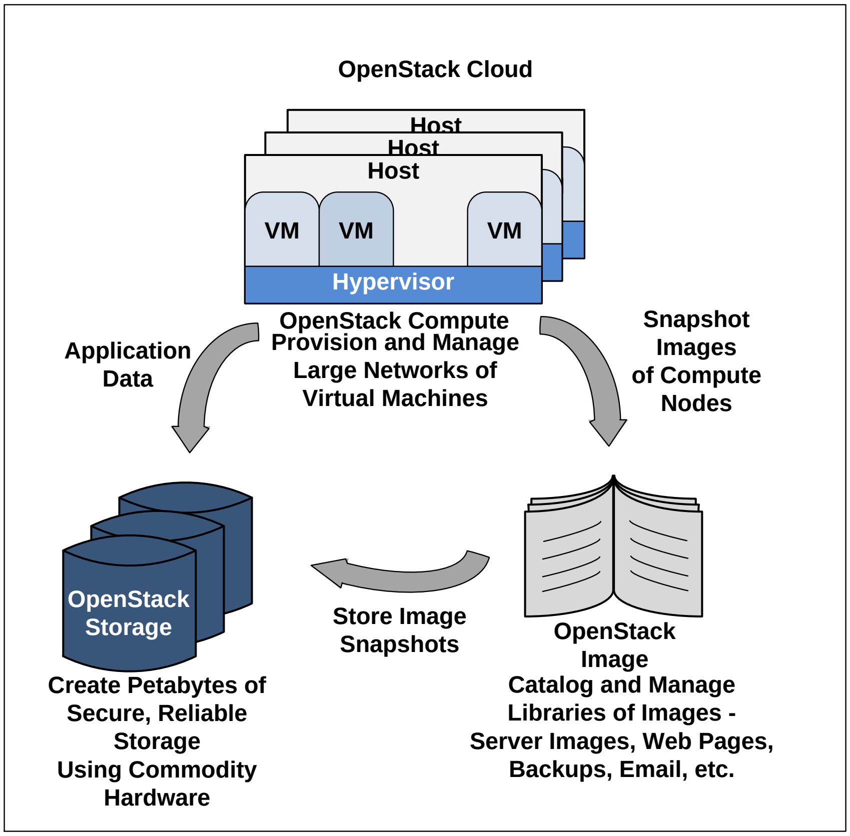

Figure 2 – The OpenStack Architecture

Today's large public clouds include Amazon Web Services (AWS), and Microsoft Azure. The problem with these current clouds is that they are proprietary. Once a user selects a cloud for their applications, they are locked in. It is very difficult to move to another cloud provider.

A move is underway to correct this situation. OpenStack is an open-source cloud suitable for private and public clouds. As the OpenStack website explains, OpenStack is a "large scale open source cloud project and community established to drive industry standards, end cloud lock-in, and speed adoption of cloud technologies by service providers, enterprises, and government agencies to build massively scalable public and private clouds using freely available Apache-licensed software." OpenStack is patterned after Amazon's AWS cloud. It consists of three major facilities (as shown in Figure 2):

OpenStack Compute includes VMs to provide compute capacity. Most major hypervisors are supported, as are Windows, Linux, and some UNIX guest operating systems. An OpenStack VM can have two IP addresses – one for internal communications and one for external communications.

OpenStack Storage is a massive block store that provides persistent storage. It is modeled after Amazon's Elastic Block Store (EBS). Storage assigned to an OpenStack VM does not survive when the VM is terminated. If the storage must be persistent, it must be directed to OpenStack Storage.

OpenStack Image uses OpenStack Storage to catalog and manage server images.

Among others, OpenStack is used by the Rackspace cloud and by HPE Helion, which allows companies to set up their own private, public, or hybrid clouds with HPE solutions and guidance. OpenStack is enjoying explosive growth with companies that are implementing their own clouds. If OpenStack becomes the de facto standard for new clouds, the problem of cloud lock-in will largely disappear.

Integrating the HPE Shadowbase Data Replication Engine with the Cloud

The HPE Shadowbase Architecture

Figure 3 shows a simplified view of the HPE Shadowbase data replication engine. The role of the Shadowbase architecture is to replicate data generated by a source environment to a target environment in real-time (as the source data is generated). The source and target environments can be on the same system, or completely different, heterogeneous systems.

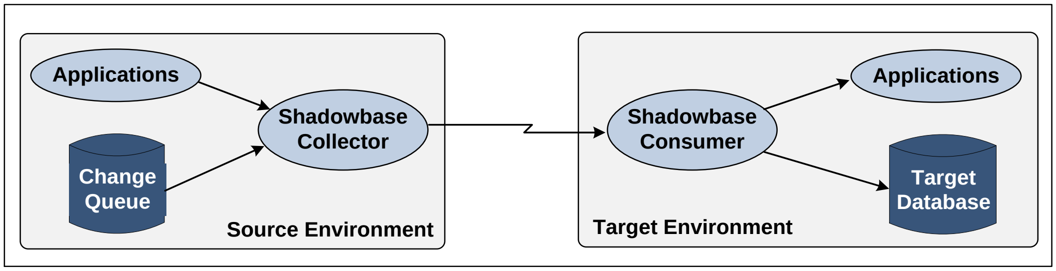

Figure 3 – The HPE Shadowbase Data Replication Engine

The Shadowbase Collector resides in the source environment. It is fed data updates from applications (the so-called *push* model), or it reads updates from a change queue of some sort (the so-called *pull* model). The change queue may be generated by the applications, or it may be a transactional data log such as those maintained by transaction processing managers, e.g., the TMF Audit Trail in HPE NonStop systems or the Redo Log in Oracle databases. Various database managers are supported by different cloud implementations. For instance, OpenStack supports MySQL, and Amazon AWS supports MySQL, Oracle, and SQL Server.

As the Shadowbase Collector receives updates, it sends them over a communication channel to the Shadowbase Consumer resident in the target environment. There, the Shadowbase Consumer applies the data updates to a target database and/or sends them to target applications. In this way, the target environment is kept synchronized with the source environment.

The HPE Shadowbase replication engine is heterogeneous. The source and target platforms may be different, as may be the source and target databases – the Shadowbase engine provides all the format conversion required between the two environments. The engine guarantees the transactional integrity of the data it is applying to the target environment. All target databases and applications are available for application use (read and write) during replication.

Replication may be uni-directional, or bi-directional (carried out in both directions). In the latter case, a change to either database is replicated to the other. Thus, an application may be active in both environments, and the two environments are kept synchronized. A transaction can be sent to either environment for processing and will achieve the same result. This architecture is known as an *active/active* system.

The HPE Shadowbase Replication Engine in the Cloud

As shown in Figure 4, the HPE Shadowbase replication engine can be integrated into the cloud to transfer data internally between cloud applications/databases as well as between cloud applications/databases and external systems. This capability depends upon the fact that for many cloud systems (OpenStack and Amazon AWS included), a VM may have two IP addresses – one for internal communications and one for external communications.

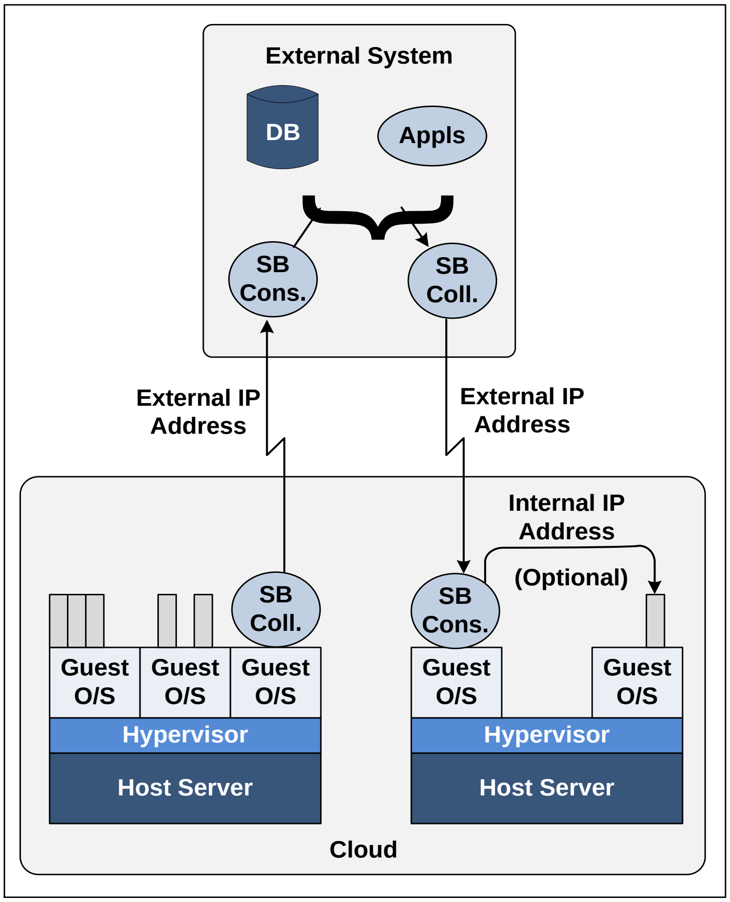

Figure 4 – HPE Shadowbase (SB) Replication Engine in the Cloud (External Interfaces)

For data transfer from the cloud to an external system, the HPE Shadowbase Collector ("SB Coll." in the figure) is configured to run inside the customer's VM in the cloud. In Figure 4, the Collector is shown running in a Linux guest operating system. However, the HPE Shadowbase replication engine can run in most guest operating systems supported by the cloud, including HPE vNonStop.

The cloud-resident Collector uses the local change queue to collect data from local cloud applications or databases. In the event-driven push model, the Collector gathers the data from the change queue as it is generated and sends the updates that are to be replicated as soon as they are ready. In the pull model, the Consumers ("SB Cons." in the figure) are configured to connect to the Collector and 'pull' the changes periodically. Regardless of the model, the Collector uses its external IP address to send updates to the HPE Shadowbase Consumer located in the external system. The Consumer applies the updates to a target database or sends them to target applications, as appropriate.

For data transfer from an external system into the cloud, the Shadowbase Consumer is configured in a VM (or inside the customer's VM). It receives updates via its external IP address from the Shadowbase Collector resident in the external system and sends these updates to applications via their internal IP addresses or applies these updates into the customer's databases directly.

Multiple instances of the HPE Shadowbase replication engine can be configured to support multiple external systems. Thus, by configuring Shadowbase components inside VMs running in the cloud, the engine can transfer data in either direction between the cloud and external systems.

HPE Shadowbase in the Cloud Use Cases

There are many ways that the HPE Shadowbase replication engine can be used to enable businesses to take advantage of the benefits of the cloud while avoiding some of the issues discussed earlier. Most of these techniques exploit the replication engine's ability to keep cloud applications and data synchronized with external systems, thereby enabling businesses to shift noncritical functions to the cloud while maintaining critical or sensitive functions in the company's internal systems (i.e., in a hybrid cloud architecture).

In the following examples, the critical systems could be HPE NonStop servers; they are the pinnacle of massively parallel, continuously available systems and are typically used for the most demanding IT business services. Or, they could be any of the many other platform and database types supported by HPE Shadowbase technology. Potential uses of HPE Shadowbase software to add resiliency to cloud applications are displayed in Figure 5.

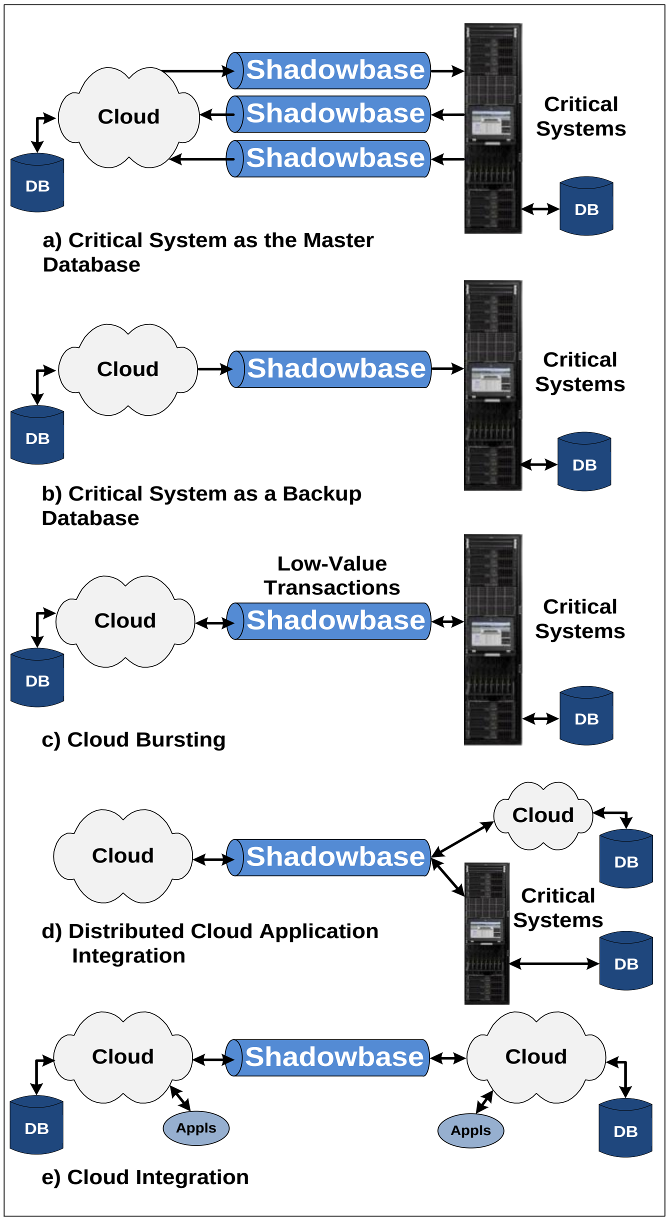

Figure 5 – HPE Shadowbase Cloud Use Cases

Critical System as the Master Database

In this example the critical system (such as an HPE NonStop system running a mission-critical application) is used as the master database with a replicated copy kept locally by all cloud applications (Figure 5a). This configuration can operate in two modes depending on the application architecture.

In the first mode, all updates are made to the critical system master database. The Shadowbase engine replicates these updates to local copies of the database in the cloud, so the cloud applications have access to the latest data and can perform fast, local, read operations on that data. This configuration is particularly useful for applications that have a high look-to-book ratio. For instance, a travel agency system spends a lot of time searching for deals and answering queries about hotel, plane, and car reservations. These are read operations and can all be executed quickly against the cloud local replicated copies of the master database. Only when a reservation is to be made is it necessary to update the master database.

In the second mode, update operations are made by applications to the cloud local database(s). The Shadowbase engine replicates these updates to the critical system master database, applies them, and can re-replicate these changes back to any other cloud-based copies of the data. In this way, the master database is kept current and comprises all updates made by applications anywhere in the cloud, it is the single, complete, database of record.

Critical System as a Backup Database

Another solution to protect critical data in the cloud is to provide backup facilities outside of the cloud. The Shadowbase engine can be used to make periodic (or continuous, streaming) backups of data that is being applied to a cloud-resident database (Figure 5b). The backup database can then be used to recover and restore data in the event that the cloud databases are lost or become corrupted.

Cloud Bursting

Cloud bursting entails sending some of the transaction load to the cloud if the critical system becomes heavily loaded. If the critical system begins to grow heavily loaded, then it can send low-value transactions to the cloud applications for processing. High-value transactions can continue to be processed by the critical server. Shadowbase technology supports cloud bursting architectures (Figure 5c) by replicating database changes between the cloud and critical system databases, thereby ensuring that applications running in either environment are operating on the same, current, and consistent data.

Distributed Cloud Application Integration

Some applications may be distributed between both the cloud and critical systems (or between clouds). When an event is generated by one application, it needs to be propagated to another application for further processing. The travel agency application is a good example. When a reservation is to be made, the cloud based “look” application may raise an event to cause the critical system “book” application to actually make the reservation. Shadowbase architecture (Figure 5d) provides several mechanisms to detect such application events when they are raised, and send them to target applications (or databases) to be acted upon. In this way, Shadowbase technology provides the event detection and distribution fabric between distributed applications, enabling applications running in the cloud to be integrated with other applications running either in the cloud or on critical systems.

Cloud Integration for High/Continuous Availability as well as Application Interoperability

Figure 5e shows Shadowbase replication being used to integrate one cloud environment with another (for high or continuous availability purposes, for example). In some cases, the cloud operators may provide sufficient capability to provide these business continuity services, but their offerings may not be sufficient, complete, able to work heterogeneously, or simply be too expensive. For these cases, the HPE Shadowbase data replication engine can replicate the data from one cloud environment uni-directionally into another (for example, to implement a highly available disaster recovery cloud-based architecture), as well as bi-directionally between two cloud environments (for example, to implement a continuously available active/active business continuity cloud-based architecture). Please refer to the companion Shadowbase white paper entitled *Choosing a Business Continuity Solution to Match Your Business Availability Requirements*[3] for more

[3]Please read the Gravic white paper, *Choosing a Business Continuity Solution to Match Your Business Availability Requirements*.

information on building high and continuous availability architectures for your applications and data (regardless of whether they are running in a cloud infrastructure or not).

In another similar example, the major clouds today are proprietary and cannot work together. However, the HPE Shadowbase replication engine can be used to connect VMs and their applications in one cloud with those of another heterogeneous cloud. In this case (Figure 5e), Shadowbase VMs are communicating with each other. This architecture allows applications to be spread among multiple clouds and still interoperate. In fact, the same application can be running in multiple clouds in an active/active mode. If one cloud fails, then all transactions are simply routed to the surviving cloud, and the application users are unaffected. Moreover, if the shift in transaction workload is significant, the surviving cloud will automatically assign additional resources to the application to maintain the application's performance levels. Shadowbase software is able to handle the necessary data reformatting between heterogeneous clouds (databases and applications). User exit processing is also available to handle more complex data transformation.

HPE Shadowbase Data Replication Software – The Ideal Way to Integrate Applications with the Cloud

The HPE Shadowbase replication engine is designed to integrate diverse applications and databases. Its application to the cloud is simply an extension of the multiple source and target environments that it already supports.

A large, cloud-based application may comprise many different systems with different missions. Each system is implemented on a "best-fit" platform. There may be a myriad of heterogeneous platforms, applications, and databases that make up the application. Add to this mix the application components that reside in one or more clouds. A powerful, flexible, fast, and reliable data distribution fabric is required to interconnect these systems. The HPE Shadowbase data replication engine fulfills this role.

Data must be passed from external information sources to the cloud applications and between the systems and cloud applications themselves. Data transfer must be secure and HPE Shadowbase software supports encrypting the communication traffic, for example, via encrypted TCP/IP sessions using a proxy server. Data transfer must be very fast, with high capacity and minimum latency. The Shadowbase engine's process-to-process architecture eliminates disk-queuing points that can slow down replication and otherwise consume extensive disk resources. Sub-second replication latency is achieved. The replication engine can be multithreaded, including the communication channels, so that any desired data transfer capacity can be attained.

In some cases, the data is being generated by an application. In other cases, it is being taken from a database. The data must be delivered to a target application or to a target database. The HPE Shadowbase engine supports heterogeneity. It can receive data as it is generated from any supported application or database and can deliver it to any supported application or database, inside or outside the cloud.

Replication between diverse applications and databases can require complex reformatting and restructuring of the data. HPE Shadowbase software includes provisions for many types of reformatting. If it does not inherently support a specific type of reformatting, it provides user exits so the user can customize extended reformatting algorithms.

The Shadowbase engine can replicate data synchronously or asynchronously. Asynchronous replication takes place after the fact and has no impact on the source application. Synchronous replication guarantees that an update is either made to both the source and target environments or that it is made to neither. Both synchronous and asynchronous replication guarantee the transactional consistency of the target database, and that the target database is available for application processing during replication.[4]

HPE Shadowbase software is architected for continuous availability. If one of its components fails, it is automatically restarted and replication continues uninterrupted. If the target environment fails, then Shadowbase architecture queues all events until the target environment is restored to service. It will then drain its queue of saved events to bring the target environment back into synchronization with the data source, and resume replication.

[4]Contact Gravic for more information about HPE Shadowbase synchronous replication.

The data distribution fabric within a complex cloud environment and between the cloud and external systems must be low latency and provide high capacity. It must be fundamentally heterogeneous and be able to deal with any application or database as a source or as a target. It must be able to reformat and restructure data on the fly as it moves data from one source to another totally different target. It must be highly reliable. All of these are attributes of Shadowbase replication: low latency, high capacity, heterogeneous, powerful message processing, flexible end points, and high availability. Integrating heterogeneous data resources is a formidable challenge; a challenge that has been solved by HPE Shadowbase software, which is positioned to form the ideal foundation for the data distribution fabric for cloud computing, public, private, or hybrid.

Summary

Tremendous incentives exist for organizations to move to cloud computing. There are no initial expenditures to purchase equipment, datacenter space, or staff. A company must only cover the costs of the compute and storage resources that it uses.

However, the track record for clouds is still spotty. Availability is an issue; major cloud failures routinely make the headlines. Security is another issue. Thus, executives are reluctant to trust their critical applications and data to the cloud. However, it is evident that the benefits of cloud-based storage and processing are here to stay.

The HPE Shadowbase data replication engine can alleviate many of these concerns. It serves as the glue that binds a company’s critical applications and data to trusted corporate internal systems while allowing the cloud to perform less critical functions, for which the cloud is well-suited. With HPE Shadowbase software solutions, an optimum compromise of application safety versus cost can be achieved, enabling businesses to take advantage of the benefits of the cloud while avoiding the pitfalls.

International Partner Information

Global

Hewlett Packard Enterprise
6280 America Center Drive
San Jose, CA 95002
USA
Tel: +1.800.607.3567
www.hpe.com

Japan

High Availability Systems Co. Ltd
MS Shibaura Bldg.
4-13-23 Shibaura
Minato-ku, Tokyo 108-0023
Japan
Tel: +81 3 5730 8870
Fax: +81 3 5730 8629
www.ha-sys.co.jp

Gravic, Inc. Contact Information

17 General Warren Blvd.
Malvern, PA 19355-1245
USA
Tel: +1.610.647.6250
Fax: +1.610.647.7958
www.shadowbasesoftware.com
Email Sales: shadowbase@gravic.com
Email Support: sbsupport@gravic.com

Hewlett Packard Enterprise Business Partner Information

Hewlett Packard Enterprise directly sells and supports Shadowbase Solutions under the name ***HPE Shadowbase***. For more information, please contact your local HPE account team or visit our website.