

the *Availability Digest*

www.availabilitydigest.com
[@availabilitydig](https://twitter.com/availabilitydig)

Cloud Resiliency December 2017

Reliable and resilient application architectures are fundamental to today's data centers. *Reliable* means that the failure of a system component is rare. *Resilient* means that if a component does fail, it can be restored to service or its services can be transferred to another operational component quickly. In today's data centers, reliability and resiliency are achieved by incorporating redundant servers whose databases are kept synchronized via data replication. In this way, the backup server is immediately ready to take over processing should the production server fail.



The emergence of cloud computing has dramatically changed the way we think about application resiliency. When an application is running in the cloud, thin provisioning and auto-scaling maximize the efficiency of the resources made available to the application. If an application's workload should suddenly increase, the cloud can immediately assign additional resources to it to handle the new workload. If the application's workload decreases, the cloud can recover unneeded resources from the application.

Furthermore, spinning up secondary and tertiary disaster recovery environments is easy, as the cloud simply assigns the resources necessary to these environments. This provides a level of resiliency not found in typical data center environments.

Cloud Services

Increasingly, today's applications reside in architectures in which cloud services are critical path components. If the cloud experiences problems, an application can fail.

What happens if a key cloud service fails? That is up to the cloud service provider. The application administrator has little say over the management of the cloud (that is the function of a properly written SLA – a Service Level Agreement). Therefore, it is important that today's cloud-service enabled applications require redundancy at the cloud services layer.

Introducing cloud service redundancy starts with managing application workload. How do you direct work across multiple cloud service providers? This is what must be done in a redundant cloud environment.

The Hybrid Cloud

Many IT departments are turning to hybrid cloud – a blend of traditional and cloud services – for greater agility and flexibility to meet changing business requirements.

With hybrid cloud architectures as a foundation, enterprise workloads can originate from a wide range of sources such as traditional data centers, private clouds, public clouds, and hybrid clouds. Furthermore, the workload can be processed in any of these environments. Therefore, organizations can run their workloads wherever they can achieve optimal efficiency.

IT Challenges

The uncontrolled growth of cloud services is creating new challenges for the IT staff:

- Environments can be difficult to integrate, manage, and secure.
- Many IT organizations are spending far more time and resources maintaining the cloud infrastructure.
- As a result, these organizations are spending less time innovating modern services for the business.

A growing dependence on capabilities supplied by third parties makes it difficult for IT staff to have direct control over all aspects of their cloud services

Enterprises must design their resiliency strategies to safeguard all of the resources required to execute their business processes:

- Data is becoming a valuable business asset.
- IT services are being delivered from a wide range of sources, including traditional data centers, private and public clouds, and hybrid clouds.
- As cloud computing becomes more IT resource intensive, organizations must now think in terms of overall business resiliency - business processes, workflows, technology, and policies that support continuous availability of products and services.

Organizations must try to avoid, rather than recover from, disruptions

Reliability and Resiliency

A reliable service is a service that works as it was designed to work, responding in a predictable fashion when it is needed. One way to improve reliability is to build a service that is resilient. If the service fails, it can be immediately restored. The service has the ability to withstand certain types of failure and remain functional.

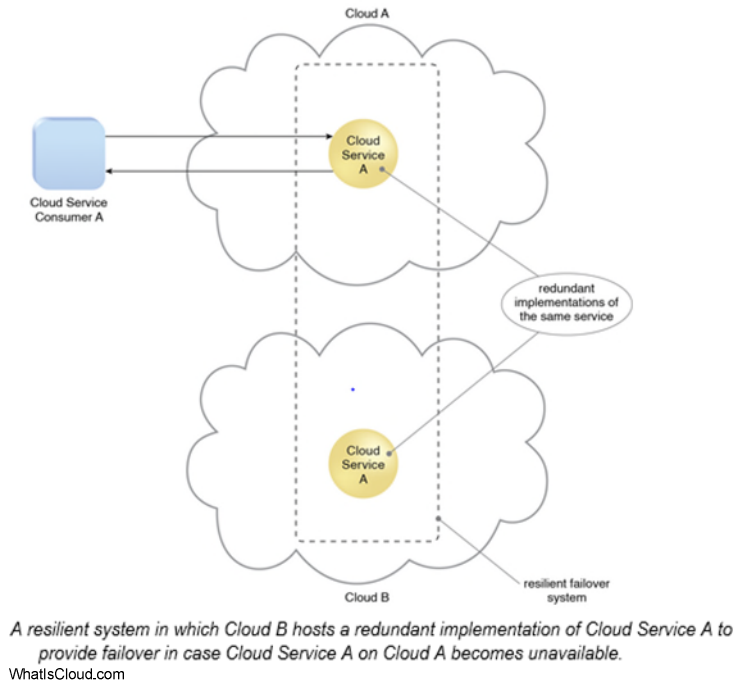
Resilient computing is a form of failover that distributes redundant implementation of IT resources across physical locations. It is configured so that if one resource becomes deficient, processing is automatically handed over to another redundant resource.

Resiliency is not architected only at the hardware level. It is the combination of intelligent infrastructure design, coupled with intelligent software design, which results in highly-reliable cloud services

In cloud computing, resiliency can refer to redundant IT resources within the same cloud (but in different physical locations) or across multiple clouds.

Many enterprises assume the cloud is reliable and immune to failure. They need to understand the risk associated with the failure of their cloud services. The cloud is a utility, and utilities fail. It is critical that cloud customers be prepared for cloud downtime.

Applications need to be designed to accommodate failure. This supports the concept of stateless design for applications. The application can pick up at any point and not have to worry about reconstructing state.



Furthermore, the application manager needs to have transparency into the cloud provider's infrastructure so that he can properly manage his applications.

Using a cloud provider with a global presence that can distribute an environment broadly can help build application resilience.

DNS

DNS (Domain Name Service) is one of the most powerful tools for managing cloud workload. With DNS, one can adjust the application workloads across cloud services. Workloads can be shifted from cloud to cloud in response to real-time conditions. Applications can be failed over to operational clouds should a cloud service be broken.

Summary

Cloud computing is an entirely new method for application processing. We are still in the early stages of using this environment. As time goes on, techniques will be further refined for cloud computing, and the management of applications running in the cloud will become more straight-forward.

Acknowledgements

Information for this article was taken from the following sources:

Resiliency in the age of cloud services, *Network World*; November 27, 2017.

Cloud services: building resiliency and business continuity, *Cloud Blogs*; January 17, 2013.

End-To-End Resiliency In A Hybrid Cloud World, *Forbes*; October 7, 2016.

Resiliency, *WhatsCloud.com*; undated

Cloud availability and resiliency: Planning for failure, *Tech Target*; July 29, 2011.