the Availability Digest

www.availabilitydigest.com
@availabilitydig

Ensuring Resiliency Within the Cloud

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 either can be restored to service or that its services can be transferred quickly to another operational component. 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. Should an application's workload increase suddenly, the cloud can immediately assign additional resources. 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.

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. To be reliable, a service should be replicated, with the replicate databases kept synchronized via data replication.

Another way to improve reliability is to build a service that is resilient. If the service fails, it can be restored quickly. The service has the ability to withstand failure and to 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.

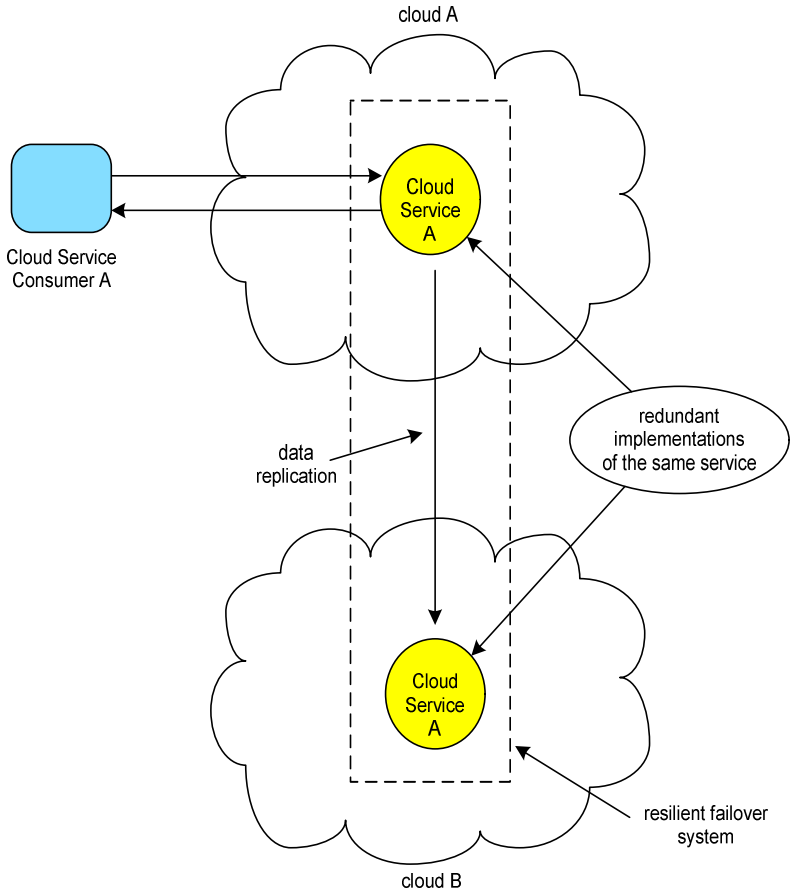

A resilient system in which Cloud B hosts a redundant implementation of Cloud Service A to provide failover in case Cloud Service A on Cloud A becomes unavailable. (WhatIsCloud.com)

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, one that can distribute an environment broadly, can help build application resilience.

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? It is important that today's cloud-service enabled applications be provided redundancy at the cloud-services layer and that recovery of a cloud-service failure be executed quickly – preferably in seconds, not minutes or hours – in order to provide resiliency of the cloud service.

This can be accomplished by managing application workload across multiple cloud-service providers. If this is done, a cloud failure may be mitigated quickly by shifting failed workload to another cloud-service provider.

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.

If an organization supports multiple capabilities for enterprise applications, resiliency can be significantly increased because there may be many systems on which a failed application may be restarted in the event of a system failure.

IT Challenges

The growth of cloud services is creating new challenges for the IT staff when it comes to achieving high resiliency:

- Environments can be difficult to integrate. In order to move an application quickly to a new environment to achieve resiliency, the environments must be integrated.

- Many IT organizations are spending far more time and resources maintaining the cloud infrastructure. This activity is made more complex if the infrastructure must be designed to support resiliency.

- As a result, these organizations are spending less time innovating modern services for the business and more time on resilient infrastructure.

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. In order to provide resiliency, the IT staff needs direct control of its infrastructure.

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. Lost or corrupted data must be recovered immediately.

- IT services are being delivered from a wide range of platforms, including traditional data centers, private and public clouds, and hybrid clouds. If any one of these fails, applications requiring high resiliency must be moved to a surviving platform.

- Data must be replicated between all of these sources so that it is available should a source fail. Doing this makes the data resilient.

- 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 the continuous availability of products and services.

Organizations must try to avoid, rather than recover, from disruptions. However, disruptions are inevitable. Fast recovery from a disruption is the key to resiliency.

DNS

DNS (Domain Name Service) is one of the most powerful tools for managing cloud workload. The DNS service assigns URLs to IP addresses so that a request made to a given URL is routed to the appropriate server. In the context of this article, the servers are located in a cloud.

With DNS, one can adjust the application workloads across cloud services by remapping a URL to other servers. 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. Only in this way can an application be made resilient.

Summary

Cloud computing is an entirely new method for application processing. We are still in the early stages of using this environment. The resiliency of applications moved to a cloud is of paramount importance. If a cloud should fail, the application must be rapidly recovered either by recovering the failed cloud or by moving the application to a surviving cloud. A high degree of resiliency results in minimum downtime and maximum availability for the application.

As time goes on, techniques will be further refined for cloud computing; and the management of applications running in the cloud will grow more straightforward. Application resiliency will improve as it becomes easier to move applications within or between clouds.

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