CloudVOTE

Photo Communications Corp. DBA Meridia Audience Response

CloudVOTE Resiliency and Backup

Last Update: 12/09/16

Table of Contents

CloudVOTE Resiliency Features

As with availability considerations, CloudVOTE has a built-in resiliency that's designed to support a robust disaster recovery. CloudVOTE is hosted in a Microsoft-managed Azure platform, which conforms to the most of certification and regulatory requirements in the industry.

Correct application of availability features and strategies is an important part of disaster-proofing your application. CloudVOTE offers these pillars of resiliency in the Azure environment:

Multiple Datacenter Regions

CloudVOTE is housed in datacenters in many regions around the country and worldwide. This infrastructure supports several disaster recovery scenarios, such as the system-provided geo-replication of Azure Storage and SQL databases to secondary regions. Your application can be easily and inexpensively deployed as a cloud service to multiple locations around the world. Deploying data and services to multiple regions helps protect CloudVOTE from major outages in a single region.

Azure Traffic Manager

When a region-specific failure occurs, CloudVOTE uses an efficient, automated process - Azure Traffic Manager for redirecting traffic to services or deployments in another region. Azure Traffic Manager automatically manages the failover of user traffic to another region in case the primary region fails.

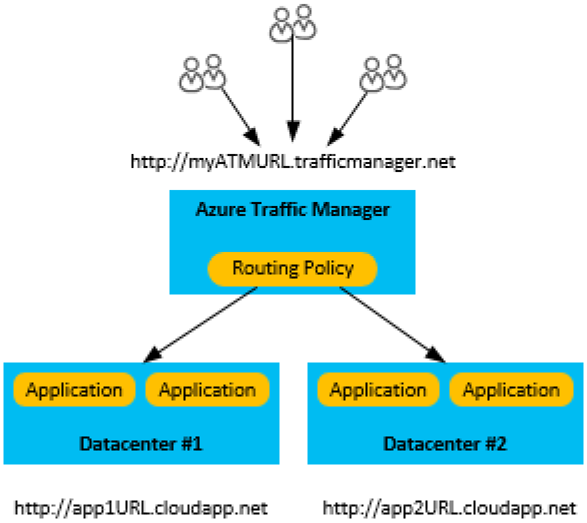

CloudVOTE Disaster Scenarios

The following sections cover several different types of disaster scenarios. Region-wide service disruptions are not the only cause of application-wide failures.

CloudVOTE is designed to consider the possible causes of failure and takes advantage of available Azure features and augments them with application-specific strategies. The chosen resiliency and backup solution is chosen based on your application's needs and objectives for business continuity.

Application Failure

Azure Traffic Manager automatically handles failures that result from the underlying hardware or operating system software in the host virtual machine. Azure Traffic Manager redirects traffic to a new instance of the role on a functioning server and adds it to the load-balancer rotation. If the number of role instances is greater than one, Azure shifts processing to the other running role instances while replacing the failed node.

Data Corruption

CloudVOTE automatically stores Azure SQL Database and Azure Storage data three times redundantly within different fault domains in the same region. Because of geo-replication, the data is stored three *additional* times in a different region. However, if your users or your application corrupts that data in the primary copy, the data quickly replicates to the other copies. Unfortunately, this results in three copies of corrupt data.

To manage potential corruption of your data, CloudVOTE offers two options. First, you can employ a custom backup strategy. CloudVOTE can store your backups in Azure or you can opt for on premise storage, depending on your business requirements or governance regulations. Another option is to use the new point-in-time restore option for recovering a SQL database.

Network Outage

When parts of the CloudVOTE network are inaccessible, you might not be able to get to your application or data. If one or more role instances are unavailable due to network issues, CloudVOTE uses the remaining available instances to run your application. If CloudVOTE can't access its data because of an outage, you can potentially run in degraded mode locally by using cached data.

CloudVOTE can store data in an alternate location until connectivity is restored. If degraded mode is not an option, the remaining options are application downtime or failover to an alternate region.

Failure of a Dependent Service

Azure provides many services that can experience periodic downtime. CloudVOTE considers what happens in your application if the dependent service is unavailable. In many ways, this scenario is similar to the network outage scenario.

Azure Redis Cache provides CloudVOTE with an option of local caching of your application data, which provides disaster recovery benefits. This helps to preserve cached data if a single node fails by maintaining duplicate copies on other nodes.

Some of the disadvantages of this approach are that high availability decreases throughput and increases latency because of the updating of the secondary copy on writes. It also doubles the amount of memory that's used for each item.

Region-Wide Service Disruption

The previous failures have primarily been failures that can be managed within the same CloudVOTE region deployment. However, CloudVOTE is also prepared for the possibility that there is a service disruption of the entire region. If a region-wide CloudVOTE disruption occurs, the locally redundant copies of your data are not available. With geo-replication, there are three additional copies of your CloudVOTE service in a different region. After the region is declared lost, CloudVOTE remaps all of the DNS entries to the geo-replicated region.

Azure-Wide Service Disruption

In disaster planning, you must consider the entire range of possible disasters. One of the most severe service disruptions would involve all CloudVOTE regions simultaneously. This will almost certainly cause a temporary downtime. Widespread service disruptions that span regions should be much rarer than isolated service disruptions that involve dependent services or single regions.

Data Strategies for Disaster Recovery

CloudVOTE automatically maintains a copy of your application's data. CloudVOTE uses this data for reference and transactional purposes at a secondary site. This allows CloudVOTE to readily deploy applications to multiple regions. These regions are geographically distributed in such a way that multiple-region service disruption should be extremely rare.

BACKUP AND RESTORE

Regular backups of application data can support some disaster recovery scenarios. Different storage resources require different techniques.

Depending on your specific CloudVOTE SQL Database tier, you can take advantage of point-in-time restore to recover your database. Another option is to use Active Geo-Replication for SQL Database. This automatically replicates database changes to secondary databases in the same Azure region or even in a different Azure region. This provides a potential alternative to some of the more manual data synchronization techniques.

CloudVOTE also supports additional, manual backup schemes, which are available upon request.

The built-in redundancy of CloudVOTE Azure Storage creates two replicas of the backup file in the same region.

REFERENCE DATA PATTERN FOR DISASTER RECOVERY

Reference data is read-only data that supports application functionality. Although CloudVOTE employs backup and restore as one of the methods to handle region-wide service disruptions, the Recovery Time Objective (RTO) is relatively long.

With CloudVOTE, you can improve the RTO by maintaining a permanent copy of the reference data in the secondary region. This eliminates the time required to restore backups in the event of a disaster. To meet the multiple-region disaster recovery requirements, CloudVOTE can deploy the application and the reference data together in multiple regions. CloudVOTE can deploy reference data to the role itself, to external storage, or to a combination of both.

The reference data deployment model within compute nodes implicitly satisfies the disaster recovery requirements. CloudVOTE also offers reference data deployment to SQL Database, which requires deployment of a copy of the reference data to each region. The same strategy applies to CloudVOTE Azure Storage.

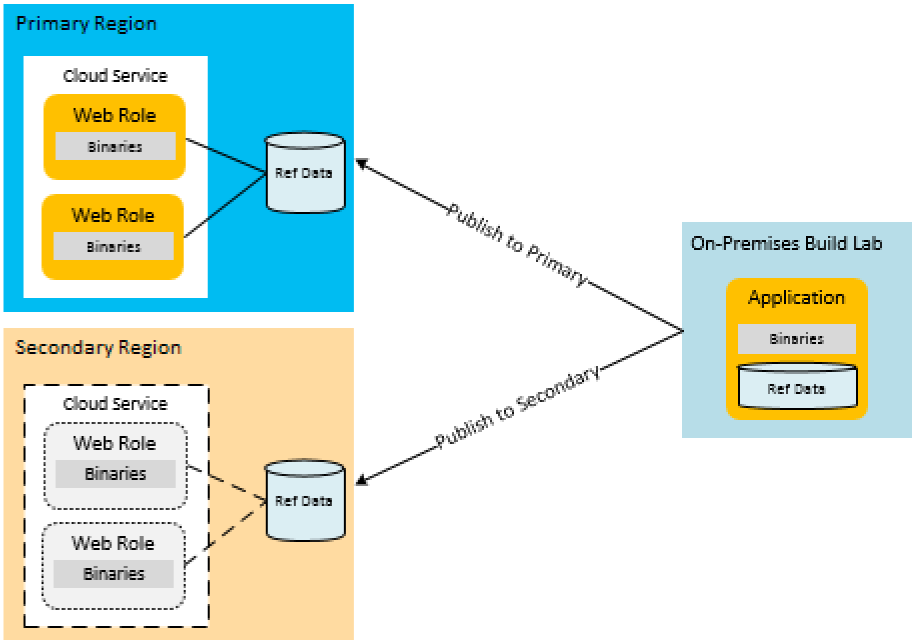

TRANSACTIONAL DATA PATTERN FOR DISASTER RECOVERY

Implementation of a fully functional disaster mode strategy requires asynchronous replication of the transactional data to the secondary region.

CloudVOTE offers different ways of handling transactional data in a failover scenario. For example, intermediate storage locations might be stored within an Azure SQL Database while the queues themselves might be either CloudVOTE Azure Storage or CloudVOTE Azure Service Bus queues. CloudVOTE server storage destinations might also vary, such as Azure tables instead of SQL Database.

Deployment Topologies for Disaster Recovery

CloudVOTE is prepared to handle your mission-critical applications in the possibility of a region-wide service disruption by incorporating a multiple-region deployment strategy into its design.

Multiple-region deployments involve processes that publish the application and reference data to the secondary region after a disaster. In case of an instant failover, the deployment process might involve an active/passive setup or an active/active setup. Such CloudVOTE deployment has existing instances of the application running in parallel in the alternate region. Azure Traffic Manager provides load-balancing and traffic routing services at the DNS level that detect service disruptions and route the users to different regions when needed.

CloudVOTE can dynamically and quickly allocate resources to a different region, thus saving you money on idle resources while waiting for the next failure to occur.

The following sections cover different CloudVOTE deployment topologies for disaster recovery.

SINGLE-REGION DEPLOYMENT

A single-region deployment is meant to provide contrast with the other architectures. CloudVOTE does not employ single-region deployments for its clients, because it does not provide any form of disaster recovery plan.

We are using the following diagram to depict an application running in a single Azure region and to demonstrate the low resiliency and insufficient data recovery capabilities.

Single-Region Deployment

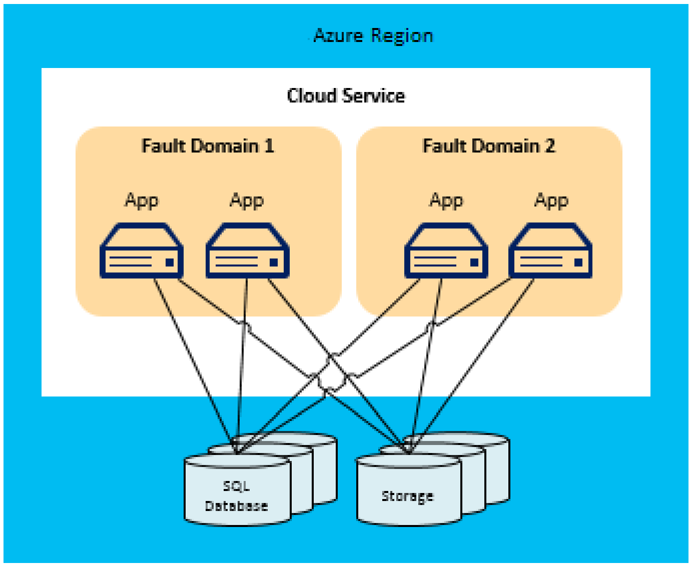

Clearly, the database is a single point of failure. Even though CloudVOTE replicates the data across different fault domains to internal replicas, this all occurs in the same region. Such application would not withstand a catastrophic failure. If the region goes down, all of the fault domains go down--including all service instances and storage resources.

RTO in this scenario may be up to 24 hours, which is only suitable for non-critical applications.

ACTIVE-PASSIVE

The active-passive pattern is the choice that many CloudVOTE clients favor. This pattern provides a great balance between improvements to the RTO and relatively small increase in cost. In this scenario, there is a primary and a secondary CloudVOTE Azure region. All of the traffic goes to the active deployment on the primary region. The secondary region is better prepared for disaster recovery because the database is running on both regions. Additionally, a synchronization mechanism is in place between them. This standby approach can involve two variations: a database-only approach or a complete deployment in the secondary region.

Active/Passive (Database Only)

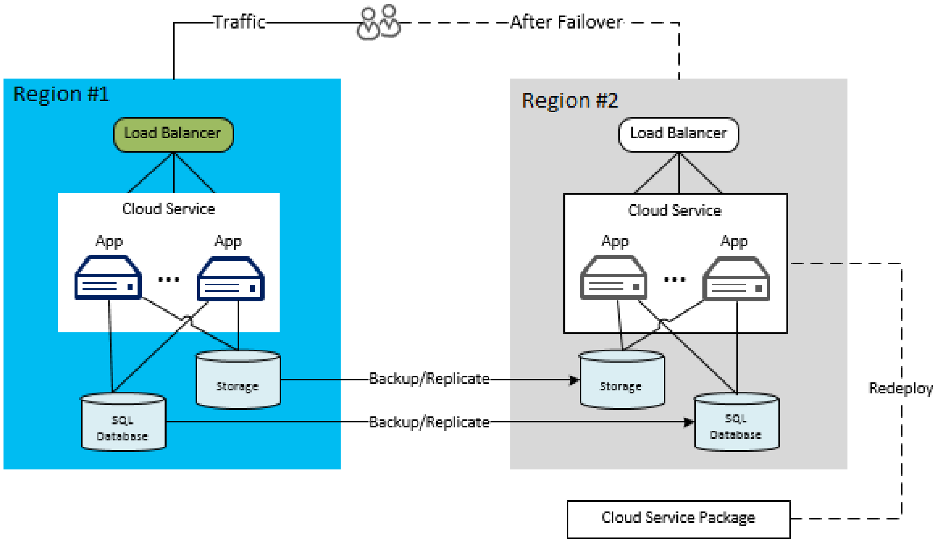

Active/Passive (Full Replica)

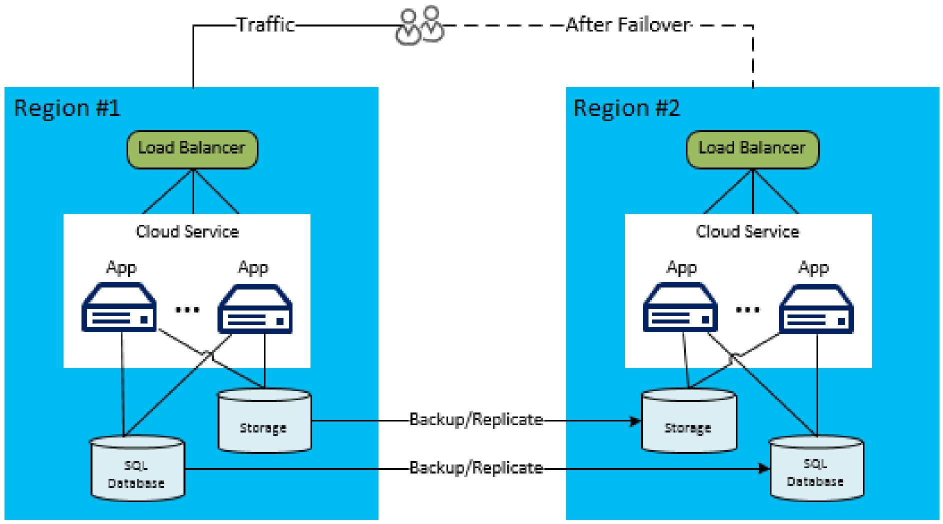

Active-Active

Decreasing the RTO increases costs and complexity. The active-active solution actually breaks this tendency with regard to cost.

In an active-active pattern, CloudVOTE services and database are fully deployed, running and receiving user traffic in parallel, in both regions.

This option yields the quickest recovery time. The services are already scaled to handle a portion of the load at each region. CloudVOTE Azure DNS services are enabled to use the secondary region. In case of failover, DNS change will route all traffic to the secondary region.

Active/Active

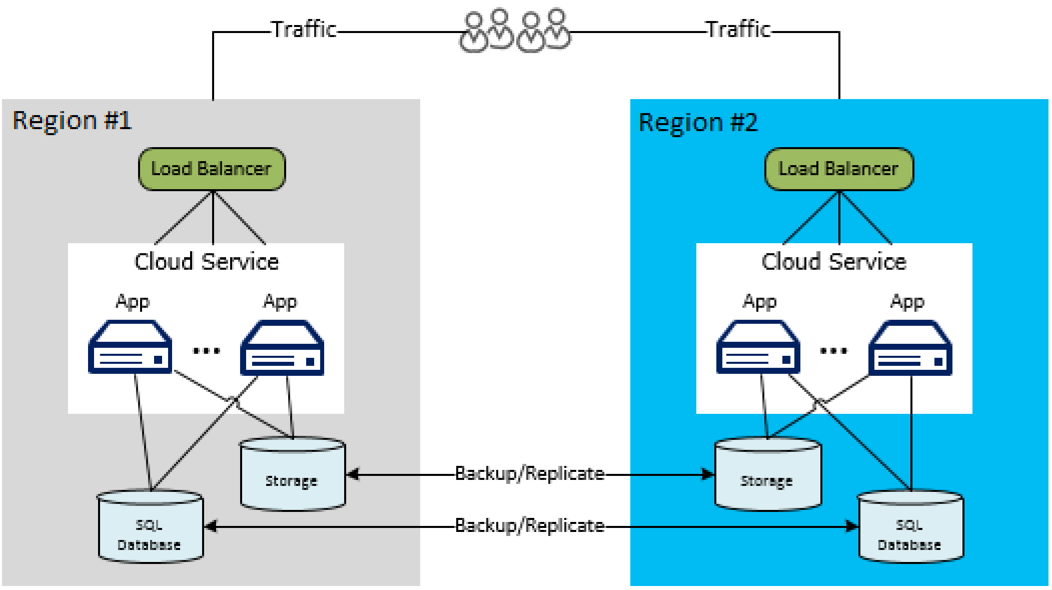

Automation

CloudVOTE handles your deployment's backup and recovery scenarios automatically. You and your team will be notified in case of emergency via our standard Disaster Recovery Plan and procedures.

Failure Detection and Monitoring

CloudVOTE is able to detect and diagnose failures automatically and manually. CloudVOTE uses all server and deployment monitoring tools available in Azure, so we can quickly know when and where a system error occurred and can remediate the situation.

CloudVOTE provides timely, informative and comprehensive reporting to our clients as defined in our standard CloudVOTE SaaS Agreement.

Disaster Simulation

CloudVOTE platform undergoes regular simulation testing, involving creating small real-life situations to observe how system, people and processes work together in disastrous situations. Simulations also show how effective the solutions and processes are in the recovery plan.

Every such simulated test is specific to a CloudVOTE customer and results are shared with the responsible parties.

Final Solution

Final CloudVOTE solution will depend on the level of resiliency your organization needs and is able to justify the cost for. It will also depend on how adverse your organization and business requirements are to downtime, outage, or loss of data.

Give us a call today to discuss how CloudVOTE can help you make your training and education better, more efficient and economical.