More info: www.agya.cloud

Cloud Continuity
How it Works

// Product Overview

Description

Cloud Continuity is a Business Continuity/Disaster Recovery (BC/DR) solution for companies seeking to protect their mission-critical servers, applications, and data. It is a real-time, Cloud-based server replication service offering continuous uptime in the event of a server or site failure.

Cloud Continuity is a subscription-based, managed service that eliminates the need for the end customer to purchase, install, oversee and support any of: hardware, operating systems, application software, BC/DR software and DR data center facilities, thereby helping customers eliminate upfront capital expenses, excessive deployment time and demands on IT resources.

Cloud Continuity helps businesses meet their continuous availability needs easily within the limits of their staff and budget, ensuring business and compliance accountability despite a challenging economic environment.

How It Works - Overview

All data from enabled servers is continuously replicated in real-time to an off-site Cloud-based data center. In the event of a server or system failure of a customer production server, businesses have the option of automatically or manually redirecting end-user traffic to the replica environment in the Cloud (a process called “Failover”), allowing end users to continue working, without interruption, for as long as is required to fix the initial problem. Once the production site issue is addressed, any and all new data is seamlessly synchronized with the production server while users continue to operate from the Cloud. When synchronization is complete, traffic is redirected to the production environment (a process called “Failback”).

// Replication

In the context of the Cloud Continuity offering, “Replication” means copying all the application data (but not the applications) from one server (primary, or Production Server) to another server (Replica Server). After the initial bulk data set-up, only data changes made to the Production Server are copied to the Replica Server, ensuring fast and efficient network throughput and minimal impact to server performance. Changes are copied to the Replica Server in real-time, ensuring that all data is up-to-date at all times.

Data Synchronization

The service commences with an initial synchronization of all databases, files, registry keys, etc. from each of the Production Servers as defined during the setup stage in the Readiness Assessment (RA). The service will be actively protecting the Production Server once the initial data synchronization has completed and replication mode has started.
Data synchronization not only occurs during the initial setup, but also after each time the replication has been stopped or interrupted for a long period of time. This ensures that any changes that might have been made to the Production Server when replication was not running are subsequently recognized. If data synchronization is interrupted before completion, then it will need to start over again the next time.
The data synchronization stage requires more resources than the ongoing replication mode as data comparison occurs at the block level between Production and Replica for the entire data set to be replicated. Once replication mode has been reached the CPU resource requirements for Cloud Continuity are minimal.

Replication Mode

Once the initial synchronization is completed, the service is now in replication mode, and a change initiated data replication process is used to ensure that any data changes made to the Production Server are also made to the Replica Server in the Data Center. The default change detection uses asynchronous, real-time, block-level data replication.

Data Spool

If the change rate of the data is too great for the connection between the customer and the replica site to send data in real-time, then data changes are spooled on the customer’s Production Server. If there is less than 5 GB free on the Production Server then the replication of data will stop and the spooled data changes will be flushed to prevent any disruption to the Production Server applications or use. When replication is restarted through the Cloud Management Platform, a resynchronization of data will start. The time to complete this resynchronization depends on the volume of data it needs to compare between the production server and the Replica Server.

Bandwidth Considerations

Replication throughput is multi-streamed and makes use of all available bandwidth by default (an option to throttle the speed can be used). In an attempt to prevent bandwidth constraint issues, customers should be set up with a minimum of 10Mbps available bandwidth to start, and then re-evaluated with respect to the ongoing data change rate once the service is running. Once data is replicating, the average daily change rate should not exceed what the bandwidth is capable of handling within any given 24-hour period to ensure the Replica Server is up to date and not trailing behind.

// Failover & Failback

Failover Process

In the event of a production server failure, Frontier’s solution can automatically detect the outage on supported applications and a failover process can be initiated whereby traffic is immediately redirected to the replica environment where users continue to operate of off the Cloud environment for as long as required to resolve the initial problem. For 3rd party applications, the Cloud Management Platform can be used to initiate the application

failover. This process is called 'failover' and takes about a minute. There is no interruption to service, and no loss of data, as the replica environment is continuously being updated in real-time with all changes occurring in the Production Server.

Failback Process

Once the initial issue with the Production Server(s) has been addressed and the on-site server is live again, any new data generated since the initial failover is first synchronized from the Replica Server back to the production server. While synchronizing occurs, the Replica Server is still available and used by all users. Once synchronization is complete, traffic is redirected to the Production Server again. This process is called 'failback' and like the failover process, it takes about a minute to complete, with no loss of service and no loss of data.

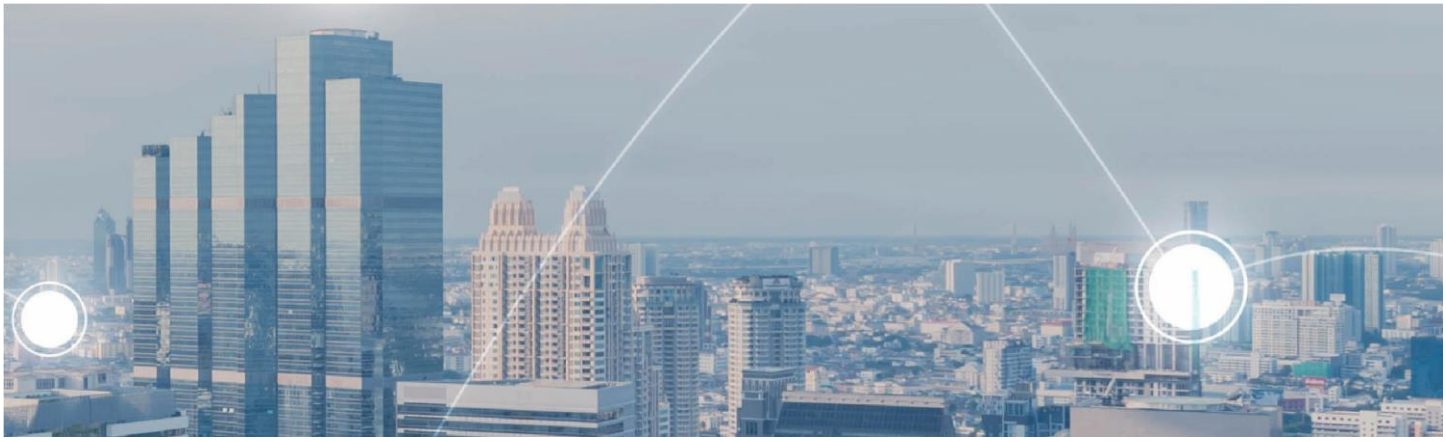

Setup Notes // Specific Setup Notes for Exchange 2013

When configuring replication services for Exchange 2013 and running in a configuration where the Client Access Server (CAS) and mailbox roles are separated, a Network Load Balancing (NLB) cluster is configured for the CAS so Outlook clients can support redirection when a failover occurs. The mailbox server that is hosting email data is replicated in real-time to the Replica Server in the remote Cloud-based Data Center. On failover, Active Directory is updated automatically so all user profiles mailboxes are set to be homed on the Replica Server. Outlook, when connecting to the CAS, is redirected to the replica automatically for a seamless failover for all users.

Externally, a secondary MX record for the email domains would be published that points to the Replica Server Cloud hosted external IP so that new emails are received.

// Specific Setup Notes for SQL

When configuring SQL replication services, a SQL replica is installed with the same version and data locations as the production. By default, the production databases and any user databases are replicated in real-time to the replica. Any changes, such as new databases, will be included automatically within the replication profile as long as they are in drive locations that are present on the replica DR server. SQL account logins and policies will also be propagated to the replicas automatically as part of the replication.

On failover, the Active Directory is updated automatically so that any database connection will be redirected to the Replica Server. The databases will be mounted automatically and ready to be used on the Replica Server.

During regular operation, any data changes in the databases would result in data being replicated. Cloud Continuity supports SQL Clustering and provides data replication from the active member node quorum and continues to replicate upon an active node switch.

// Specific Setup Notes for 3rd Party Applications

3rd party applications are defined as applications not specifically listed as a “Supported Application” for the Frontier Cloud Continuity service. These applications require the licensing, installation, and management to be the responsibility of the customer or Partner, and thereby impact the standard delivery of the Cloud Continuity service.

Servers running 3rd party applications will be protected with real-time replication of data between the Production and Replica Servers. In these cases, the failover process will result in the data replication stopping and Active Directory being updated so that users can continue to access the server by name, as they will be directed to the replica.

While Frontier will be responsible for the deployment of the replication server, the customer is responsible for the configuration, installation, and any licensing during the appropriate provisioning phase, as communicated by Frontier. Ongoing monitoring and maintenance of the Applications will be the responsibility of the customer. Alternatively, the solution can be delivered as a fully integrated service for additional fees.

// Specific Setup Notes for SAP Deployments

SAP deployments are supported with Cloud Continuity with a combination of replicated and standalone Replica Server(s). Successful deployments combine the replication of multiple back-end data systems including databases, reporting servers, and archived file data to the Replica Server(s). In addition to the replicated servers there may be additional stand-alone SAP application servers in the DR environment pre-configured to a disaster occurring that will be used in the event of a disaster.

Due to the network latency constraints many SAP applications have, we suggest above-minimal throughput for replicating data, and we suggest not just failing over a single server but rather all servers associated with the SAP application in the event of a disaster or outage.

Many SAP applications perform regular re-indexing of databases so that change to the data on the production system should be taken into account when estimating time to synchronize and data spool sizes.

// Specific Setup Notes for Oracle Deployments

When configuring Oracle replication services, an Oracle replica is installed with the same application version, Oracle SID, and data locations as the production. By default, the production databases and any user databases are replicated in real-time to the replica. Any changes, such as new databases, will be included automatically within the replication profile as long as they are in drive locations that are present on the replica DR server. Oracle account logins and policies will also be propagated to the replicas automatically as part of the replication.

On failover, the Active Directory is updated automatically so that any database connection will be redirected to the Replica Server. The databases will be mounted automatically and ready to be used on the Replica Server. During regular operation, any data changes in the databases would result in data being replicated.

Cloud Continuity does not support configurations already running Oracle RAC.

// Specific Setup Notes for Active Directory

Description

When Active Directory is present, an administrator account with Domain Admin privileges must be created for the Cloud Continuity service that is a member of the domain that it will be protecting. This provides the account with the necessary access within the domain to facilitate the tasks required to perform failover and failback, and allows the service to make the appropriate service changes and to access the data to replicate. If servers are in a Workgroup or are Linux/Unix, then a Local Administrator account must be created on each of the Production servers.

// Solution Components

At the End Customer Site

Production Servers (physical or virtual as defined by the end customer for replication), Router or Firewall supporting a Standard IPSec VPN tunnel, Internet connectivity.

At the Hosting Data Center

Production Servers (physical or virtual as defined by the end customer for replication), Router or Firewall supporting a Standard IPSec VPN tunnel, Internet connectivity.

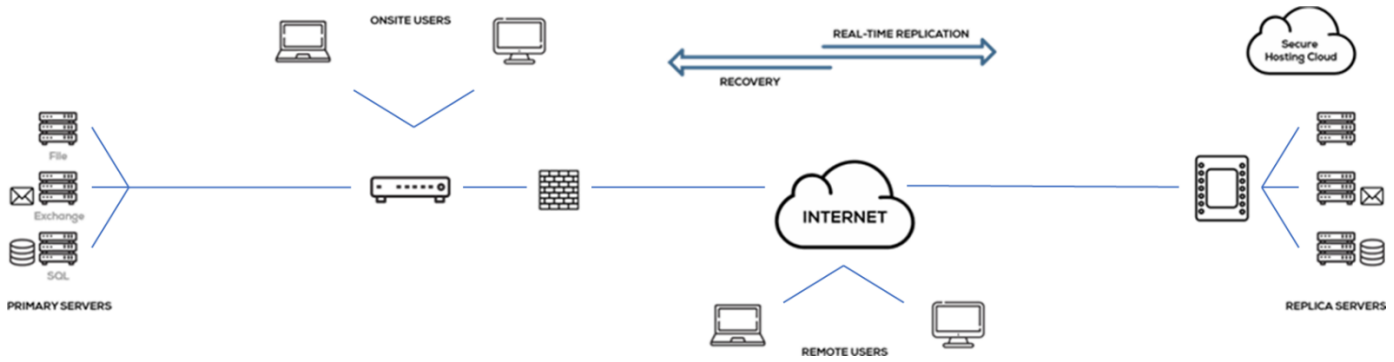

// Management Interface

Management of the entire solution is possible through an integrated web-based user interface: the Cloud Management Portal. Customers and Partners have different login access and privileges allowing complete management in a top down model.

Cloud Continuity

The Control Panel is used by end customers to check the status of their replicated environment, test failover processes to ensure that they are protected, and to manage the replication environment and servers. From anywhere at any time, end-customers can instantly initiate a failover and failback with the click of a mouse, without engaging their supplier. Support teams may also have access to the Control Panel to initiate operational requests on a customer's behalf, and to initiate failovers and failbacks when needed.

Virtual NOC

The Virtual NOC is used to provide detailed information about the devices within the Production and Replica networks.

Metering and Service Health

The Cloud Management Portal allows for the metering and monitoring of devices and services to ensure your environment is protected and ready to be used in the event of a disaster or outage.