catalyst cloud

White paper

High availability on the Catalyst Cloud

Features and techniques to improve availability, resiliency and business continuity of web applications hosted on the Catalyst Cloud

3 February 2016

catalyst

open source technologists

Level 6, Catalyst House, 150 Willis Street
PO Box 11053, Manners Street, Wellington 6142, New Zealand
+64 4 499 2267 // info@catalyst-nz.net // www.catalyst-nz.net

Table of Contents

1. Introduction

This document outlines the physical infrastructure and software features that make the Catalyst Cloud highly available and resilient. It covers built in features that are inherited by every project and services that can be used to enhance the availability of web applications and web sites hosted on the Catalyst Cloud.

2. Physical Infrastructure

The Catalyst Cloud is hosted on multiple regions or geographical locations. Regions are completely independent and isolated from each other, providing fault tolerance and geographic diversity. This section describes the physical attributes of our data centres.

Security

Catalyst Cloud regions have the following security attributes:

- Comprehensive access control systems with multiple perimeters
- Closed Circuit TV with cameras covering entrance and data centre
- Fire suppression mechanisms

Power and Cooling

Catalyst Cloud regions have the following power supply and cooling attributes:

- High Capacity UPS
- Backup Diesel Generators
- A+B power to UPS and racks (Porirua and Hamilton)
- N+1 process coolers

Network

The physical networks backing the Catalyst Cloud have the following properties:

- Diverse Fibre paths linking cloud regions
- Diverse Fibre providers for fibre paths
- Diverse ISPs
- Diverse International Upstream Providers
- High capacity intra cloud bandwidth

3. Catalyst Cloud features that enhance availability and resiliency

Compute

If a physical compute node fails, our monitoring systems will detect the failure and trigger an "evacuate" process that will restart all affected virtual compute instances on a healthy physical server. This process usually takes between 5 to 20 minutes which allows us to meet our 99.9% availability SLA for individual compute instances.

Customers that require more than 99.9% availability can combine multiple compute instances within the same region using anti affinity groups. Anti affinity groups ensure that compute instances belonging to the same group are hosted on different physical servers. This reduces the risk and probability of multiple compute instances failing at the same time due to the loss of a single server.

Our documentation provides instructions on how to deploy highly available compute instances using anti-affinity filters and keepalived: http://docs.catalystcloud.io/tutorials/deploying-highly-available-instances-with-keepalived.html

Customers that require their applications to survive the loss of an entire data centre can launch compute instances in different regions (assuming your application supports operating in this way). The following tutorial explains how to use the Fastly CDN as a fail-over mechanism between regions: http://docs.catalystcloud.io/tutorials/region-failover-using-the-fastly-cdn.html

Block Storage

We run a distributed storage system that by default retains three copies of your data on different servers spread across a region (a datacenter). We can afford to lose many disks and multiple storage nodes without losing any data. As soon as a disk or storage node fails, our storage solution begins recovering the data from an existing copy, always ensuring that three replicas are present.

The storage solution is self managing and self healing and places your data in optimal locations for data survival and resiliency. It runs automated error checks in the background that can detect and recover a single bit of data wrong, by comparing the three copies of the data and ensuring they are identical.

The solution is designed and implemented with very high availability and data resiliency in mind. It has no single points of failure.

Object Storage

Currently Object Storage is backed by the same technology and physical infrastructure as block storage. Consequently it offers the same protections and guarantees as Block storage does.

In the first half of 2016 we will start replicating object storage data across three

regions, providing 99.999999999% durability and 99.99% availability for data over a given year.

Virtual Routers

In the same way that if a compute instance fails, if a physical network node fails our monitoring systems will detect the failure and trigger the evacuate process that will ensure all affected virtual router instances are restarted on a healthy server. This process usually takes between 5 to 20 minutes.

We are working on a new feature that launches two virtual routers on separate network nodes responding on the same IP address. Once this is complete the failover between routers will take milliseconds which will most likely not be noticed. Meanwhile customers requiring higher availability are advised to combine compute instances from multiple regions where possible.

Monitoring

The catalyst cloud has robust fine grained monitoring systems in place. These systems are monitored 24x7.

Roadmap

The following features are expected to be available on the Catalyst Cloud in 2016 and will improve resiliency and availability even further:

- Highly Available redundant routers (using VRRP)
- Object storage replicated across three regions
- Load Balance as a Service (LBaaS)