PRIVATE CLOUD

POWERED BY OPENSTACK

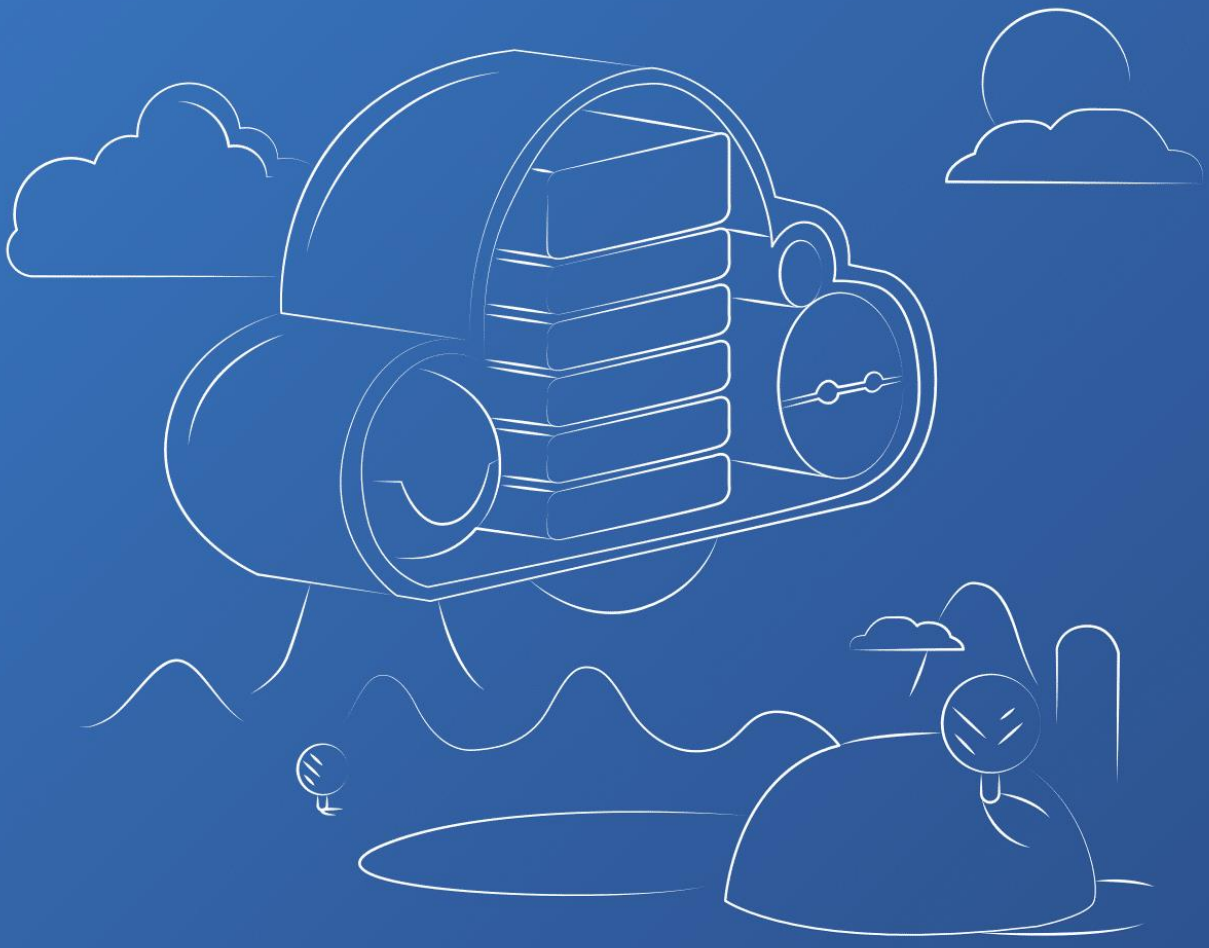

Globo.Tech | © GloboTech Communications
sales@globo.tech

EXECUTIVE SUMMARY

The way people consume content has drastically changed over the last decade. People now watch television, read news, listen to music, and share content, over the Internet. Internet adoption opened doors for new business opportunity. Platforms like Netflix, Spotify, Twitter, and Facebook are examples of emerging Software as a Service, accessible over the Web. Cloud computing is one of the main drivers of Software as a Service, providing the infrastructure and the computing power over the Internet; giving businesses the resources and power to tap this opportunity, while keeping their IT budget under control. Cloud computing as recently transformed how business is conducted, delivering better performance, greater flexibility, and greater elasticity to their IT infrastructure. There are 2 main types of cloud technologies, delivering platform as a service; Public Cloud and Private Cloud. Public Cloud delivers computing power to a group of users sharing the same pool of server's resources, while the private cloud is a proprietary environment.

OpenStack is currently the most popular[1] open source software stack for creating and deploying cloud computing solutions. The stack consists of multiple projects/services allowing businesses to control their computing resources, network, and storage. With OpenStack, you can manage large pool of servers through a web-based dashboard, or using the OpenStack API. This document aims to explain in simple terms the OpenStack solution: its advantages, services, and infrastructure.

Find out why thousands of companies trust OpenStack for their private cloud environments.

OVERVIEW

The OpenStack initiative was incepted in 2010. Its mission was, and still is, "to produce the ubiquitous Open Source Cloud Computing platform that will meet the needs of public and private clouds regardless of size, by being simple to implement and massively scalable". Since then, OpenStack his used by hundreds of leaders in all industries, such as BMW, Volkswagen, BestBuy, and more, to power and manage their hosting infrastructure.

RELEASE HISTORY [3]

Over the years, OpenStack has evolved rapidly to become the most popular open source software stack for cloud computing. Since the launch of the OpenStack Foundation, in 2012, OpenStack continued to grow as more Individuals and corporations started investing in the project, discovering more flaws, and rapidly filling the gaps. OpenStack has official releases every six months, usually in October and April. Here's how OpenStack evolved around the years to become what it is today.

WHY OPENSTACK

OpenStack promote a collaborative approach to cloud computing. With its open-source model, developers around the world can work together to improve services delivered by the foundation. In addition, corporations with the means can adapt one or many of the OpenStack projects to improve their business operations. OpenStack incorporate several services, core and optional, that can be combined to better suit certain type of use, such as big data, web applications, video processing, web hosting, public cloud, and more. With its active community, free to use, open source approach, and its 7 years of developments; OpenStack is a convenient solution for most uses.

Active Community

Enjoy fast growing open source communities, with thousands of active developers working on multiple projects to deliver the best solutions for virtualization.

Increased Flexibility

With its modular services, you can customize the solution to your liking, allowing you to get the perfect platform to fulfill your needs.

Open-Source

OpenStack is free to use, distribute and modify, which means it's usually cheaper than its proprietary counterpart.

Open API

OpenStack open API can use to launch instances, resize them to your liking and more. Many third-party service providers support OpenStack open API and it's compatible with AWS EC2 API.

In addition to the benefits provided by any cloud computing service, OpenStack offers several advantages over the competition. First, it is free for download, and supported by several notable vendors, including Red Hat, Canonical (Ubuntu) and many others. Furthermore, because it is an open source project with ongoing contributions, it is always improving and offering new products to users.

OPENSTACK SERVICES

OpenStack incorporate various projects, each designed for specific use. The core services of OpenStack manage the compute, network, and storage of your server. These services allow virtualization, the process of creating Virtual Machines by dividing resources from your servers. These services can then be combined with optional services to automate provisioning of virtual machines, manage DNS, manage databases, manage containers and more.

PROJECT OVERVIEW

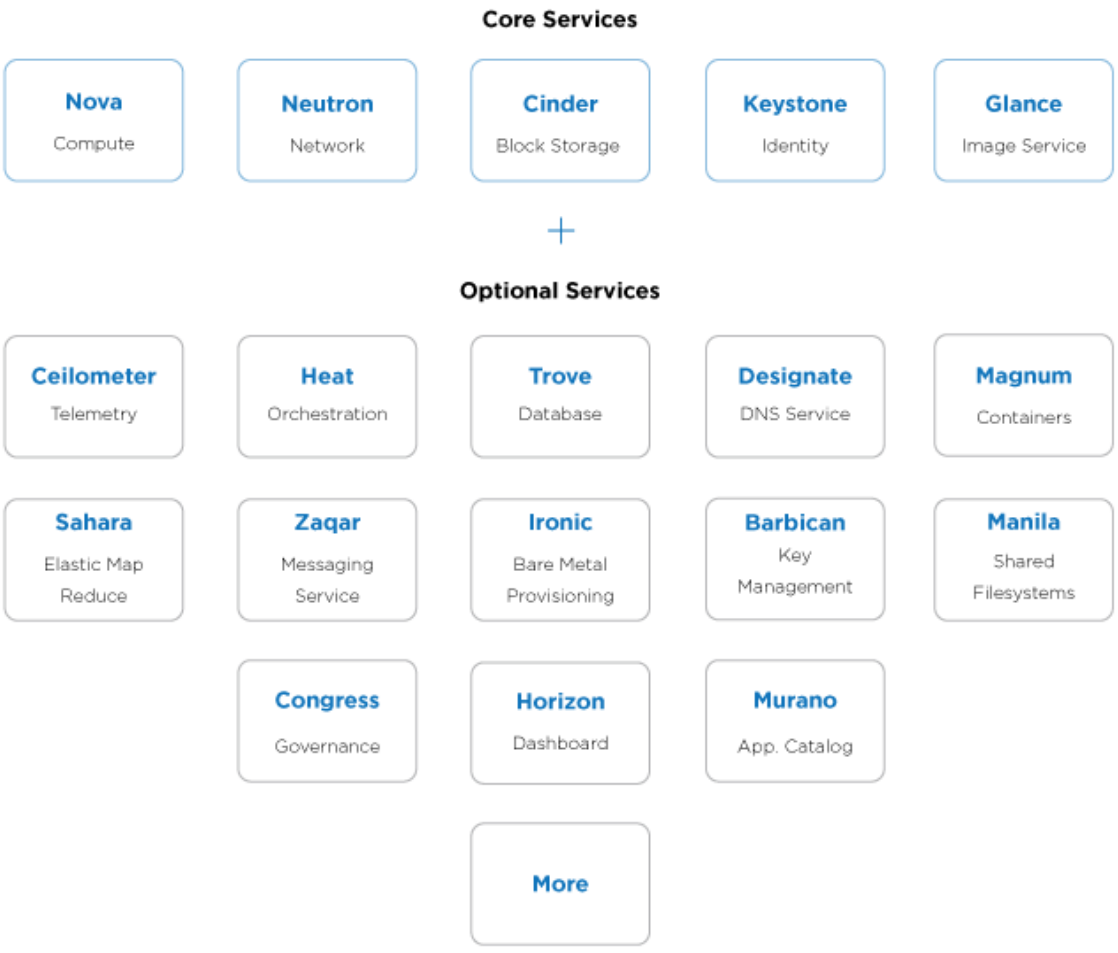

To get more information about each of the individual OpenStack projects, and their level of maturity, visit the [OpenStack Project Navigator]

CORE SERVICES

Compute (code name Nova [4]) – offers virtual servers to meet the demands of the user. Compute is one of the integral components in an OpenStack cloud, and it is compatible with most third-party technologies.

Network (code name Neutron [5]) – offers Network as a service, allowing you to design and build your own network. Different networking models are available, and you can customize the system to meet your needs with additional services such as VPNs, firewalls and more.

Block Storage (code name Cinder [6]) – provides guest virtual machines with block storage services. This system can leverage many different storage platforms by connecting to them using plug-ins.

Identity (code name Keystone) – authorizes and authenticates all other OpenStack Services. Identity is compatible with multiple types of authentication, and it maintains a comprehensive directory of all services deployed within the cloud.

Image (code name Glance [7]) – is a component that stores and organizes virtual disk images. Images can also be used to store backups. OpenStack Image can be used with simple file systems, as well as with other components of OpenStack, such as Object Store.

Object Store (code name Swift) – offers distributed file storage and retrieval. Because storage and processing are distributed, this system is both flexible and reliable. Object Store is highly scalable and can handle significant amounts of data. Swift can be used regardless of Openstack.

Other Services – Additional third party services for OpenStack operations such as databases (MySQL/MariaDB/etc.) and queue services (RabbitMQ/etc.) are required for all OpenStack services to work properly.

INFRASTRUCTURE [2]

The choice of your infrastructure is relative to the amount of resources required by your business. Since private cloud solutions are completely modular, it is possible to customize the solution according to the specific needs of your projects/business. The factors influencing the choice of the required equipment are related to; the number of virtual machines required, the need for clustering, the number of active users, the level of security required, the need for high availability, the bandwidth required by your platform, etc.

VIRTUALIZATION ARCHITECTURE

A good practice in designing your infrastructure is to have your OpenStack private cloud separated in different layers: Controller, Computing, Networking, and Storage

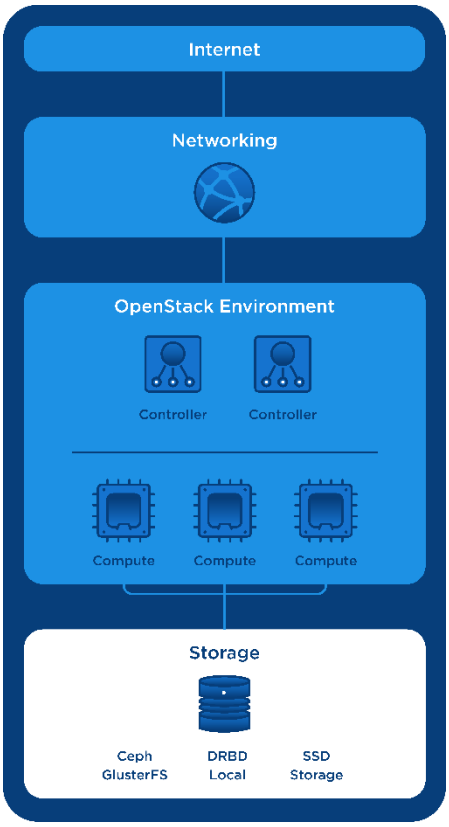

CONTROLLER

In the preferred model, the controller node holds all OpenStack services and third-party services required to orchestrate Virtual Machines provisioned on the Compute Nodes. This node supplies APIs, dashboard, storage management, network management and identity management to your private cloud. If the controller node becomes inaccessible, only the OpenStack services will be affected. Your Virtual Machines will still run, but it won't be possible to resize, close or modify them in ways. It is highly recommended to have two controllers with failover capabilities for a production environment. A 3-controller model is also possible and a 1-controller model can be acceptable for pre-production of for a Proof of Concept.

COMPUTE

A compute node is a server system used to provision and host Virtual Machines (VMs). Each compute nodes runs their own operating system and connects to the controller using network connectivity. It's good practice to separate the compute from the controller, not only to avoid availability issues, but also to improve computing power. Each compute nodes are independent from each other, meaning that when one compute is down, it won't affect the other computes, but the VMs running on the faulty compute will stop functioning. Nova is agnostic and fully compatible with most popular hypervisors, such as KVM, XEN, Hyper-V and VMWare.

NETWORK

There are several methods to design your network within OpenStack. You can use OpenStack's native solution: Neutron or third-party solutions; Hardware, such as Cisco Nexus, Arista or other network or Software (Software Defined Networks), such as (open)Contrail, OpenDayLight, etc. There are two preferred types of networking design natively available with Neutron, FLAT network, and DVR network. Both designs allow for your compute node to handle the routing locally. Small businesses usually use FLAT networks, since the approach is cost effective and requires low maintenance. A FLAT network is the process of connecting multiple computers or servers to a single switch. It is directly managed in Linux Kernel, using Open vSwitch or Linux Bridge. It's a straight forward and proven method but limits what a user can do in terms of features and performances since its using Linux kernel. The main issues with this approach are the lack of redundancy and the limited scalability. In opposition, Distributed Virtual Routing (DVR) adds distribution to your routing by attributing the same gateway IP/MAC address across all your VMs, using a virtual router (qRouter) for the public network. Neutron DVR is an OpenStack project, maintained by the neutron community, that serves the purpose of distributing virtual routers through the compute nodes, isolating them in case unavailability from a node (the other compute nodes won't be affected), thereby mitigating the risk of unavailability at the network level.

STORAGE

There are 2 types of storage that can be used in a private cloud infrastructure; local storage and centralized storage. Local storage is just like the storage you have on a single server. If the server with the local storage become unavailable, your Virtual Machines will be lost and/or temporarily unavailable, since the data is not available to the other compute. Centralized storage is storage accessible through the network and made available to all compute nodes via standard protocols, such as iSCSI, NFS, RBD (Ceph) or other third party storage solutions. Centralized storage is usually recommended over internal storage, since all servers can access the same pool of storage space; delivering scalability, manageability, but can introduce a higher level of latency to your infrastructure. OpenStack's virtualization projects must constantly communicate with a storage solution to function properly. In OpenStack, Nova will allocate local storage to your WMs, Glance will use your centralized storage for your VMs images and Cinder will be able to create block storage.

MANAGEMENT NODE

A management node is a node where the upkeep of your OpenStack will be complete. You can use it to setup a VPN to secure your connection to OpenStack Component and servers. Run your automation deployment such as Ansible (or your favorite automation tool) script to scale your OpenStack Deployment or Centralized and repatriate all OpenStack Component Logs within one repository using tools such as Kibana and Logstash.

INFRASTRUCTURE ARCHITECTURE

Infrastructure architecture is a turning point for the successful implementation of a private cloud. As previously stated, good practice in designing your infrastructure is to have your OpenStack private cloud separated in different layers: Controller, Computing, Networking, and Storage. We will now look more closely at 3 different examples of infrastructure, to understand their benefits and limitations.

Developer Infrastructure and Proof of Concept

When exploring new technologies, such as OpenStack, some users may need to rapidly deploy a Proof of Concept, to test different functionality. The OpenStack community developed a project called DevStack, which allows you to install all the basic components of Openstack on a single server, to deploy the solution and quickly be operational. This simple deployment model can also be used in a larger scheme in which your DevOps will have their own environment to work with OpenStack

Benefits	Limitations
-Low Cost -Fast Deployment -Easy to setup	-Limited Resources -Stability (Multiple single point of failure) -Internal Storage

Core Services: Nova, Neutron, Glance
Optional Services: Nagios, SNMP

Non-High-Availability Infrastructure

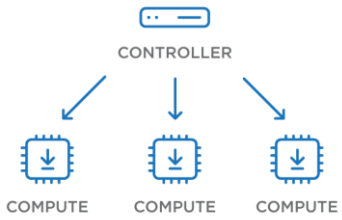

Some businesses may require a more stable infrastructure to mitigate the risk of unavailability in case of failure at the controller, storage, or network level. To mitigate such risk, it is recommended to separate your infrastructure in several layers. In this case, we separated the controller from the compute to make sure that, in the event of a failure from the controller, your compute node and your VMS will still be running, but OpenStack API will be degraded or unavailable. It's important to access which resources are critical for the well-being of your operations, to conceive such infrastructure tailored for your business.

Benefits	Limitations
-Affordable -Mitigating at least 1 single point of failure	-Limited Resources -Stability (Remaining single point of failure) -Internal Storage

Core Services: Nova, Neutron, Glance
Optional Services: Cinder, Ceilometer

High Availability Infrastructure

For production purpose, it's imperative that your infrastructure stay up at all time, since downtime usually means loss of profit. That's why high availability is key, since it provides redundancy at every level; controller and storage and it includes fail over capabilities. In case of failure from the controller or the storage, the backup node will take its place, and continue delivering your service. To eliminate every single point of failure, networking switch, network provider and power supply must also be redundant, with fail over capabilities. Usually these are taken care by your hosting provider. It's important to note that compute nodes fail over isn't a service provided by OpenStack. If this functionality is required for your project, VMware may be a more suitable solution.

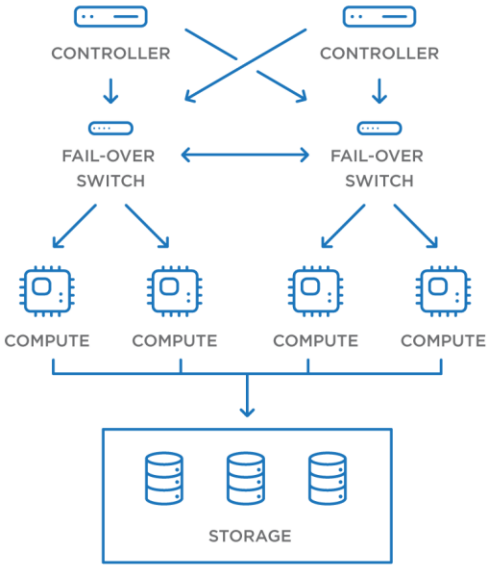

Benefits	Limitations
-Centralized Storage -High Availability -Increased Performance -	-Cost -Increased Infrastructure Complexity

Core Services: Nova, Neutron, Glance
Optional Services: Cinder, Ceilometer

Infrastructure Summary

Private cloud infrastructure are usually separated in different layers; controller, compute, network, storage, and management. Unless it's for a Proof of Concept, redundancy and high availability should be considered while designing your private cloud infrastructure. Virtualization is complex and requires a team of professionals to conceive, deploy and maintain the infrastructure.

OVERVIEW

At GloboTech, we know that your business relies on its online presence for success and we strive to provide you with the most reliable IT infrastructure to deliver it. Choosing, implementing, and maintaining a private cloud infrastructure can be an arduous process for some businesses. That's why we provide you with the help of a dedicated technical executive to guide you through the process. With 17 years of experience in the hosting industry and thousands of satisfied customers, you can rely on GloboTech Communications for your hosting needs.

COMMITMENTS & SERVICE LEVEL AGREEMENTS

Our prime focus, at GloboTech, is to build and strengthen long-term relationships with all our clients. Our servers are backed by one of the best SLAs in the hosting industry, with clear and straightforward guarantees.

- 99.9% Private Cloud Uptime
- 100% Power Uptime
- 1 Hour Hardware Replacement
- 100% Network Uptime
- 100% Infrastructure Uptime
- First support ticker response within 15 minutes

ADVANTAGES

DEDICATED TECHNICAL EXECUTIVE

GloboTech Dedicated Technical Account Manager will work with you and your team to plan, deploy, and monitor your private cloud architecture, to help your business overcome the challenges of the future.

INFRASTRUCTURE MIGRATION

Migrating a website or a server from one host to another can be challenging, let GloboTech staff assist you by creating a migration plan and performing the migration for you.

ENTREPRISE GRADE HARDWARE

Our servers are designed exclusively with enterprise-grade hardware, resulting in lower unrecoverable data errors, better speed, and stability.

INFRASTRUCTURE & HIGH AVAILABILITY

By eliminating all single points of failure from both the server and network level, GloboTech offers a high-availability solution that provides the maximum uptime for your web-based content and applications.

OpenStack is an open source hypervisor governed by the OpenStack foundation and backed by thousands of companies and individuals. The popular stack is constituted of multiple projects offering a series of functionality enabling virtualization for your private cloud infrastructure. OpenStack allows resource allocation and management to your server for an increased flexibility and faster provisioning of your virtual machines.

GloboTech provides you with the expertise and the resources to insure the success of your online presence. GloboTech differentiates itself from the competition by delivering uncompromised support, professional assistance and straightforward SLAs to businesses around the world. Our infrastructure and network are fully redundant, delivering 100% availability. You can trust GloboTech to provide you with the most reliable private cloud infrastructure.

For more information about OpenStack services, infrastructure or to get a quote, please contact one of our certified professional at sales@globo.tech. You can also read our various start-up guides and technical documents in the Learning Center.

GET IN TOUCH WITH OUR SALES TEAM

We strive to provide the best pre-sales assistance. Chat now and experience the GloboTech difference.

Toll free: 1-888-482-6661

Montreal: 514-907-0050

By email: sales@globo.tech

CONTACT OUR SUPPORT TEAM

Our support team is available 24/7/365 for all your technical support issues. Use the GloboTech Support portal or use the email address and send us an email.

Toll free: 1-888-482-6661

Montreal: 514-907-0050

By email: support@globo.tech

Montreal Office

630 Sherbrooke W. Suite 200

Montreal, Quebec H3A 1E4, Canada

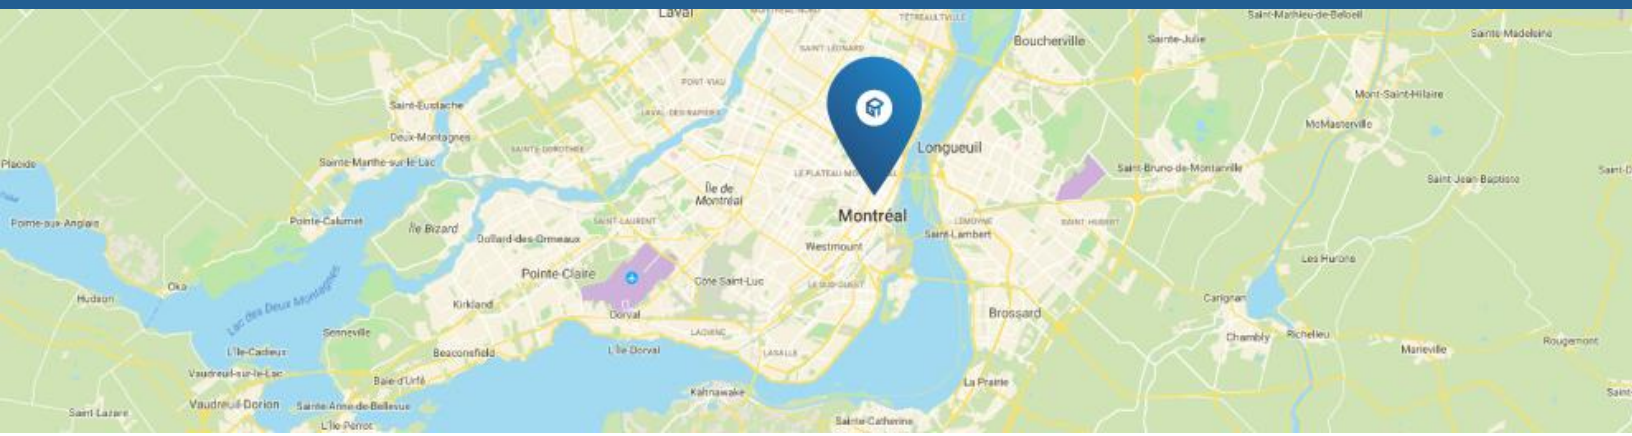

1. http://superuser.openstack.org/articles/survey-says-openstack-and-docker-most-popular-open-source-projects/

2. OpenStack Architecture

   http://docs.openstack.org/developer/nova/architecture.html

3. Release History

   https://en.wikipedia.org/wiki/OpenStack

4. OpenStack Nova

   https://www.openstack.org/software/releases/mitaka/components/nova

5. OpenStack Neutron

   https://www.openstack.org/software/releases/mitaka/components/neutron

6. OpenStack Cinder

   https://www.openstack.org/software/releases/mitaka/components/cinder

7. OpenStack Glance

   https://www.openstack.org/software/releases/mitaka/components/glance

8. OpenStack Horizon

   https://www.openstack.org/software/releases/mitaka/components/horizon