Cloud Services and Platforms

Cloud Reference Model

• Infrastructure & Facilities Layer

Includes the physical infrastructure such as datacenter facilities, electrical and mechanical equipment, etc.

• Hardware Layer

Includes physical compute, network and storage hardware.

• Virtualization Layer

Partitions the physical hardware resources into multiple virtual resources that enabling pooling of resources.

• Platform & Middleware Layer

Builds upon the IaaS layers below and provides standardized stacks of services such as database service, queuing service, application frameworks and run-time environments, messaging services, monitoring services, analytics services, etc.

• Service Management Layer

Provides APIs for requesting, managing and monitoring cloud resources.

• Applications Layer

Includes SaaS applications such as Email, cloud storage application, productivity applications, management portals, customer self-service portals, etc.

• Infrastructure & Facilities Layer

Includes the physical infrastructure such as datacenter facilities, electrical and mechanical equipment, etc.

• Hardware Layer

Includes physical compute, network and storage hardware.

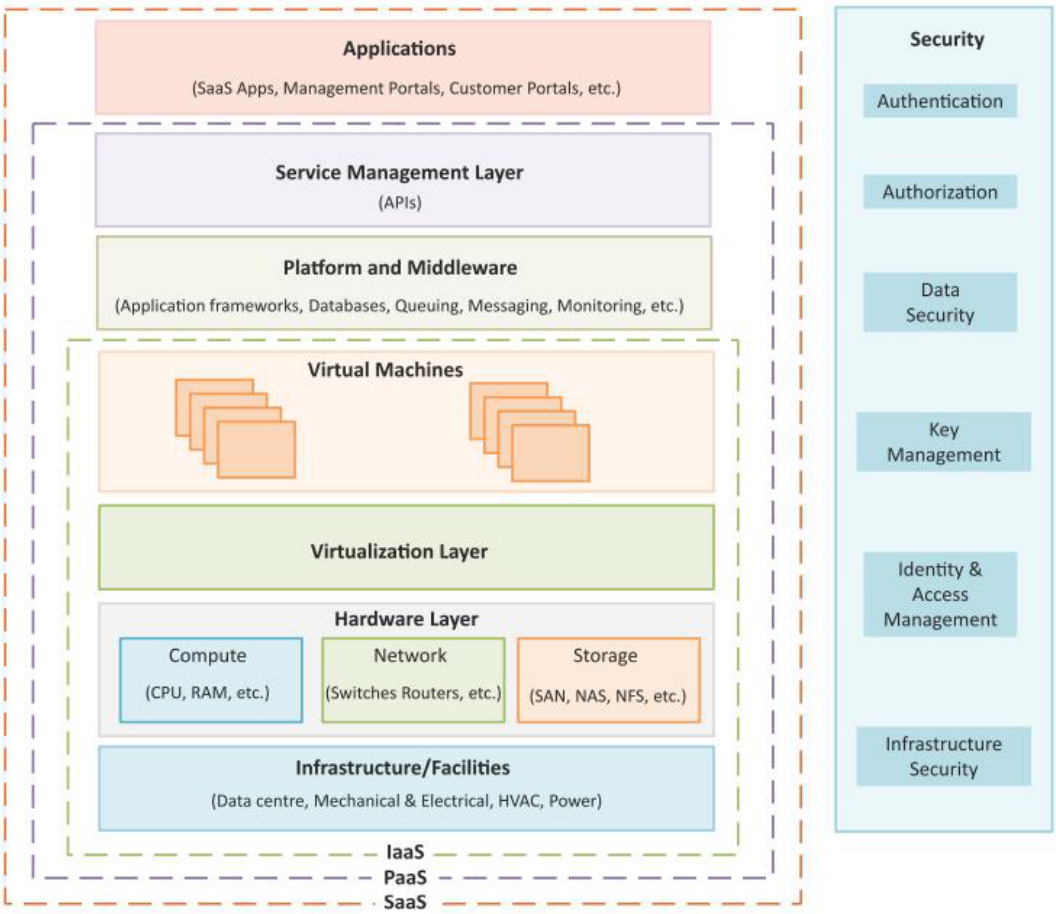

Compute Service

- Compute services provide dynamically scalable compute capacity in the cloud.
- Compute resources can be provisioned on-demand in the form of virtual machines. Virtual machines can be created from standard images provided by the cloud service provider or custom images created by the users.
- Compute services can be accessed from the web consoles of these services that provide graphical user interfaces for provisioning, managing and monitoring these services.
- Cloud service providers also provide APIs for various programming languages that allow developers to access and manage these services programmatically.

Compute Service - Amazon EC2

- Amazon Elastic Compute Cloud (EC2) is a compute service provided by Amazon.
- Launching EC2 Instances

  To launch a new instance click on the launch instance button. This will open a wizard where you can select the Amazon machine image (AMI) with which you want

to launch the instance. You can also create their own AMIs with custom applications, libraries and data. Instances can be launched with a variety of operating systems.

- Instance Sizes

  When you launch an instance you specify the instance type (micro, small, medium, large, extra-large, etc.), the number of instances to launch based on the selected AMI and availability zones for the instances.

- Key-pairs

  When launching a new instance, the user selects a key-pair from existing keypairs or creates a new keypair for the instance. Keypairs are used to securely connect to an instance after it launches.

- Security Groups

  The security groups to be associated with the instance can be selected from the instance launch wizard. Security groups are used to open or block a specific network port for the launched instances.

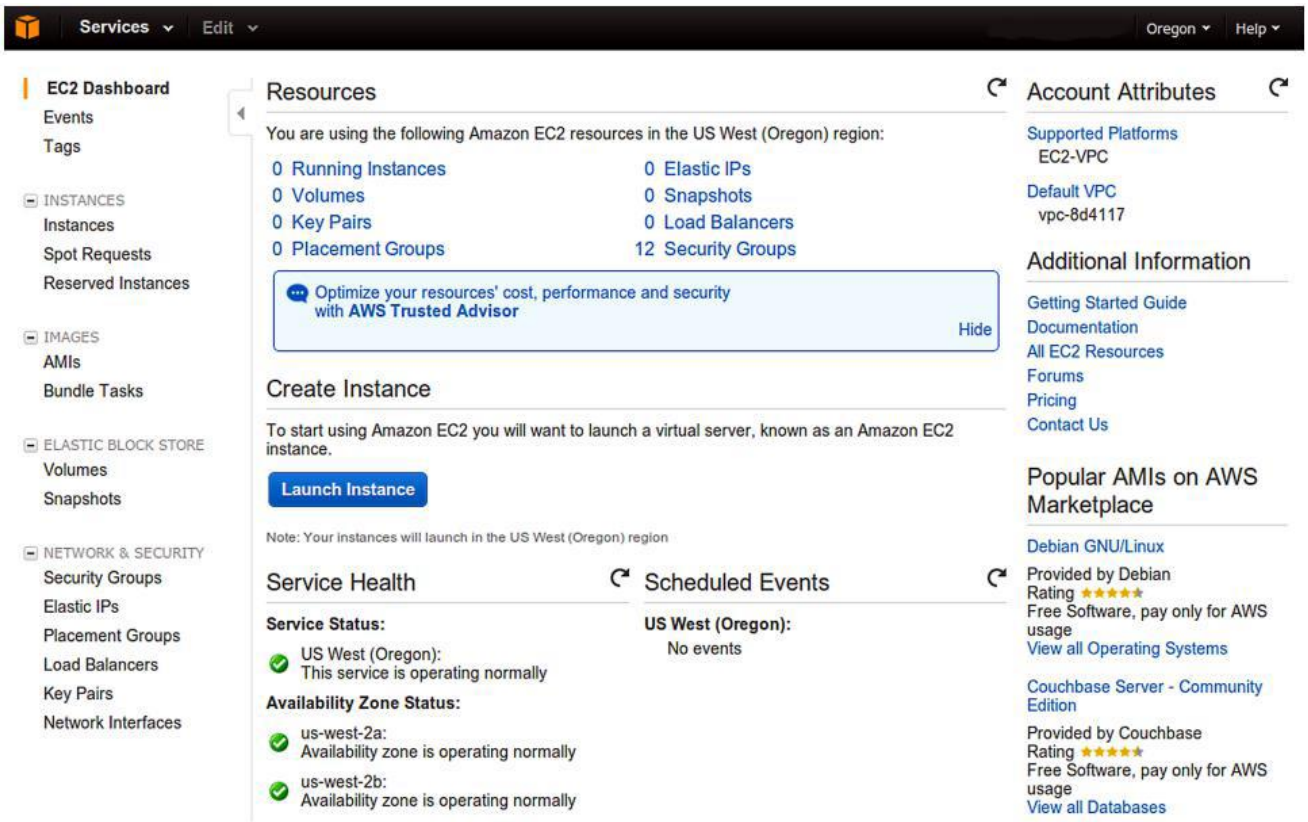

Compute Services – Google Compute Engine

- Google Compute Engine is a compute service provided by Google.

- **Launching Instances**

To create a new instance, the user selects an instance machine type, a zone in which the instance will be launched, a machine image for the instance and provides an instance name, instance tags and meta-data.

- **Disk Resources**

Every instance is launched with a disk resource. Depending on the instance type, the disk resource can be a scratch disk space or persistent disk space. The scratch disk space is deleted when the instance terminates. Whereas, persistent disks live beyond the life of an instance.

- **Network Options**

Network option allows you to control the traffic to and from the instances. By default, traffic between instances in the same network, over any port and any protocol and incoming SSH connections from anywhere are enabled.

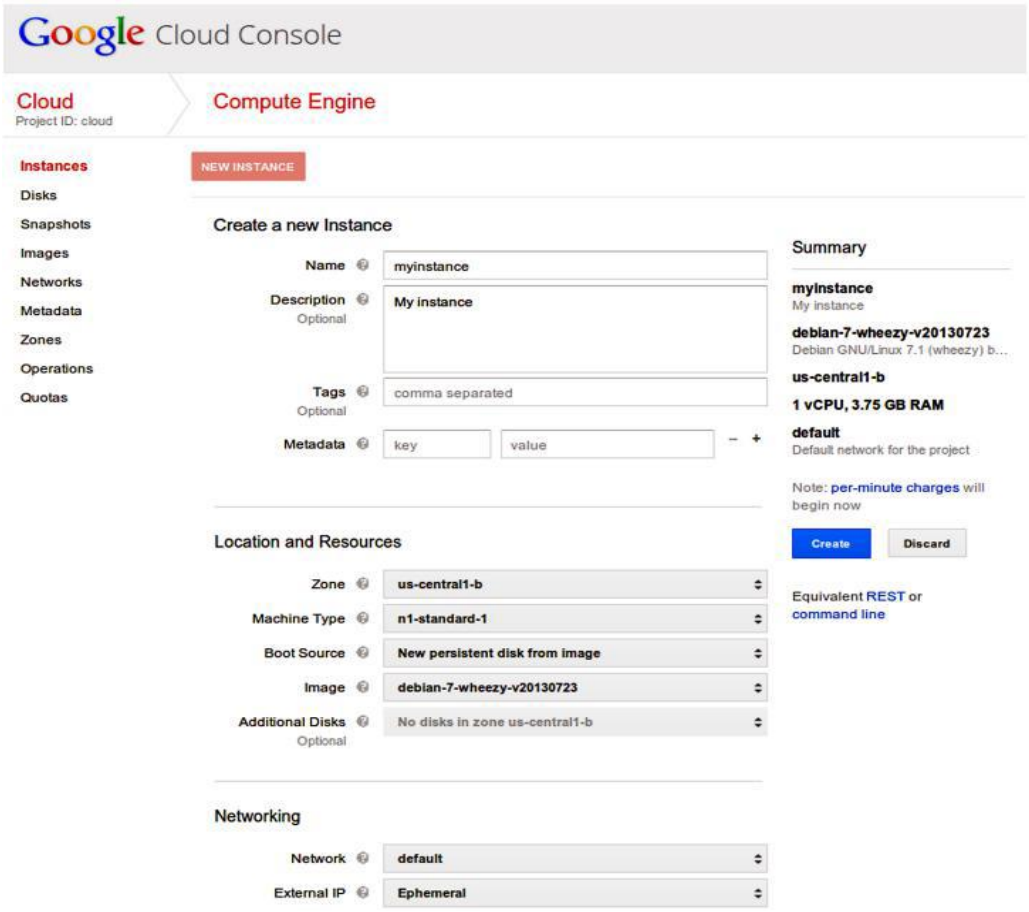

Compute Services – Windows Azure VMs

- Windows Azure Virtual Machines is the compute service from Microsoft.
- Launching Instances:
  - To create a new instance, you select the instance type and the machine image.
  - You can either provide a user name and password or upload a certificate file for securely connecting to the instance.
  - Any changes made to the VM are persistently stored and new VMs can be created from the previously stored machine images.

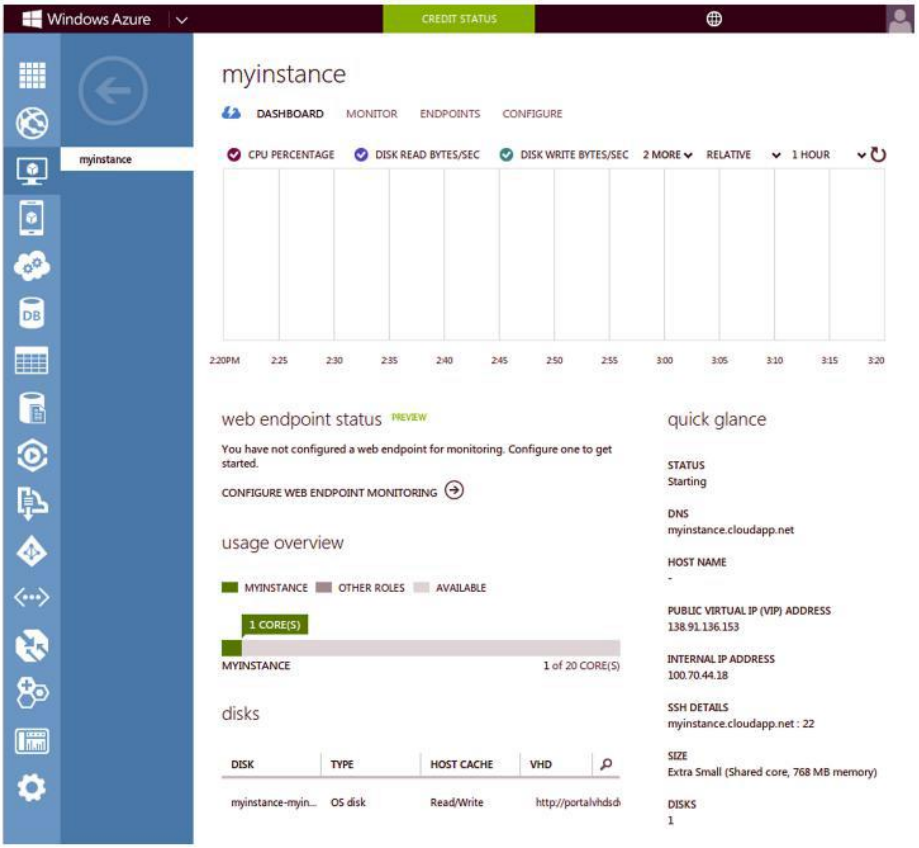

Storage Services

- Cloud storage services allow storage and retrieval of any amount of data, at any time from anywhere on the web.
- Most cloud storage services organize data into buckets or containers.

- Scalability

Cloud storage services provide high capacity and scalability. Objects upto several tera-bytes in size can be uploaded and multiple buckets/containers can be created on cloud storages.

- Replication

When an object is uploaded it is replicated at multiple facilities and/or on multiple devices within each facility.

- Access Policies

Cloud storage services provide several security features such as Access Control Lists (ACLs), bucket/container level policies, etc. ACLs can be used to selectively grant access permissions on individual objects. Bucket/container level policies can also be defined to allow or deny permissions across some or all of the objects within a single bucket/container.

- Encryption

Cloud storage services provide Server Side Encryption (SSE) options to encrypt all data stored in the cloud storage.

- Consistency

Strong data consistency is provided for all upload and delete operations. Therefore, any object that is uploaded can be immediately downloaded after the upload is complete.

Storage Services – Amazon S3

- Amazon Simple Storage Service(S3) is an online cloud-based data storage infrastructure for storing and retrieving any amount of data.
- S3 provides highly reliable, scalable, fast, fully redundant and affordable storage infrastructure.
- **Buckets**
  - Data stored on S3 is organized in the form of buckets. You must create a bucket before you can store data on S3.

- **Uploading Files to Buckets**
  - S3 console provides simple wizards for creating a new bucket and uploading files.
  - You can upload any kind of file to S3.
  - While uploading a file, you can specify the redundancy and encryption options and access permissions.

Storage Services – Google Cloud Storage

- GCS is the Cloud storage service from Google
- Buckets

  Objects in GCS are organized into buckets.
- Access Control Lists

  ACLs are used to control access to objects and buckets. ACLs can be configured to share objects and buckets with the entire world, a Google group, a Google-hosted domain, or specific Google account holders.

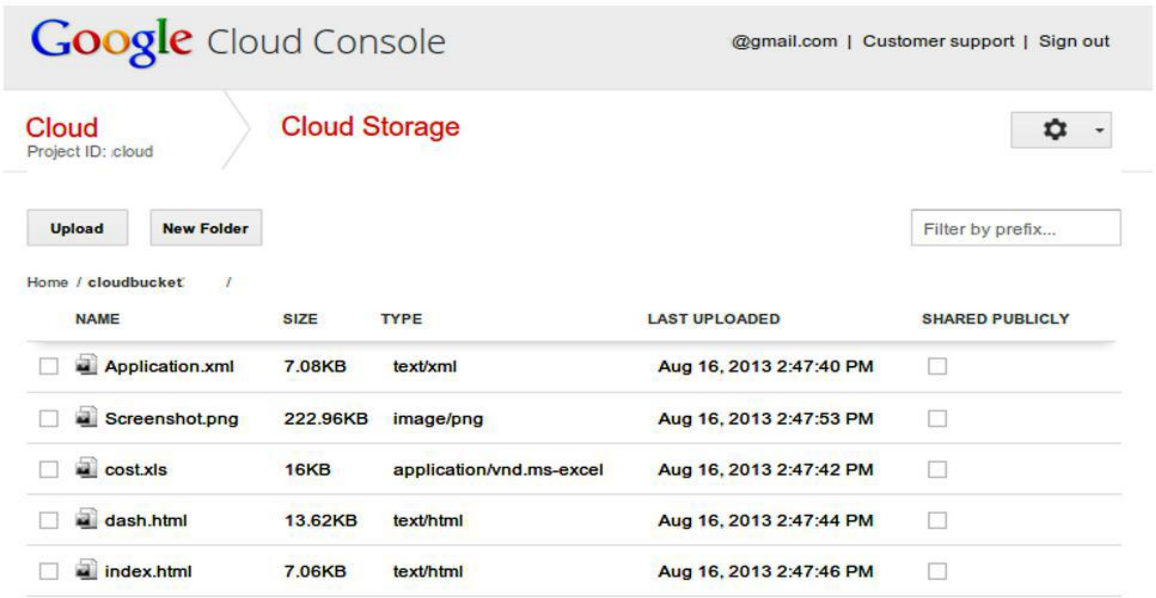

Storage Services – Windows Azure Storage

- Windows Azure Storage is the cloud storage service from Microsoft.
- Windows Azure Storage provides various storage services such as blob storage service, table service and queue service.
- Blob storage service
    - The blob storage service allows storing unstructured binary data or binary large objects (blobs).
    - Blobs are organized into containers.
    - Block blobs - can be subdivided into some number of blocks. If a failure occurs while transferring a block blob, retransmission can resume with the most recent block rather than sending the entire blob again.
    - Page blobs - are divided into number of pages and are designed for random access. Applications can read and write individual pages at random in a page blob.

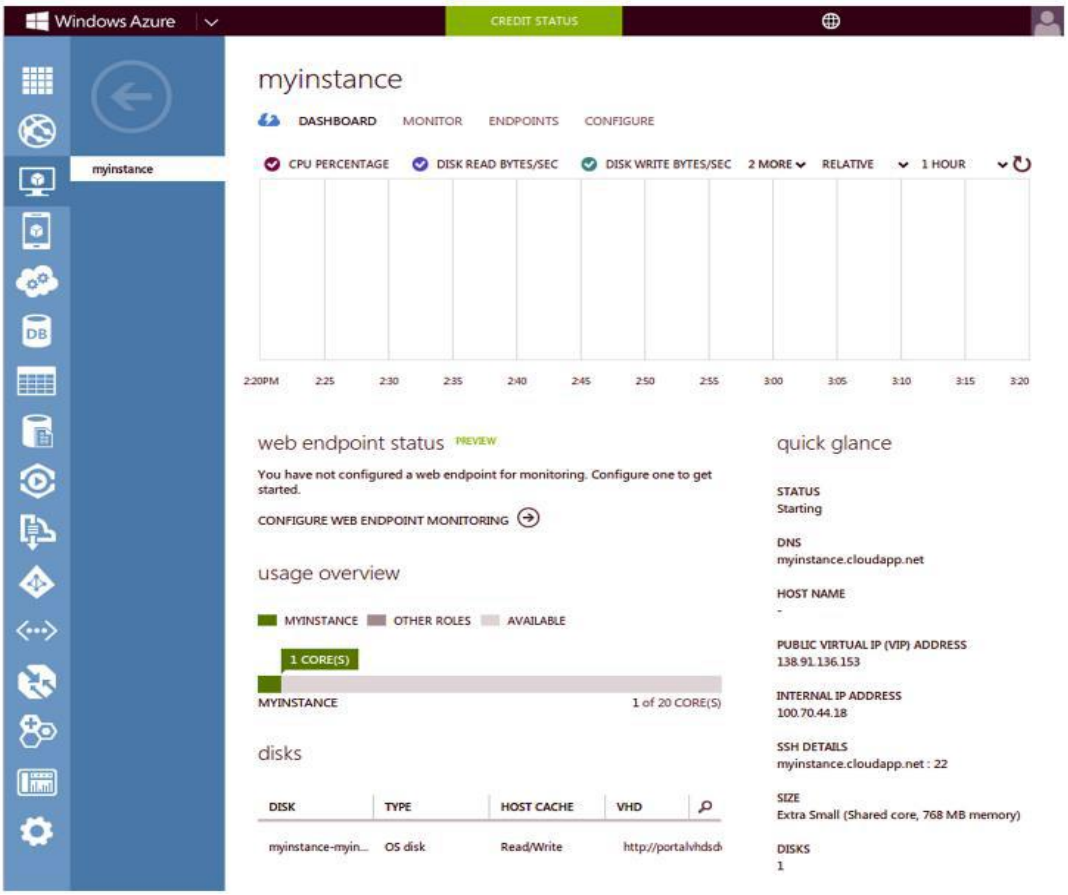

Application Runtimes & Frameworks

- Cloud-based application runtimes and frameworks allow developers to develop and host applications in the cloud.
- Support for various programming languages

  Application runtimes provide support for programming languages (e.g., Java, Python, or Ruby).
- Resource Allocation

  Application runtimes automatically allocate resources for applications and handle the application scaling, without the need to run and maintain servers.

Google App Engine

- Google App Engine is the platform-as-a-service (PaaS) from Google, which includes both an application runtime and web frameworks.
- Runtimes
  - App Engine provides runtime environments for Java, Python, PHP and Go programming language.
- Sandbox
  - Applications run in a secure sandbox environment isolated from other applications.
  - The sandbox environment provides a limited access to the underlying operating system.

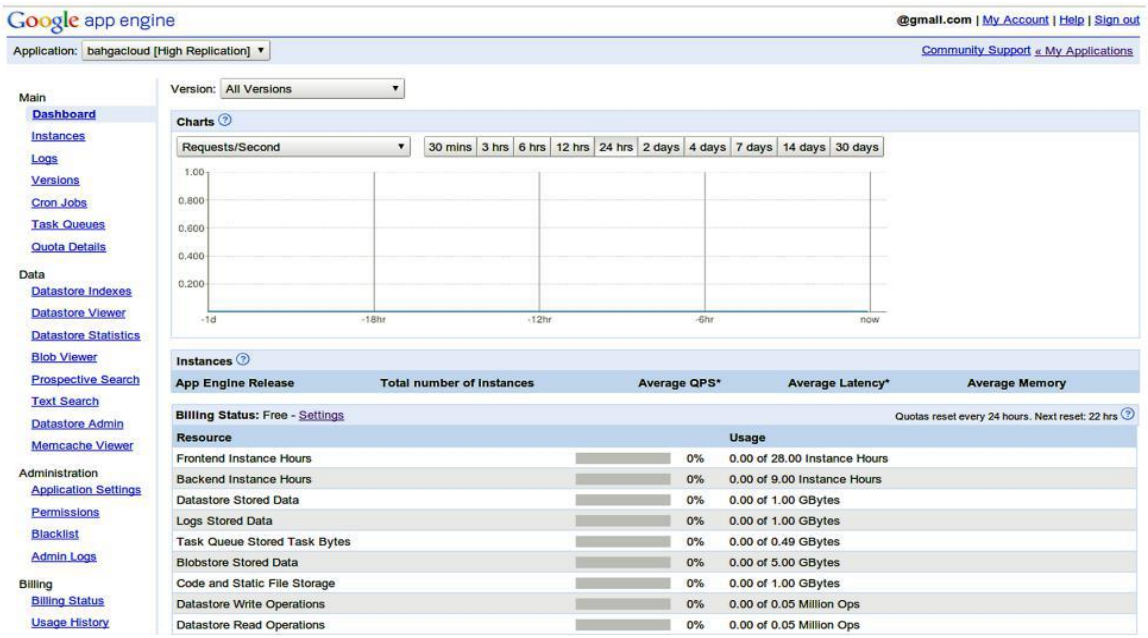

- Web Frameworks
  - App Engine provides a simple Python web application framework called webapp2. App Engine also supports any framework written in pure Python that speaks WSGI, including Django, CherryPy, Pylons, web.py, and web2py.
- Datastore
  - App Engine provides a no-SQL data storage service
- Authentication
  - App Engine applications can be integrated with Google Accounts for user authentication.
- URL Fetch service
  - URL Fetch service allows applications to access resources on the Internet, such as web services or other data.
- Other services
  - Email service
  - Image Manipulation service
  - Memcache
  - Task Queues
  - Scheduled Tasks service

Windows Azure Web Sites

- Windows Azure Web Sites is a Platform-as-a-Service (PaaS) from Microsoft.
- Azure Web Sites allows you to host web applications in the Azure cloud.
- Shared & Standard Options.
  - In the shared option, Azure Web Sites run on a set of virtual machines that may contain multiple web sites created by multiple users.
  - In the standard option, Azure Web Sites run on virtual machines (VMs) that belong to an individual user.
- Azure Web Sites supports applications created in ASP.NET, PHP, Node.js and Python programming languages.
- Multiple copies of an application can be run in different VMs, with Web Sites automatically load balancing requests across them.