

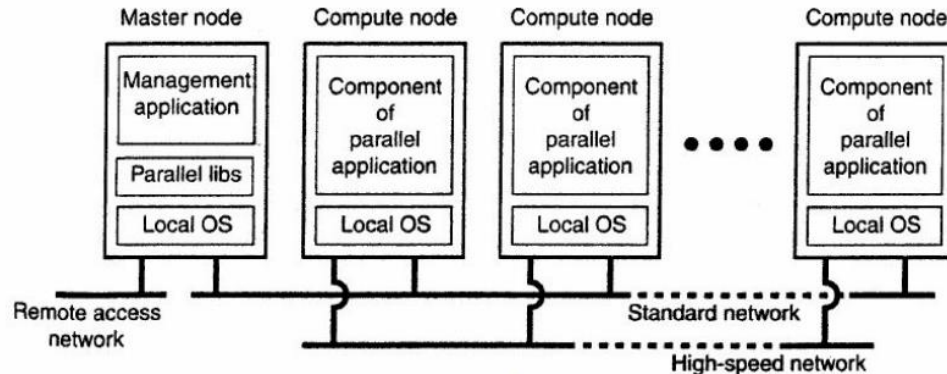
Computing Clusters & Management

Md Anisul Islam
anisul.iut@gmail.com



What is Computing Cluster?

- A computing cluster is a parallel or distributed computer system
 - Collection of interconnected stand-alone computers
 - Under one administration
 - Working together as a single integrated computing resource



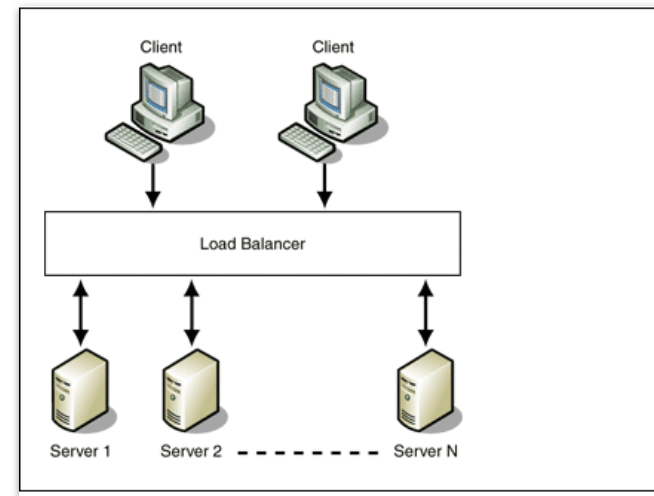
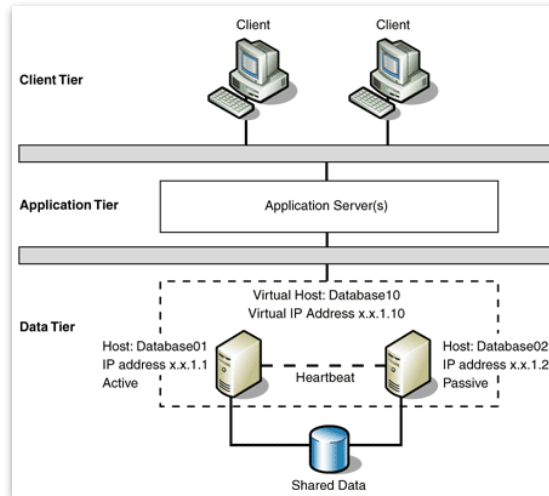


Components of Cluster

- Multiple standalone computers
- Operating system
- High-performance interconnects
- Middleware
- Nodes

Types of Computing Clusters

- Failover clusters
- High-performance clusters
- Load balancing clusters





Advantages

- Better availability and reliability
- Scalability
- Enhanced network performance
- Easy troubleshooting

Disadvantage

- Maintenance
- Resource consumption
- Application incompatibility




Cluster Management & Tools

Cluster management includes:

- Monitoring nodes
- Resource allocation
- Failure recovery

Popular cluster management tools:

- Swarm
- Fleet
- Google Kubernetes
- Apache Mesos

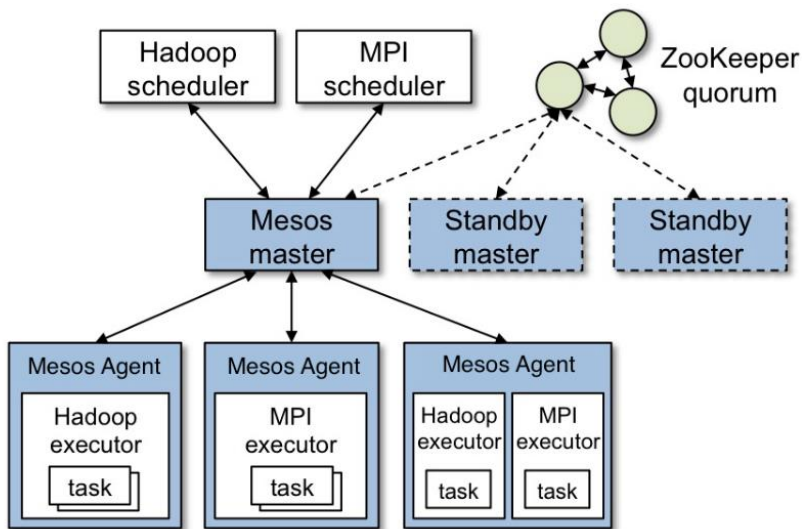


kubernetes
by Google



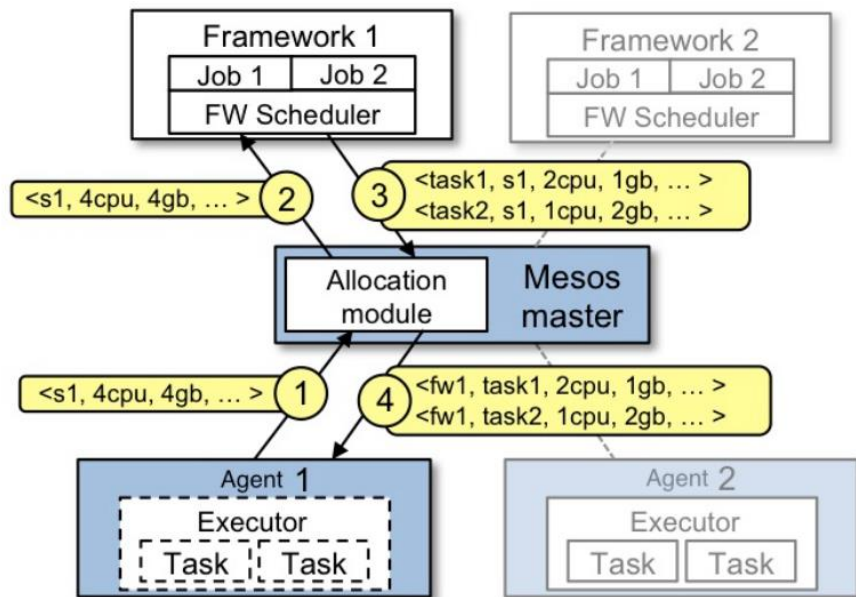
Apache
MESOS™

Apache Mesos: Architecture



- Master daemons manages agent-daemon running in nodes
- Master decides how many resources to offer to each framework according to a given organizational policy, such as fair sharing or strict priority
- A framework running on top of Mesos consists of 2 components: a scheduler and an executor
- Scheduler registers with the master to be offered resources and executor process that is launched on agent nodes to run the framework's tasks
- Master determines how many resources are offered to each framework, the frameworks' schedulers select which of the offered resources to use

Apache Mesos: Resource Offer



- Agent 1 reports to the master it has 4 CPUs and 4 GB of memory free. Master then invokes the allocation policy module, which tells Framework 1 should be offered all available resources.
- The master sends a resource offer describing what is available on Agent 1 to Framework 1
- The framework's scheduler replies to the master with information about two tasks to run on the agent, using `<2 CPUs, 1 GB RAM>` for the first task, and `<1 CPUs, 2 GB RAM>` for the second task.
- Finally, the master sends the tasks to the agent, which allocates appropriate resources to the framework's executor,



Google Kubernetes Demo

[Demo](#)



Google Kubernetes Demo

Google Cloud Platform Walletz

Kubernetes clusters [+ CREATE CLUSTER](#) [REFRESH](#) [DELETE](#) [SHOW INFO PANEL](#)

Filter by label or name

Kubernetes clusters

<input type="checkbox"/>	Name ^	Location	Cluster size	Total cores	Total memory	Notifications	Labels	
<input type="checkbox"/>	cluster-1	us-central1-a	3	3 vCPUs	11.25 GB			Connect



Google Kubernetes Demo

Google Cloud Platform Walletz

Kubernetes clusters EDIT DELETE CONNECT

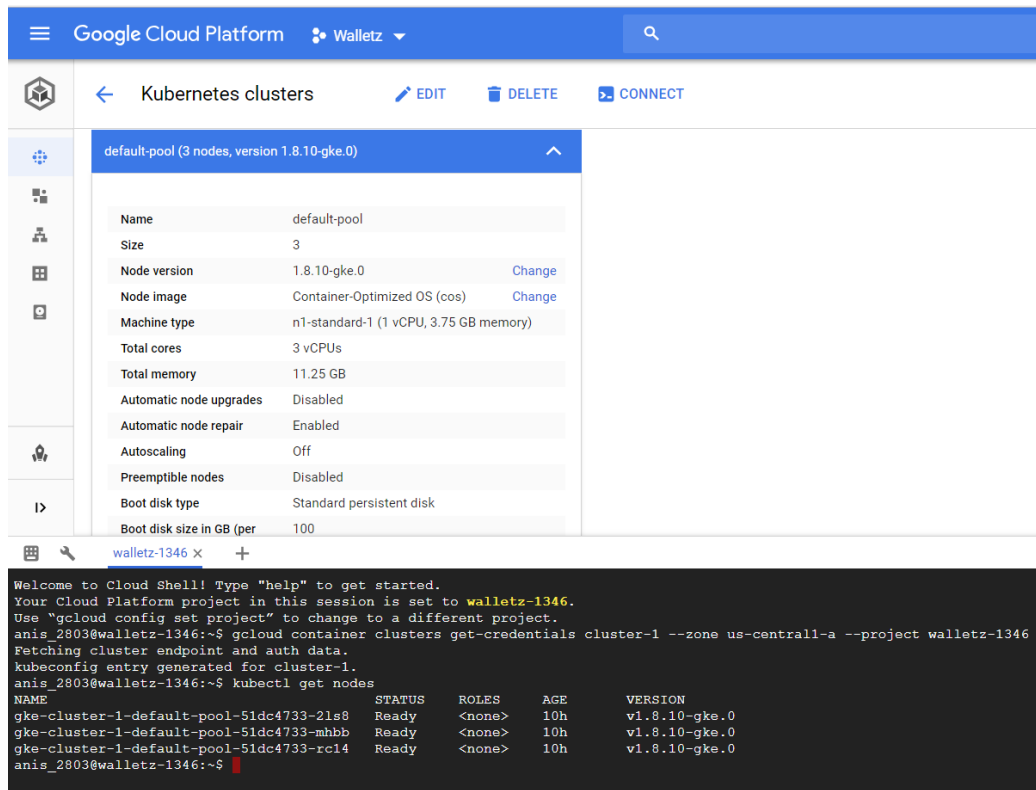
cluster-1

[Details](#) [Storage](#) [Nodes](#)

Cluster

Master version	1.8.10-gke.0	Upgrade available
Endpoint	35.232.85.41	Show credentials
Client certificate	Enabled	
Kubernetes alpha features	Disabled	
Total size	3	
Master zone	us-central1-a	
Node zones	us-central1-a	
Network	default	
Subnet	default	
VPC-native (alias IP)	Disabled	
Container address range	10.8.0.0/14	
Stackdriver Logging	Enabled	
Stackdriver Monitoring	Enabled	
Private cluster	Disabled	
Master authorized networks	Disabled	
Network policy	Disabled	
Legacy authorization	Disabled	
Maintenance window	Any time	

Google Kubernetes Demo



The screenshot displays the Google Cloud Platform console interface. At the top, the navigation bar shows "Google Cloud Platform" and the project name "Walletz". Below this, the "Kubernetes clusters" page is active, showing a cluster named "default-pool (3 nodes, version 1.8.10-gke.0)".

The cluster details are as follows:

Property	Value	Action
Name	default-pool	
Size	3	
Node version	1.8.10-gke.0	Change
Node image	Container-Optimized OS (cos)	Change
Machine type	n1-standard-1 (1 vCPU, 3.75 GB memory)	
Total cores	3 vCPUs	
Total memory	11.25 GB	
Automatic node upgrades	Disabled	
Automatic node repair	Enabled	
Autoscaling	Off	
Preemptible nodes	Disabled	
Boot disk type	Standard persistent disk	
Boot disk size in GB (per node)	100	

Below the cluster details, a terminal window is open for the project "walletz-1346". The terminal output shows the following commands and results:

```
Welcome to Cloud Shell! Type "help" to get started.
Your Cloud Platform project in this session is set to walletz-1346.
Use "gcloud config set project" to change to a different project.
anis_2803@walletz-1346:~$ gcloud container clusters get-credentials cluster-1 --zone us-central1-a --project walletz-1346
Fetching cluster endpoint and auth data.
kubeconfig entry generated for cluster-1.
anis_2803@walletz-1346:~$ kubectl get nodes
```

NAME	STATUS	ROLES	AGE	VERSION
gke-cluster-1-default-pool-51dc4733-21s8	Ready	<none>	10h	v1.8.10-gke.0
gke-cluster-1-default-pool-51dc4733-mhbb	Ready	<none>	10h	v1.8.10-gke.0
gke-cluster-1-default-pool-51dc4733-rc14	Ready	<none>	10h	v1.8.10-gke.0

```
anis_2803@walletz-1346:~$
```

Google Kubernetes Demo

The screenshot shows the Google Cloud Pricing page for Kubernetes Engine. The page is divided into two main sections: configuration options on the left and a cost estimate on the right.

Navigation: The top navigation bar includes "Why Google", "Products", "Solutions", "Launcher", "Pricing" (selected), "Security", "Customers", "Documentation", "Support", "Partners", and "CONTACT SALES". A search bar and "CONSOLE" link are also present.

Product Selection: A row of icons represents various Google Cloud products: COMPUTE ENGINE, APP ENGINE, KUBERNETES ENGINE (highlighted), CLOUD STORAGE, NETWORKING, BIGQUERY, CLOUD DATASTORE, and CLOUD DATAPROC.

Kubernetes Engine Configuration:

- Number of nodes:** A text input field with a help icon.
- What are these nodes for?** A text input field with a help icon.
- Instance type:** A dropdown menu showing "n1-standard-1 (vCPUs: 1, RAM: 3.75 GB)" with a help icon.
- Add GPUs:** A checkbox that is currently unchecked with a help icon.
- Local SSD:** A dropdown menu showing "0" with a help icon.
- Datcenter location:** A dropdown menu showing "Iowa (us-central1)" with a help icon.
- Average hours per day each node is running:** A dropdown menu showing "24" hours, with a help icon.
- Average days per week each node is running:** A dropdown menu showing "7" with a help icon.

Persistent Disk Configuration:

- Location:** A dropdown menu showing "Iowa (us-central1)" with a help icon.
- Persistent disk storage:** A dropdown menu showing "GB" with a help icon.
- Snapshot storage:** A dropdown menu showing "GB" with a help icon.

Cost Estimate (Estimate 1):

- Persistent Disk:** Iowa, Storage: 10 GB, Price: \$0.40.
- Kubernetes Engine:** 3 x, 2,190 total hours per month, Instance type: n1-standard-1, Region: Iowa, GCE instance Cost: \$72.82, Container Engine Cost: \$0.00.
- Discounts:** Sustained Use Discount: 30%, Effective Hourly Rate: \$0.033.
- Estimated Component Cost:** \$72.82 per 1 month.
- Total Estimated Cost:** \$73.22 per 1 month (USD).
- Adjust Estimate Timeframe:** A timeline slider with markers for 1 day, 1 week, 1 month (selected), 1 quarter, 1 year, and 3 years.
- Buttons:** "EMAIL ESTIMATE" and "SAVE ESTIMATE".

Additional Elements: A grey "ADD TO ESTIMATE" button is located at the bottom of the configuration section. A "CONTACT SALES" button is in the top right navigation bar.