Distributed Databases

Anoop Kumar Malav – Y08uc029
Chirag Jain – Y08uc041
Manohar Kuse – Y08uc073

If one person cannot complete a job, then employ two. This is what cluster computing means!

Apr, 2010

Clusters

- A group of the same or similar elements gathered or occurring closely together is called cluster.
- A computer cluster is a group of linked computers, working together closely so that in many respects they form a single computer.
- Ex : Super Computers.
- Cluster adds a functionality called scalability.

Introduction

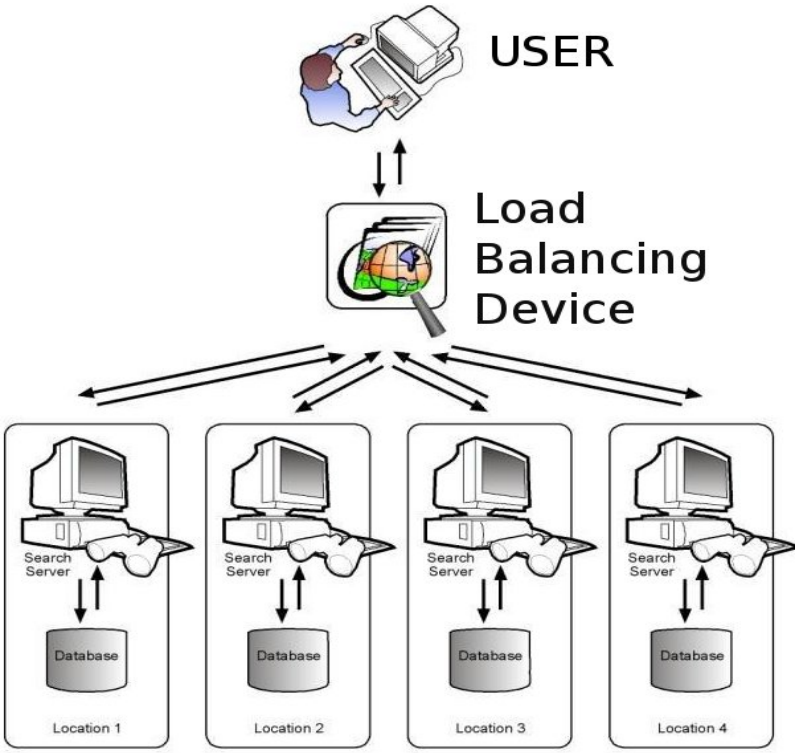

1. More Database servers can be connected to this Load Balancing Device.
2. 1 pc site has a limitation on storage, memory and computing device.
   While distributed is scalable.

Motivations

- Increased Availability :-
  - If we consider a site which uses a cluster database, even if one of the storage server goes down, the entire data is still available, which may be retrived automatically from another storage server.
  - For a single PC (or a single server) site, if that goes down, the entire data is unavailable.
  - Also if, the load on your server increases to something like a few thousand requests per second, your server will not be able to tolerate such high rate of request.
- Distributed Access To Data :-
  Any organization may have branches in several cities.
- Parallel Processing .
- Scalability

Set up

- I am going to talk about three servers :
- Load Balancer 172.22.18.116
- Storage1 172.22.18.117
- Storage2 172.22.18.118

STAGE 1: Install MySQL on the two Storage Servers

- Complete the following steps on both Storage1 and Storage2 :
- cd /usr/local/
- Wget http://site/package
- groupadd mysql
- useradd -g mysql mysql
- tar -zxvf package
- ln -s package mysql
- cd mysql

STAGE 1: Install MySQL on the two Storage Servers Cont ...

- ./scripts/mysql_install_db –user=mysql
- chown -R root .
- chown -R mysql data
- chgrp -R mysql .
- cp support-files/mysql.server /etc/rc.d/init.d/
- chmod +x /etc/rc.d/init.d/mysql.server
- chkconfig –add mysql.server

STAGE 2: Install and Configure the Managment Server(172.22.18.116)

- mkdir /usr/src/mysql-mgm
- cd /usr/src/mysql-mgm
- wget package
- tar -zxvf package

STAGE 2: Install and Configure the Managment Server(172.22.18.116) Cont ...

- cp bin/ndb_mgm .
- cp bin/ndb_mgmd .
- chmod +x ndb_mg*
- mv ndb_mg* /usr/bin/
- cd
- rm -rf /usr/src/mysql-mgm

STAGE 2: Install and Configure the Managment Server(172.22.18.116) Cont ...

- set up the **config file for this management server:**
- **mkdir /var/lib/mysql-cluster**
- **vi [or emacs or gedit or any other editor] /var/lib/mysql-cluster/config.ini**
- **Now, insert the following (changing the bits as indicated):**

STAGE 2: Management Server Configuration File (172.22.18.116) Cont ...

```
[NDBD DEFAULT]
NoOfReplicas=2
[MYSQLD DEFAULT]
[NDB_MGMD DEFAULT]
[TCP DEFAULT]
# Managment Server
[NDB_MGMD]
HostName=172.22.18.116 # the IP of THIS SERVER
# Storage Engines
[NDBD]
HostName=172.22.18.117 # the IP of the FIRST SERVER
DataDir=/var/lib/mysql-cluster
[NDBD]
HostName=172.22.18.118 # the IP of the SECOND SERVER
DataDir=/var/lib/mysql-cluster
[MYSQLD]

[MYSQLD]
```

STAGE 2: Management Server Configuration File (172.22.18.116) Cont ...

Now, **start the managment server:**
ndb_mgmd
This is the MySQL managment server, not maganment console.
You should therefore not expect any output (we will start the console later).

STAGE 3: Configure the Storage/SQL Servers and Start MySQL Cluster

- On each of the two storage/SQL servers (172.22.18.117 and 172.22.18.118) enter the following (changing the bits as appropriate):
- vi /etc/my.cnf
- Insert this on both servers (changing the IP address to the IP of the managment server that you set up in stage 2)

STAGE 3: Configure the Storage/SQL Servers and Start MySQL Cluster, Contd ...

```
[mysqld]
Ndbcluster
ndb-connectstring=172.22.18.116 # the IP of the MANAGMENT (THIRD) SERVER
[mysql_cluster]

ndb-connectstring=172.22.18.116 # the IP of the MANAGMENT (THIRD) SERVER
```

STAGE 3: Configure the Storage/SQL Servers and Start MySQL Cluster, Contd ...

Now, we make the data directory and start the storage engine:
mkdir /var/lib/mysql-cluster
cd /var/lib/mysql-cluster
/usr/local/mysql/bin/ndbd –initial
/etc/rc.d/init.d/mysql.server start
If you have done one server now go back to the start of stage 3 and repeat exactly the same procedure on the second server.
Note: you should ONLY use –initial if you are either starting from scratch or have changed the config.ini file on the managment.

STAGE 4: Check its Working

- You can now return to the managment server (mysql3) and enter the managment console:
- /usr/local/mysql/bin/ndb_mgm
- Enter the command SHOW to see what is going on. A sample output looks like this:

STAGE 4: Check its Working, Contd ...

[root@mysql3 mysql-cluster] #/usr/local/mysql/bin/ndb_mgm
– NDB Cluster – Management Client
ndb_mgm¿ show
Connected to Management Server at: localhost:1186
Cluster Configuration
———————
[ndbd(NDB)] 2 node(s)
id=2 @172.22.18.117 (Version: 4.1.9, Nodegroup: 0, Master)
id=3 @172.22.18.118 (Version: 4.1.9, Nodegroup: 0)
[ndb_mgmd(MGM)] 1 node(s)
id=1 @172.22.18.116 (Version: 4.1.9)
[mysqld(API)] 2 node(s)
If you are OK to here it is time to test mysql.
[root@mysql1] # mysql
use test;
CREATE TABLE ctest (i INT) ENGINE=NDBCLUSTER;
INSERT INTO ctest () VALUES (1);
SELECT * FROM ctest;
If this works, now go to the other server and run the same SELECT and see what you get.
Insert from that host and go back to host 1 and see if it works.

If it works then congratulations.

Web Application Deployment on Cluster

- There is no difference in application development.
- The only change is that, you need to change your storage engine.
- *CREATE TABLE ctest CREATE TABLE ctest (i INT) ENGINE=NDBCLUSTER;*
  INSERT INTO ctest () VALUES (1);

  *SELECT * FROM ctest;*
- And then ofcourse host your application on our cluster.

Bench-Marks

- A benchmark is the act of running a computer program, a set of programs, or other operations, in order to assess the relative performance of an object, normally by running a number of standard tests and trials against it.
- We have borrowed the bench-mark result from
  http : //blogs.sun.com/hasham/entry/mysql_cluster_7_performance_benchmark.
- In benchmarking a database system, we have looked primarily into computing power of the cluster.
- Popular softwares for bench-marking database clusters are OSDB (Open Source Database Benchmark), DBT2.

System Configuration

- **Load Balancer** – X4240 with 2 CPU, 8 Cores, 8GB RAM
- **Storage Nodes** – X4600 with 8 CPU, 32 Cores, 64 GB RAM

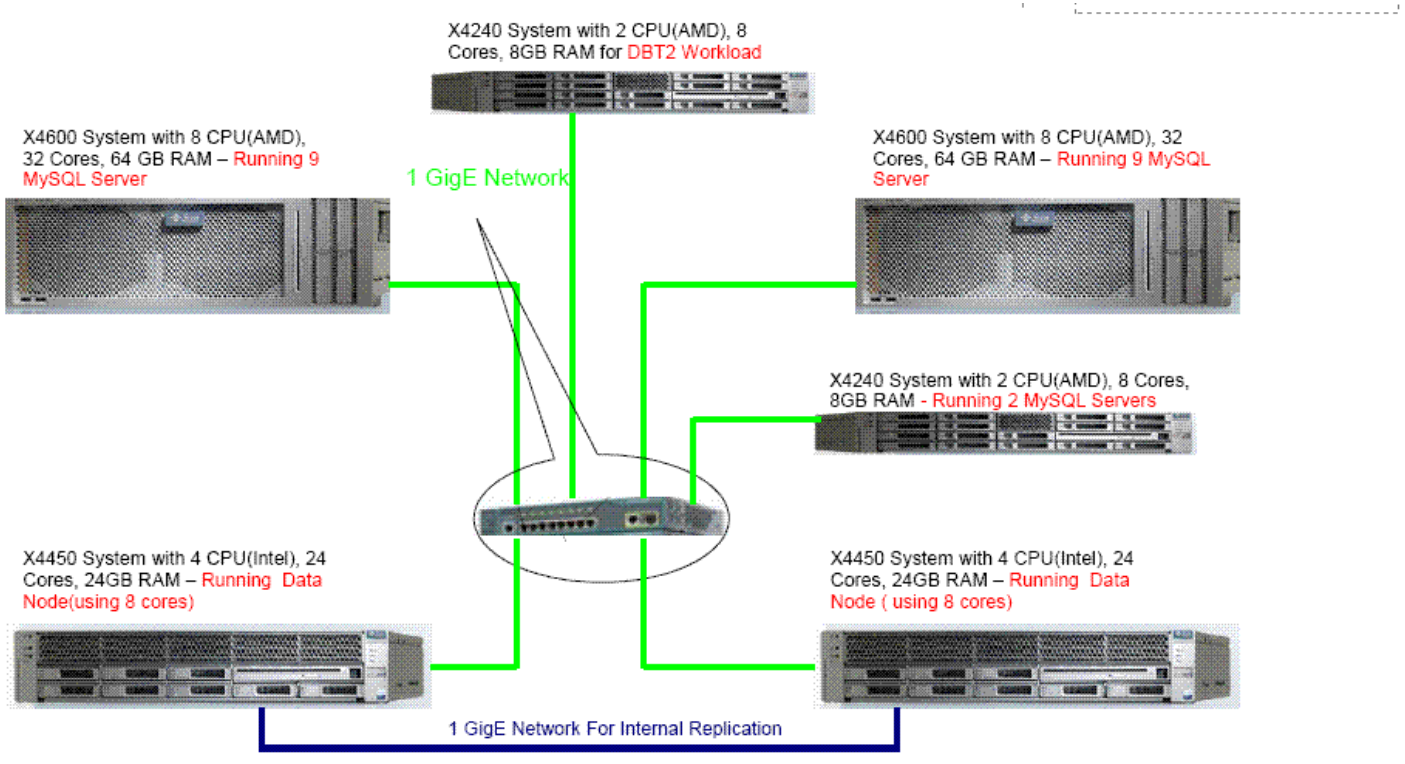

Quantitative Results

- Here, since the servers are powerful, they are running multiple mysql servers on one physical computer.
- Each mysql server is bound to three cores.
- Increasing cores per mysql increases response time, but decreases overall throughput (Performance in unit time).

Quantitative Performance

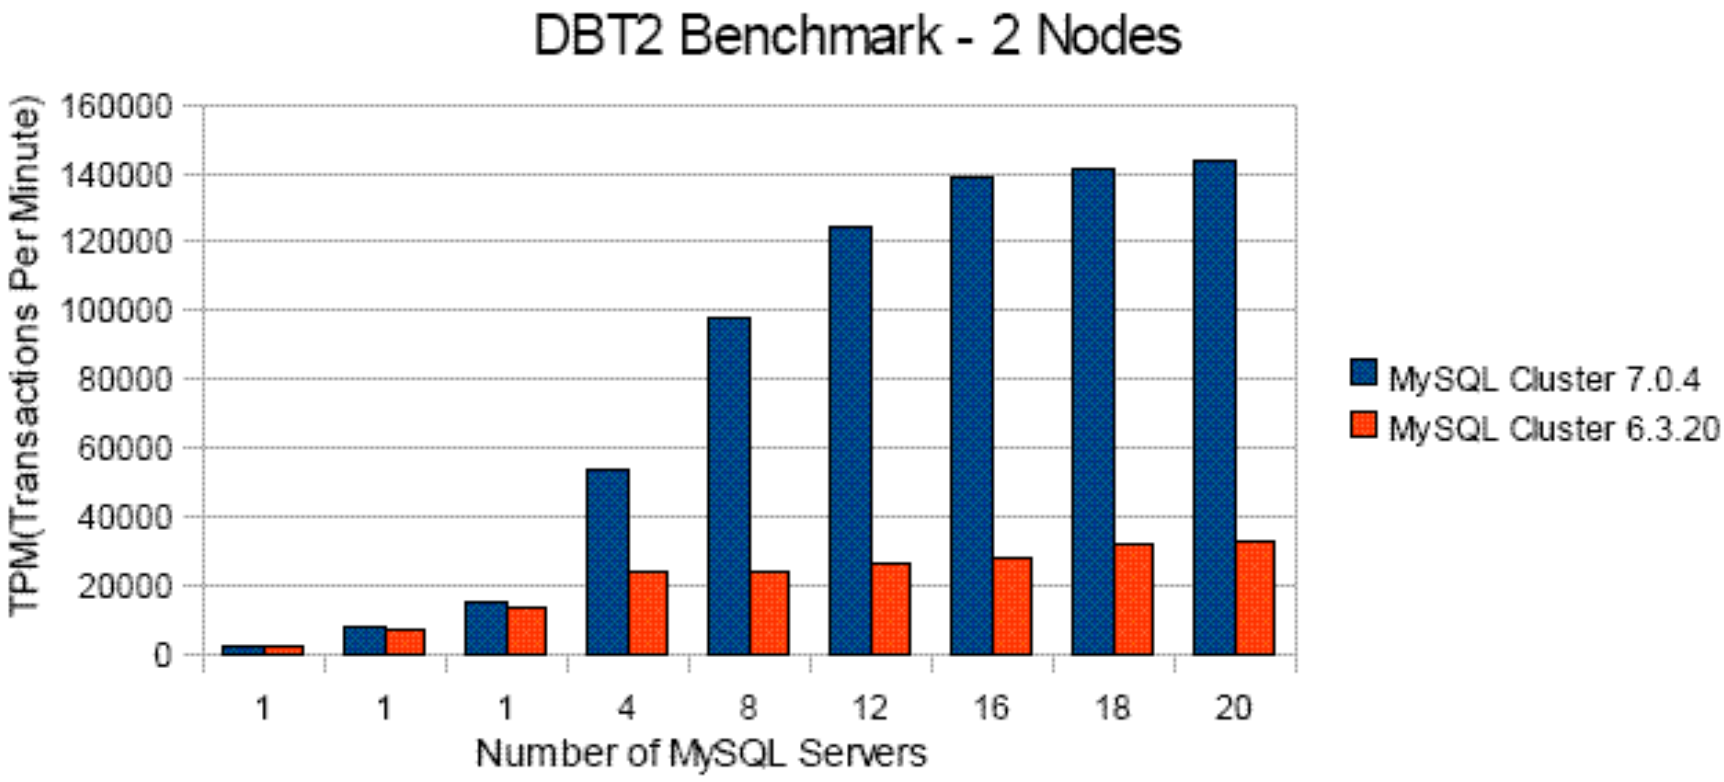

Databases as a CASH Cow

- This person was a DB Admin at 'Telco World', he was responsible for setting up the cluster at his company.
- Requirements: 33,000 queries per second.
- 4 node cluster (performance shown above) can do about 4000 queries per second.
- Estimated cost of 4 nodes about 10 lac.
- Required nodes for the task , about 32.
- So total cost of computers , whooping 1 crore !!!!

Users of Cluster

- **Nokia**, using MySQL Cluster to maintain real-time information about mobile network users.
- **Wikipedia**, more than 200 million queries and 1.2 million updates per day with peak loads of 11,000 queries per second.
- **Google**, for the search engine's AdWords program.
- **NASA**, converted an Oracle-based acquisition system to MySQL.
- **flickr**, Using MySQL in a Scale-out architecture to manage millions of photos and users.
- **Vodafone**, uses MySQL cluster for a range of internal Applications.
- **Slashdot** - with around 50 million page views per day.
- **Bredbandsbolaget**, largest ISP in Europe, uses MySQL Cluster for storage of customer data.

References

1. My love affair with MySQL Cluster, *Hingo*, *http* : *//openlife.cc/blogs/2009/june/my − love − affair − mysql − cluster − contains − benchmark − stories*.
2. MySQL Cluster Databases, *http* : *//en.wikipedia.org/wiki/MySQL_Cluster*.
3. Setting Up MySQL Cluster on Linux, *http* : *//dev.mysql.com/tech − resources/articles/mysql − cluster − for − two − servers.html*.
4. Cassandra, *http* : *//cassandra.apache.org/*.
5. Cluster Database Concept, *http* : *//www.stanford.edu/dept/itss/docs/oracle/10g/server*.101/*b*10739/*ds_concepts.htm*.
6. Database Management Systems, *Ramakrishnan and Gehrke, Mcgraw Hill*, pg. 725-738.
7. Prominent Users of MySQL Cluster, *http* : *//en.allexperts.com/e/m/my/mysql.htmhd*11

Thank - You :)