

# MySQL Cluster Tutorial

O'Reilly MySQL Conference & Expo 2010, Apr. 12 2010.

Andrew Hutchings (Oracle)

Andrew Morgan (Oracle)

Geert Vanderkelen (Oracle)

This document is a handout for the MySQL Cluster Tutorial. Please also check the slides which were shown during the tutorial.

<http://en.oreilly.com/mysql2010/public/schedule/detail/12438>

|  |           |
|--|-----------|
| <b>Introduction</b>                            | <b>5</b>  |
| Virtual Machine                                | 5         |
| Topics   | 5         |
| Speakers                                       | 6         |
| MySQL Cluster in short                         | 7         |
| Transactions                                   | 7         |
| <b>Installation and Configuration</b>          | <b>9</b>  |
| Release Model and Versioning                   | 9         |
| Download                                       | 9         |
| Installation                                   | 9         |
| Locations                                      | 10        |
| Configuration                                  | 11        |
| <b>Starting &amp; Using MySQL Cluster</b>      | <b>13</b> |
| Starting MySQL Cluster                         | 13        |
| Start Management Node(s)                       | 13        |
| Start Data Nodes                               | 14        |
| Start MySQL Server(s)                          | 14        |
| Create a Cluster table                         | 15        |
| Exercise:                                      | 16        |
| <b>Administer MySQL Cluster</b>                | <b>17</b> |
| Common Commands                                | 17        |
| Data Node Logs                                 | 20        |
| <b>MySQL Cluster Manager</b>                   | <b>23</b> |
| MySQL Cluster Manager – Architecture and Use   | 23        |
| MySQL Cluster Manager Model & Terms            | 25        |
| Using MySQL Cluster Manager – a worked example | 26        |
| Single host exercise                           | 33        |
| <b>Fault tolerance</b>                         | <b>34</b> |
| MySQL Server                                   | 34        |
| Heartbeats                                     | 34        |
| <b>Online Backup</b>                           | <b>36</b> |
| Tools  | 36        |
| Backing up the data, online                    | 36        |
| Backing up meta data                           | 36        |
| Restoring using ndb_restore                    | 37        |
| ndb_restore can do more                        | 38        |
| <b>NDB Info</b>                                | <b>39</b> |

|  |           |
|--|-----------|
| ndbinfo Data Node Statistics                       | 39        |
| ndbinfo.counters                                   | 39        |
| ndbinfo.logbuffers                                 | 40        |
| ndbinfo.logspaces                                  | 40        |
| ndbinfo.memoryusage                                | 41        |
| ndbinfo.nodes                                      | 41        |
| ndbinfo.transporters                               | 42        |
| Exercise   | 42        |
| <b>NDB API</b>                                     | <b>43</b> |
| NDB API Overview                                   | 43        |
| Example NDB API Code                               | 44        |
| <b>MySQL Cluster Connector for Java</b>            | <b>49</b> |
| Technical Overview                                 | 49        |
| ClusterJ   | 50        |
| ClusterJPA   | 52        |
| Pre-requisites for Tutorial                        | 54        |
| ClusterJ Tutorial                                  | 55        |
| Compiling and running the ClusterJ tutorial code   | 61        |
| OpenJPA/ClusterJPA Tutorial                        | 61        |
| Compiling and running the ClusterJPA tutorial code | 67        |
| Exercise   | 67        |
| <b>Schema considerations</b>                       | <b>68</b> |
| Develop for MySQL Cluster                          | 68        |
| Re-normalization                                   | 68        |
| Denormalization                                    | 69        |
| Primary Keys and Unique Indexes                    | 70        |
| Historical Data                                    | 70        |
| <b>Scaling and Performance</b>                     | <b>71</b> |
| MySQL Nodes  | 71        |
| NDBAPI   | 71        |
| Data Nodes   | 71        |
| Other Issues                                       | 71        |
| Online Add Node                                    | 72        |
| <b>Geographical Replication</b>                    | <b>73</b> |
| Binary Log Format                                  | 73        |
| Enabling Binary Logging                            | 74        |
| The LOST_EVENT incident & solution                 | 75        |
| Setting up Replication between Clusters            | 76        |

|                               |           |
|-------------------------------|-----------|
| Handling LOST_EVENTS          | 77        |
| Switching Replication Channel | 78        |
| <b>Security</b>               | <b>81</b> |
| MySQL Authentication          | 81        |

# Introduction

MySQL Cluster is a tool which could help make your data Highly Available. This tutorial will help you run a MySQL Cluster, show how to manage it and discuss various topics such as performance, backups and schema considerations.

Before going any further we need to setup the Virtual Machine (VM) running under VirtualBox. You can install MySQL Cluster yourself following instructions found in section **Installation and Configuration**, but we strongly suggest to stick to the filesystem layout and configuration files (found on the DVD).

## Virtual Machine

You have been given a DVD which contains VirtualBox and a Virtual Machine. The VM will boot Ubuntu (Linux Distribution) with all software pre-installed and configured.

To get you going, do the following:

1. Mount or open the DVD
2. Install (or upgrade) VirtualBox. The latest version is included on the DVD in the folder software/.
3. Copy the `clustervm/` and `config/` folder to your hard drive. Location does not matter, but make sure you copy the complete folder and all its contents.
4. Start VirtualBox: from the File-menu choose 'Import Appliance'
5. The 'Appliance Wizard' will show. Locate the `Ubuntu 9.10.ovf` file you copied from the DVD and follow the steps. No options should be changed.

## Topics

### ***Installation and Configuration***

What to download, how to install and configure MySQL Cluster.

### ***Running Nodes and Your First Table***

Starting MySQL Cluster and creating your first NDB table.

### ***Administer MySQL Cluster***

Managing and monitoring MySQL Cluster.

### ***MySQL Cluster Manager***

We'll introduce a new tool to manage MySQL Cluster.

### ***Fault Tolerance***

Explains what happens when some node fails.

### ***Online Backup***

How to backup your data and meta data.

### ***NDB Info***

Getting information out of MySQL Cluster made easy.

### ***NDBAPI***

Coding for Cluster using NDB API, and 'No SQL'.

### ***MySQL Cluster Connector for Java***

Introduction and talking to Cluster directly using Java.

### ***Schema Considerations***

A few tips when planning to develop for or convert to MySQL Cluster

### ***Scaling and Performance***

How you can scale and get more performance.

### ***Geographical Replication***

Making your MySQL Cluster itself highly available.

### ***Security***

Discusses how you can secure your MySQL Cluster

## **Speakers**

### ***Andrew Hutchings***

MySQL Support Engineer

“Andrew Hutchings is a MySQL Support Engineer for Oracle Corporation specialising in MySQL Cluster and C/C++ APIs. He is based in the United Kingdom and has worked for MySQL/Sun since 2008. Before joining Sun he was the Technical Architect, Senior Developer and DBA for a major UK magazine publisher. In his spare time Andrew develops various bug fixes and features for MySQL and MySQL Cluster.”

### ***Andrew Morgan***

MySQL Product Manager

“Andrew is the MySQL Product Manager responsible for High Availability Solutions – in particular MySQL Cluster and replication. He is based in United Kingdom and has worked for MySQL/Sun/Oracle since February 2009. Before joining MySQL he was responsible for delivering High Availability telecoms applications which is where he became exposed to MySQL Cluster – replacing proprietary and other 3rd party databases. His primary roles in MySQL are working with engineers to make sure that MySQL Cluster & replication evolve to meet the needs of their users as well as spreading the word on the what people can get from these technologies.”

### ***Geert Vanderkelen***

MySQL Support Engineer

“Geert is a member of the MySQL Support Team at Sun Microsystems. He is based in Germany and has worked for MySQL AB since April, 2005. Before joining MySQL he worked as developer, DBA and SysAdmin for various companies in Belgium and Germany. Today Geert specializes in MySQL Cluster and works together with colleagues around the world to ensure continued support for both customers and community. He’s also the maintainer of Sun’s MySQL Connector/Python.”

**This is a sample, click download link to get the full Tutorial**

**CLICK BELOW**

