



Linux High Availability Clustering Red hat / CentOS / Ubuntu

Description

A computer cluster is a single logical unit consisting of multiple computers that are linked through a LAN or WAN. The networked computers essentially act as a single, much more powerful machine. A computer cluster provides much faster processing speed, larger storage capacity, better data integrity, superior reliability and wider availability of resources.

Organizations often use computer clusters to maximize processing time, increase database storage and implement faster data storing & retrieving techniques.

There are many types of computer clusters, including:

- Load-balancing clusters
- High availability (HA) clusters
- High performance (HP) clusters

The major advantages of using computer clusters are clear when an organization requires large scale processing. When used this way, computer clusters offer:

- **Cost efficiency:** The cluster technique is cost effective for the amount of power and processing speed being produced. It is more efficient and much cheaper compared to other solutions like setting up mainframe computers.
- **Processing speed:** Multiple high-speed computers work together to provide unified processing, and thus faster processing overall.
- **Improved network infrastructure:** Different LAN or WAN topologies are implemented to form a computer cluster. These networks create a highly efficient and effective infrastructure that prevents bottlenecks.
- **Flexibility:** Unlike mainframe computers, computer clusters can be upgraded to enhance the existing specifications or add extra components to the system.
- **High availability of resources:** If any single component fails in a computer cluster, the other machines continue to provide uninterrupted processing. This redundancy is lacking in mainframe systems.



DRBD is a distributed replicated storage system for the Linux platform. It is implemented as a kernel driver, several user space management applications, and some shell scripts. DRBD is traditionally used in high availability (HA) computer clusters, but beginning with DRBD version 9, it can also be used to create larger software defined storage pools with a focus on cloud integration

Audience

This Course is designed for people who have experience with Linux or UNIX, System administrators, developers, architects, decision makers can all benefit from the content covered in this class, especially if they are looking to work with High availability & Redundancy cluster computing.

Duration: 16 hours

Requirements

This Course includes hands on labs training. Users should have access to physical or virtual environment with RHEL 7.x / CentOS 7.x / Ubuntu 16.x installation. Or Internet access to Tracston labs.



Module 1 – Day 1

- **What is Linux Cluster?**
 - Introduction
 - Overview
 - High Availability and Redundancy
 - Popular Linux Cluster software
- **Hardware Requirements**
- **Network Considerations**
- **Linux Cluster Pacemaker & Corosync**
 - Overview
 - Cluster Installation
 - Cluster software structure
 - Management
- **Shared Storage**
 - Overview
 - Replicated Storage
 - Storage area network (SAN)
 - Network attached storage (NAS)
 - Storage server
 - Cloud storage
- **DRBD Filesystem (Distributed Replicated Block Device)**
 - Installing DRBD (Ubuntu/CentOS)
 - Configuring DRBD
 - Making your data
 - Configuring DNS
- **Utilities**
 - HA Components
 - Clustered Resources
 - Configuration Tools
 - Scripting
- **Summary**



Module 2 – Day 2

- **Advanced Linux Cluster Pacemaker & Corosync**
 - Split Brain Prevention with Fencing
 - Automating Cluster Tasks
 - Cluster Security
 - Troubleshooting
- **Linux Cluster Keepalived Cluster**
 - Overview
 - Cluster Installation
 - Cluster software structure
 - Management
- **Advanced Linux Cluster keepalived**
 - Split Brain Prevention with Fencing
 - User Environment
 - Troubleshooting
- **Security Principles**
- **Monitoring**
- **System Performance**
- **Cluster Scripting**
- **Development Tools**
 - SDK
 - IDE tools
 - CLI
- **Summary**