

Course Description



Course: SLE321v15 Deploying and Administering SUSE Linux Enterprise High Availability 15



Training Level:

- ❑ Advanced

Duration:

- ❑ 3 days

Course Overview

This three day course teaches students to deploy and administer the SUSE Linux Enterprise High Availability 15 product. The course introduces the product features and implementation planning for both on premise and the public cloud. A SLE15 HA cluster is deployed and then applications and clustered storage are configured, tested and administered.

Key Objectives

During this course you will learn to:

- ❑ Plan and prepare to implement a SLE HA cluster
- ❑ Install SLE HA and create a Corosync/Pacemaker cluster
- ❑ Administer a cluster using Web and CLI tools
- ❑ Configure node fencing using the BMCs and SBD
- ❑ Cluster basic resources, such as IP addresses and services
- ❑ Cluster storage in a cluster safe way
- ❑ Configure the cluster behaviour using constraints
- ❑ Perform basic cluster troubleshooting

Audience

This course is designed for existing Linux administrators who want to configure highly available services using SUSE Linux Enterprise HA. This course provides a foundation for deploying SAP on SLE 15 HA.

Prerequisites

Students require a good knowledge of SLES15. Some familiarity with the basic concepts of clustering for HA would be useful but not required.



Course Outline

- ❑ Section 1: Introduction
 - ❑ Agenda
 - ❑ Course Overview
 - ❑ Lab Environment
- ❑ Section 2: Introduction to SUSE Linux Enterprise HA
 - ❑ Overview of SLE15 HA
 - ❑ The Cluster Story
 - ❑ Cluster Terminology
 - ❑ SLE15 HA Architecture
 - ❑ SLE HA in the Public Cloud
 - ❑ Overview of Public Cloud
 - ❑ Supported Public Cloud Environments
- ❑ Section 3: Cluster Implementation Requirements
 - ❑ Overview of SLE15 HA Implementation Process
 - ❑ Collecting the Required Parameters
 - ❑ Determine Expectations
 - ❑ Planning Storage
 - ❑ Differences between 2 node and 2+N node clusters
 - ❑ Differences between SLEHA12 and SLEHA15
 - ❑ Designing Test Cases
 - ❑ Testing and Documentation
- ❑ Section 4: Cluster Node Preparation
 - ❑ Local Filesystems
 - ❑ Preparing the Infrastructure
 - ❑ Cluster Networking
 - ❑ Bonding Configuration
 - ❑ Time Synchronization
 - ❑ Name Resolution
 - ❑ Name Resolution using /etc/hosts
 - ❑ Name Resolution using DNS
 - ❑ User Management Best Practices
 - ❑ Software Management for the HA Environment
 - ❑ Updating Cluster Nodes
- ❑ Section 5: Deploy a Cluster
 - ❑ Preparation Checklist
 - ❑ Cluster setup step by step
 - ❑ Secure Cluster Communications
- ❑ Section 6: Administer the HA Cluster
 - ❑ Overview of the Cluster Administration Tools
 - ❑ Introduction to Hawk
 - ❑ Command Line Tools
 - ❑ Configure and Synchronize files with csync2
- ❑ Section 7: Split Brain Avoidance
 - ❑ Understand Split Brain
 - ❑ Understand Fencing
 - ❑ Understand STONITH
 - ❑ Fencing using BMCs
 - ❑ Fencing using SBD
 - ❑ Split brain Avoidance in 2 node vs 2+N node clusters

SUSE Training

Information about SUSE Training can be found at:

<https://training.suse.com>



Contact suse-training@suse.com with any questions.



- Fencing in the Public Cloud
- Section 8: Cluster Resources
 - Introduction to Cluster Resources
 - Resource Types
 - Resource Agents
- Section 9: Controlling Cluster Behavior using Constraints
 - Introduction to Constraints
 - Location Constraints
 - Order Constraints
 - Colocation Constraints
- Section 10: Cluster Managed Storage
 - Overview of Storage for a HA Cluster
 - Cluster Managed Storage
 - Overview of cLVM
 - Overview of OCFS2
 - Overview of Cluster Managed MDRAID
 - Overview of DRBD
 - Cluster Managed NFS
- Section 11: Access Control Lists
 - Using Access Control Lists
- Section 12: Basic HA Troubleshooting
 - General Troubleshooting
 - CLI Troubleshooting Tools
 - Log Files
 - Updating the Epoch
 - SBD startup behaviour
 - Working with the CIB

