

Red Hat Cluster Tutorial

Article Number: 211 | Rating: Unrated | Last Updated: Sun, Jun 3, 2018 8:58 PM

Contents

- 1 What's New?
 - 1.1 A Note on Terminology
 - 1.2 Why Should I Follow This (Lengthy) Tutorial?
 - 1.3 High-Level Explanation of How HA Clustering Works
- 2 The Task Ahead
 - 2.1 A Note on Patience
 - 2.2 Technologies We Will Use
 - 2.3 A Note on Hardware
 - 2.4 System Requirements
 - 2.5 Recommended Hardware; A Little More Detail
 - 2.5.1 The Most Important Consideration - Storage
 - 2.5.1.1 Extra Security - LSI SafeStore
 - 2.5.2 RAM - Preparing for Degradation
 - 2.5.3 Never Over-Provision!
 - 2.5.4 CPU Cores - Possibly Acceptable Over-Provisioning
 - 2.5.4.1 A Note on Hyper-Threading
 - 2.5.5 Six Network Interfaces, Seriously?
 - 2.5.5.1 A Note on Dedicated IPMI Interfaces
 - 2.5.6 Network Switches
 - 2.5.7 Why Switched PDUs?
 - 2.5.8 Network Managed UPSes Are Worth It
 - 2.5.9 Dashboard Servers
 - 2.6 What You Should Know Before Beginning
 - 2.7 A Word on Complexity
- 3 Overview of Components

- 3.1 Component; Cman
- 3.2 Component; Corosync
 - 3.2.1 A Little History
 - 3.2.2 The Future of Corosync
- 3.3 Concept; Quorum
- 3.4 Concept; Virtual Synchrony
- 3.5 Concept; Fencing
 - 3.5.1 Is "Fencing" the same as STONITH?
- 3.6 Component; Totem
- 3.7 Component; Rgmanager
 - 3.7.1 What about Pacemaker?
- 3.8 Component; Qdisk
- 3.9 Component; DRBD
- 3.10 Component; DLM
- 3.11 Component; Clustered LVM
- 3.12 Component; GFS2
- 3.13 Component; KVM
- 4 Node Installation
 - 4.1 Node Host Names
 - 4.2 Foundation Pack Host Names
 - 4.3 OS Installation
 - 4.4 Network Security Considerations
 - 4.5 SELinux Considerations
- 5 Network
 - 5.1 A Map!
 - 5.2 Subnets
 - 5.2.1 A Note on STP
 - 5.3 Setting Up the Network
 - 5.3.1 Planning The Use of Physical Interfaces
 - 5.3.2 Connecting Fence Devices
- 6 Let's Build!
 - 6.1 Why so Much Duplication of Commands?
 - 6.2 Red Hat Enterprise Linux Specific Steps
 - 6.3 Add the Alteeve's Niche! Repo
 - 6.4 Update the OS
 - 6.5 Installing Required Programs
 - 6.6 Switch Network Daemons
 - 6.7 Altering Which Daemons Start on Boot

- 6.8 Network Security
- 6.9 Configuring iptables
- 6.10 Mapping Physical Network Interfaces to ethX Device Names
 - 6.10.1 Making Sure All Network Interfaces are Started
 - 6.10.2 Finding Current Names for Physical Interfaces
 - 6.10.3 Building the MAC Address List
 - 6.10.4 Changing the Interface Device Names
 - 6.10.5 Test the New Network Name Mapping
- 6.11 Configuring our Bridge, Bonds and Interfaces
 - 6.11.1 Creating New Network Configuration Files
 - 6.11.2 Configuring the Bridge
 - 6.11.3 Creating the Bonded Interfaces
 - 6.11.4 Alter the Interface Configurations
- 6.12 Loading the New Network Configuration
 - 6.12.1 Verifying the New Network Config
- 6.13 Adding Everything to /etc/hosts
- 6.14 What is IPMI
 - 6.14.1 Reading IPMI Data
 - 6.14.2 Finding our IPMI LAN Channel
 - 6.14.3 Reading IPMI Network Info
 - 6.14.4 Find the IPMI User ID
 - 6.14.5 Testing the IPMI Connection From the Peer
- 6.15 Setting up SSH
 - 6.15.1 Create the RSA Keys
 - 6.15.2 Populate known_hosts
 - 6.15.3 Copy Public Keys to Enable SSH Without a Password
- 6.16 Setting Up UPS Monitoring
 - 6.16.1 Installing apcupsd
 - 6.16.2 Configuring Apcupsd For Two UPSes
 - 6.16.3 SELinux and apcupsd
 - 6.16.4 Testing the Multi-UPS apcupds
- 6.17 Monitoring Storage
 - 6.17.1 Monitoring LSI-Based RAID Controllers with MegaCli
 - 6.17.1.1 Installing MegaCli
 - 6.17.1.2 Checking Storage Health with MegaCli64
 - 6.17.1.3 Managing MegaSAS.log
- 7 Configuring The Cluster Foundation
 - 7.1 Keeping Time in Sync

- 7.2 Alternate Configuration Methods
- 7.3 The First cluster.conf Foundation Configuration
 - 7.3.1 Name the Cluster and Set the Configuration Version
 - 7.3.2 Configuring cman Options
 - 7.3.3 Defining Cluster Nodes
 - 7.3.4 Defining Fence Devices
 - 7.3.5 Using the Fence Devices
 - 7.3.6 Giving Nodes More Time to Start and Avoiding "Fence Loops"
 - 7.3.7 Configuring Totem
 - 7.3.8 Validating and Pushing the /etc/cluster/cluster.conf File
 - 7.3.9 Setting up ricci
 - 7.3.10 Starting the Cluster for the First Time
- 7.4 Testing Fencing
 - 7.4.1 Using Fence_check to Verify our Fencing Config
 - 7.4.2 Crashing an-a05n01 for the First Time
 - 7.4.3 Cutting the Power to an-a05n01
 - 7.4.4 Hanging an-a05n02
 - 7.4.5 Cutting the Power to an-a05n02
- 8 Installing DRBD
 - 8.1 Option 1 - Fully Supported by Red Hat and Linbit
 - 8.2 Option 2 - Install From AN!Repo
 - 8.3 Option 3 - Install From Source
 - 8.3.1 Hooking DRBD into the Cluster's Fencing
 - 8.3.2 The "Why" of our Layout - More Safety!
 - 8.4 Creating The Partitions For DRBD
 - 8.4.1 Block Alignment
 - 8.4.2 Determining Storage Pool Sizes
 - 8.4.3 Creating the DRBD Partitions
 - 8.5 Configuring DRBD
 - 8.5.1 Configuring DRBD Global and Common Options
 - 8.5.2 Configuring the DRBD Resources
 - 8.6 Initializing the DRBD Resources
 - 8.7 Loading the drbd Kernel Module
- 9 Initializing Clustered Storage
 - 9.1 Clustered Logical Volume Management
 - 9.1.1 Configuring Clustered LVM Locking
 - 9.1.2 Testing the clvmd Daemon
 - 9.1.3 Initialize our DRBD Resource for use as LVM PVs

- 9.1.4 Creating Cluster Volume Groups
 - 9.1.5 Creating a Logical Volume
 - 9.2 Creating the Shared GFS2 Partition
 - 9.3 Adding /shared to /etc/fstab
 - 9.3.1 Stopping All Clustered Storage Components
- 10 Managing Storage In The Cluster
 - 10.1 A Note on Daemon Starting
 - 10.1.1 Defining the Resources
 - 10.1.2 Creating Failover Domains
 - 10.1.3 Creating Clustered Storage and libvirtd Service
 - 10.2 Validating and Pushing the Changes
 - 10.3 Checking the Cluster's Status
 - 10.4 Managing Cluster Resources
 - 10.5 Stopping Clustered Storage - A Preview to Cold-Stopping the Cluster
 - 10.6 Starting Clustered Storage
- 11 Testing Network Redundancy
 - 11.1 What we will be Watching
 - 11.1.1 Understanding '/proc/net/bonding/{bcn,sn,ifn}_bond1'
 - 11.1.2 Understanding '/etc/init.d/drbd status'
 - 11.1.3 Understanding 'cman_tool nodes'
 - 11.2 Network Testing Terminal Layout
 - 11.3 How to Know if the Tests Passed
 - 11.4 Breaking things!
 - 11.4.1 Failing a Bond's Primary Interface
 - 11.4.2 Failing the Network Switches
- 12 Provisioning Virtual Machines
 - 12.1 Before We Begin - Building a Dashboard
 - 12.2 A Note on the Following Server Installations
 - 12.3 Provision Planning
 - 12.4 Provisioning vm01-win2008
 - 12.4.1 Creating vm01-win2008's Storage
 - 12.4.2 Creating vm01-win2008's virt-install Call
 - 12.4.3 Initializing vm01-win2008's Install
 - 12.5 Provisioning vm02-win2012
 - 12.5.1 Creating vm02-win2012's Storage
 - 12.5.2 Creating vm02-win2012's virt-install Call
 - 12.5.3 Initializing vm02-win2012's Install
 - 12.6 Provisioning vm03-win7

- 12.6.1 Creating vm03-win7's Storage
 - 12.6.2 Creating vm03-win7's virt-install Call
 - 12.6.3 Initializing vm03-win7's Install
 - 12.7 Provisioning vm04-win8
 - 12.7.1 Creating vm04-win8's Storage
 - 12.7.2 Creating vm04-win8's virt-install Call
 - 12.7.3 Initializing vm04-win8's Install
 - 12.8 Provisioning vm05-freebsd9
 - 12.8.1 Creating vm05-freebsd9's Storage
 - 12.8.2 Creating vm05-freebsd9's virt-install Call
 - 12.8.3 Initializing vm05-freebsd9's Install
 - 12.9 Provisioning vm06-solaris11
 - 12.9.1 Creating vm06-solaris11's Storage
 - 12.9.2 Calculating Free Space; Converting GiB to MB
 - 12.9.3 Creating vm06-solaris11's virt-install Call
 - 12.9.4 Initializing vm06-solaris11's Install
 - 12.10 Provisioning vm07-rhel6
 - 12.10.1 Creating vm07-rhel6's Storage
 - 12.10.2 Creating vm07-rhel6's virt-install Call
 - 12.10.3 Initializing vm07-rhel6's Install
 - 12.10.4 Making sure RHEL 6 reboots after panic'ing
 - 12.11 Provisioning vm08-sles11
 - 12.11.1 Creating vm08-sles11's Storage
 - 12.11.2 Creating vm08-sles11's virt-install Call
 - 12.11.3 Initializing vm08-sles11's Install
- 13 Making Our VMs Highly Available Cluster Services
 - 13.1 Creating the Ordered Fail-Over Domains
 - 13.2 Making vm01-win2008 a Highly Available Service
 - 13.2.1 Dumping the vm01-win2008 XML Definition File
 - 13.2.2 Creating the vm:vm01-win2008 Service
 - 13.2.3 Testing vm01-win2008 Management With clusvcadm
 - 13.2.4 Solving vm01-win2008 "Failure to Enable" Error
 - 13.2.5 Testing vm01-win2008 Live Migration
 - 13.3 Making vm02-win2012 a Highly Available Service
 - 13.3.1 Dumping the vm02-win2012 XML Definition File
 - 13.3.2 Creating the vm:vm02-win2012 Service
 - 13.3.3 Testing vm02-win2012 Management With clusvcadm
 - 13.4 Making vm03-win7 a Highly Available Service

- 13.4.1 Dumping the vm03-win7 XML Definition File
 - 13.4.2 Creating the vm:vm03-win7 Service
 - 13.4.3 Testing vm03-win7 Management With clusvcadm
 - 13.5 Making vm04-win8 a Highly Available Service
 - 13.5.1 Dumping the vm04-win8 XML Definition File
 - 13.5.2 Creating the vm:vm04-win8 Service
 - 13.5.3 Testing vm04-win8 Management With clusvcadm
 - 13.6 Making vm05-freebsd9 a Highly Available Service
 - 13.6.1 Dumping the vm05-freebsd9 XML Definition File
 - 13.6.2 Creating the vm:vm05-freebsd9 Service
 - 13.6.3 Testing vm05-freebsd9 Management With clusvcadm
 - 13.7 Making vm06-solaris11 a Highly Available Service
 - 13.7.1 Dumping the vm06-solaris11 XML Definition File
 - 13.7.2 Creating the vm:vm06-solaris11 Service
 - 13.7.3 Testing vm06-solaris11 Management With clusvcadm
 - 13.8 Making vm07-rhel6 a Highly Available Service
 - 13.8.1 Dumping the vm07-rhel6 XML Definition File
 - 13.8.2 Creating the vm:vm07-rhel6 Service
 - 13.8.3 Testing vm07-rhel6 Management With clusvcadm
 - 13.9 Making vm08-sles11 a Highly Available Service
 - 13.9.1 Dumping the vm08-sles11 XML Definition File
 - 13.9.2 Creating the vm:vm08-sles11 Service
 - 13.9.3 Testing vm08-sles11 Management With clusvcadm
- 14 Setting Up Alerts
 - 14.1 Alert System Overview
 - 14.2 AN!CM Requirements
 - 14.3 Setting Up Your Dashboard
 - 14.4 Testing Monitoring
 - 14.5 Enabling Monitoring
 - 14.6 We're Done! or are We?
- 15 Testing Server Recovery
 - 15.1 Controlled Migration and Node Withdrawal
 - 15.1.1 Withdraw an-a05n01
 - 15.1.2 Load Testing in a Degraded State
 - 15.1.3 Rejoin an-a05n01
 - 15.1.4 Withdraw an-a05n02
 - 15.1.5 Rejoin an-a05n02
 - 15.2 Out-of-Cluster Server Power-off

- 15.3 Crashing Nodes; The Ultimate Test
 - 15.3.1 Crashing an-a05n01
 - 15.3.2 Degraded Mode Load Testing
 - 15.3.3 Recovering an-a05n01
 - 15.3.4 Crashing an-a05n02
 - 15.3.5 Recovering an-a05n02
- 15.4 Done and Done!
- 16 Troubleshooting
 - 16.1 SELinux Related Problems
 - 16.1.1 Password-less SSH doesn't work, but ~/.ssh/authorized_keys is fine
 - 16.1.2 Live-Migration fails with '[vm] error: Unable to read from monitor: Connection reset by peer'
 - 16.1.3 Attempting to Live Migrate Fails with 'Host key verification failed.'
 - 16.2 Other Tutorials
 - 16.3 Older Issues From Previous Tutorials

Posted - Sun, Jun 3, 2018 8:58 PM. This article has been viewed 2758 times.

Online URL: <http://kb.ictbanking.net/article.php?id=211>