

# How To Make Databases on Linux on System z Highly Available

Mike Friesenegger (mikef@suse.com)  
SUSE

Wednesday, August 14, 2013  
Session Number 13437



# Agenda

- Clarify the term "Availability"
- What is High Availability
- Minimize a Database Startup Time after a Failure
- Making an IBM DB2 Database Highly Available
- Making an Oracle Database Highly Available
- Demo

# Questions to ask about the importance of a database

- How much uptime does the business expect?
  - 24 x 7 x 365
  - Occasional downtime acceptable
  - Not important
- How quickly does the database need to recover from a failure?
  - Recovery time is unacceptable
  - Some recovery time is acceptable
  - Not important

# Determine the level of “Availability”

- Maximum<sup>1</sup> or Continuous<sup>2</sup> Availability
  - 24 x 7 x 365
  - Recovery time is unacceptable
- High Availability
  - Occasional downtime acceptable
  - Some recovery time is acceptable
- No Availability
  - Not important

<sup>1</sup> “High Availability Overview”, [http://www.oracle.com/pls/db112/portal.portal\\_db?selected=14&frame=](http://www.oracle.com/pls/db112/portal.portal_db?selected=14&frame=)

<sup>2</sup> “High Availability and Disaster Recovery Options for DB2 for Linux, UNIX, and Windows”, <http://www.redbooks.ibm.com/abstracts/sg247363.html>

# The cost of “Availability”

- Maximum or Continuous Availability
  - Additional hardware/software costs
  - Most complex configuration
  - Largest amount of time to plan and implement
- High Availability
  - Additional hardware/software costs
  - Complex configuration
  - Time required to plan and implement
- No Availability
  - Normal cost of running a single instance database



# The products that assist with “Availability”



- Maximum or Continuous Availability
  - Oracle Real Application Cluster (RAC)
  - IBM DB2 pureScale®
- High Availability
  - An OS based HA Clustering solution
    - IBM PowerHA for AIX
    - Microsoft Cluster Server for Windows
    - Veritas Cluster Server
    - SUSE Linux Enterprise High Availability Extension
  - Optional components
    - Database replication
    - Fast Failover capability



# A note about the “Availability” products

- Maximum or Continuous Availability
  - Availability is only one aspect of these products
  - Flexibility to scale and increase processing capacity is the biggest goal of these products
  - Automatically load balance application connections is a benefit of the scalability
  - May address disaster recovery
- High Availability
  - Primary focus on increasing availability
  - May address disaster recovery

# Clarify the term “Availability”

- Depends on your business requirements
  - Maybe you do not need any availability
- If you do
  - Maximum or Continuous Availability provides the most options and features with the most complexity and cost
  - High Availability may be all you need!
- The remainder of this session will focus on providing High Availability
  - Focused on Oracle and DB2 databases
  - Using the SUSE Linux Enterprise High Availability Extension



# SUSE® Linux Enterprise High Availability Extension<sup>1</sup>




- Service availability 24/7
  - Policy driven clustering
- Shared and Scaled data-access
  - Cluster file system
- Stretch & Geo<sup>2</sup> Clustering
  - Cluster across unlimited distance
- Virtualization Agnostic
  - Platform independent setup
- Disaster tolerance
  - Data replication via IP
  - Node recovery
- Scale network services
  - IP load-balancing
- User friendly tools
  - Graphical user interface
  - Unified command line interface
- Free Resource Agents
  - Oracle

<sup>1</sup> The SLEHA Extension is free with a purchased SLES for System z subscription


<sup>2</sup> Geo Clustering is a purchasable feature

# What is High Availability


# A Simple HA Cluster




# An Active/Passive Database Cluster




# Migrating a Database to a New Node




# Migrated Database on a New Node



# Node Failure in the Cluster

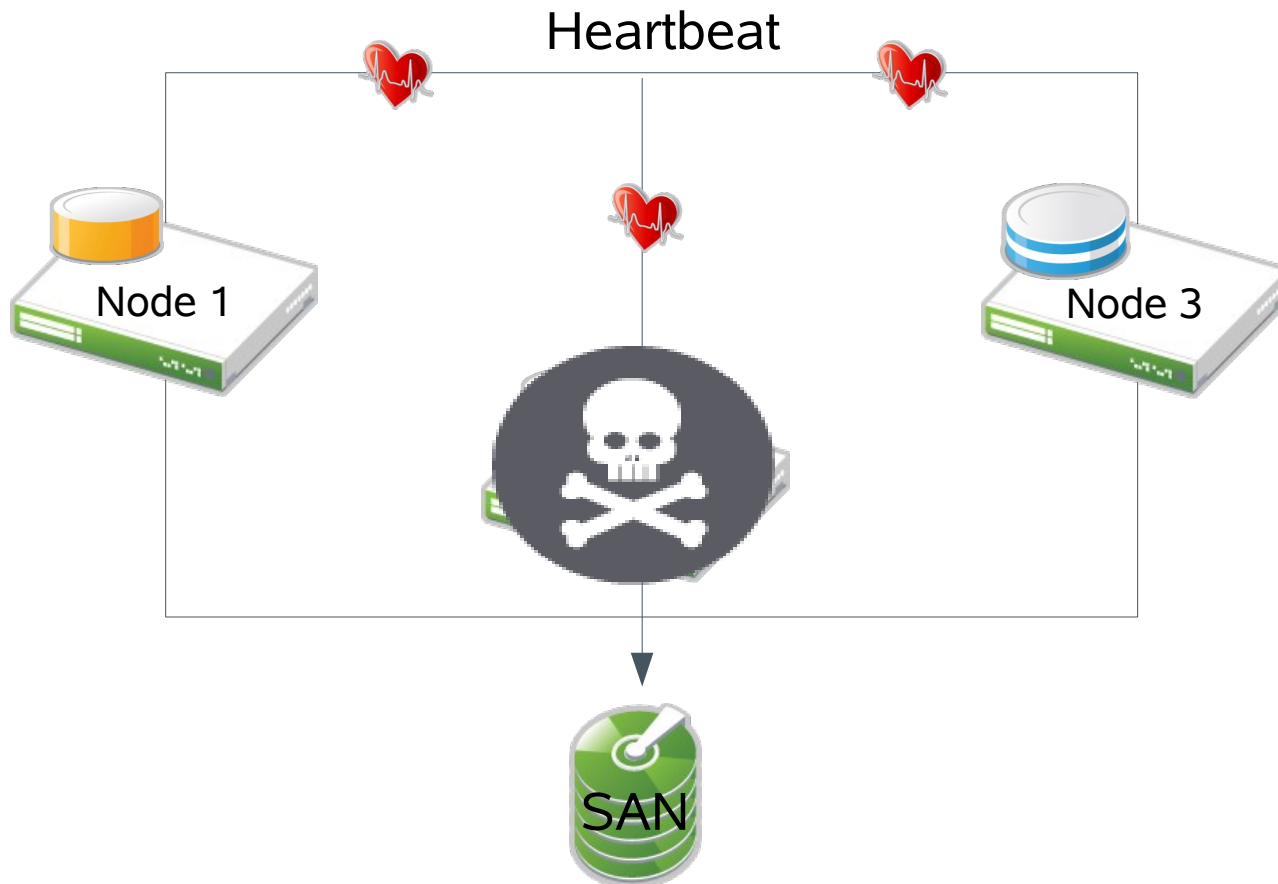


# Kick the Failed Node Out of the Cluster





# Database Automatically Restarted on another Node



# Minimize a Database Startup Time

# Migration and Failover can be Time Consuming


## Steps in migration

- Clean database shutdown
- Release shared storage from current node
- Attach shared storage to new node
- Start database and listener


## Steps for failover

- Detection of failed node
- Kick the failed node out of cluster
  - Issue node eviction
  - Wait for node eviction
  - Guarantee that node is kicked out
- Attach shared storage to new node
- Start database and listener
  - Verify no database corruption


# Replication to Standby Database via Log Shipping




# Node Failure in the Cluster



# Kick the Failed Node Out of the Cluster



# Standby Database Automatically Started on another Node



# Making an IBM DB2 Database Highly Available



# Components in a Highly Available IBM DB2 Database



- IBM DB2 10.1 for Linux
- IBM DB2 High Availability and Disaster Recovery (HADR)
  - Provides log shipping feature
  - Recommended to use a third-party clustering solution
- SUSE SLE High Availability Extension 11 SP3
  - Manages start/stop/failover of resources
    - Shared storage
    - Service IP Address
    - IBM DB2 Database and HADR
  - Provides node monitoring
  - IP address failover



# Making an Oracle Database Highly Available

# Components in a Highly Available Oracle Database



- Oracle Enterprise 11gR2 for Linux
- Oracle Data Guard
  - Provides log shipping feature
  - Oracle recommends using Oracle Clusterware
    - Other third-party HA clustering solutions will work
- SUSE SLE High Availability Extension 11 SP3
  - Manages start/stop/failover of resources
    - Shared storage
    - Service IP Address
    - Oracle Database and Data Guard
  - Provides node monitoring
  - IP address failover



# Thank You!!

