

# High Availability for Highly Reliable Systems

Mike Friesenegger  
SUSE

Wednesday, February 6, 2013  
Session: 12367



# Agenda

- What is a high availability (HA) cluster?
- What is required to build a HA cluster using SLES?
- Let's build a two node cluster...

# Challenge

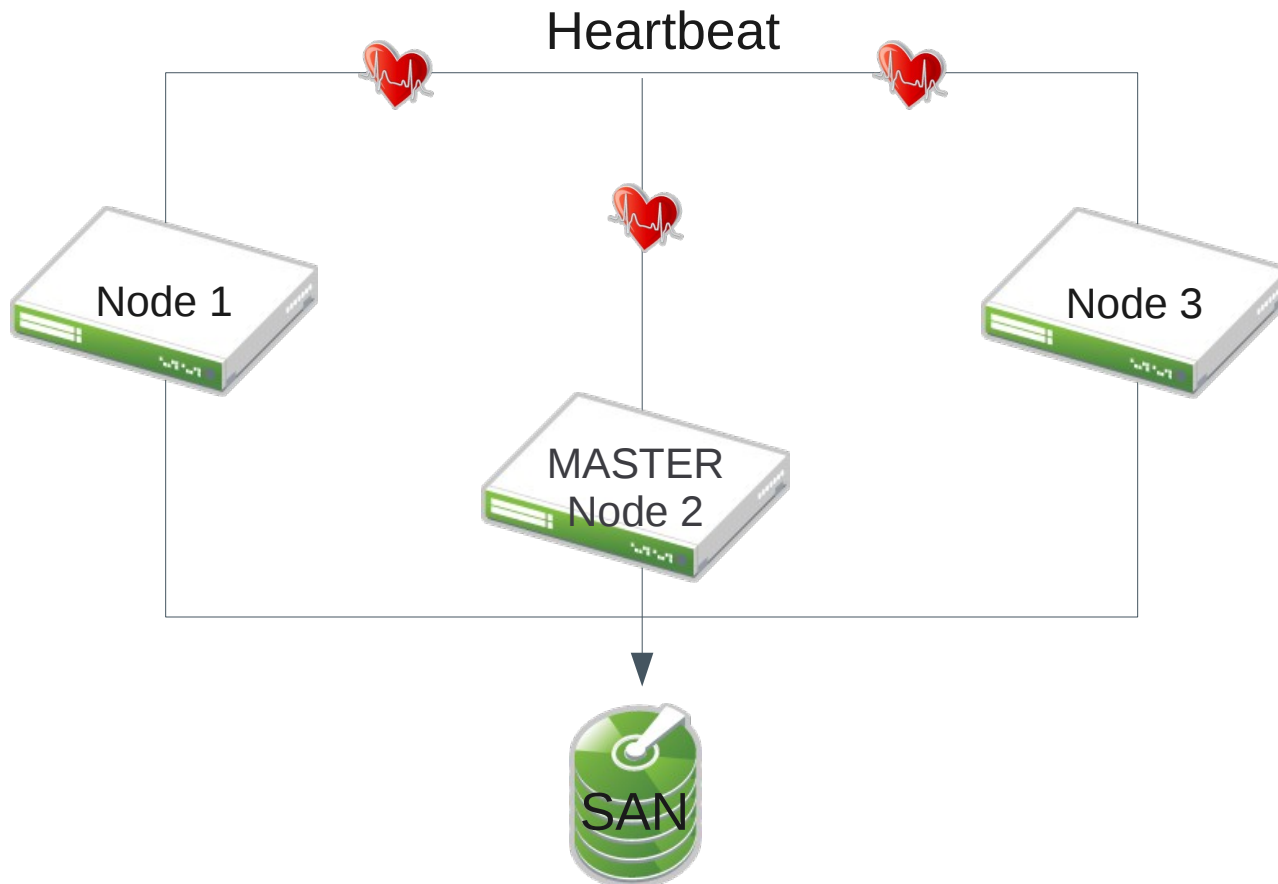
## Murphy's Law is Universal

- Faults will occur
  - Hardware crash, flood, fire, power outage, earthquake?
- Can you afford a service outage or worse, loss of data?
  - You might afford a five second blip, but can you afford a longer outage?
- How much does downtime cost?


**Can you afford low availability systems?**

# What is a high availability cluster?

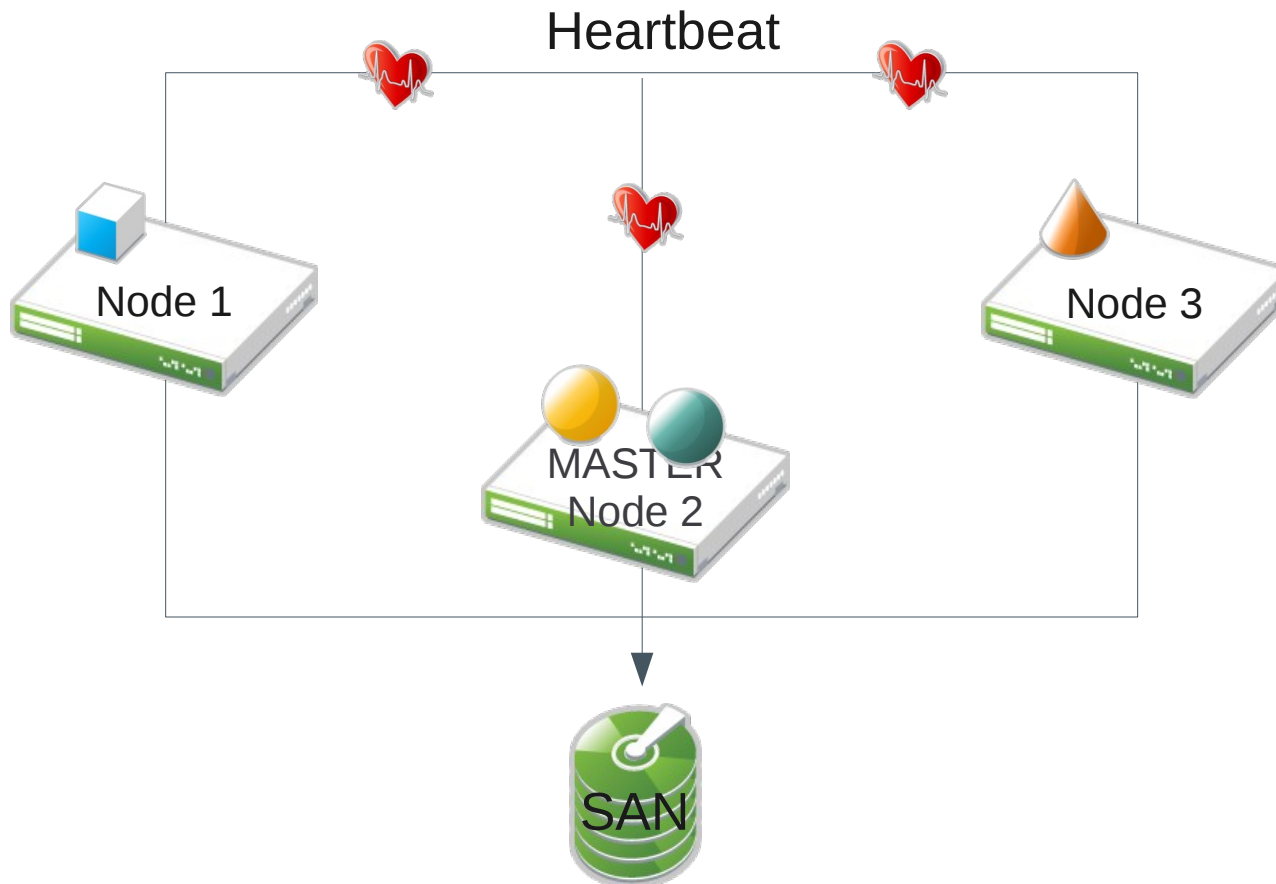
# A Simple HA Cluster




# Resources Running in the Cluster



# Migrating a Resource




# Node Failure in the Cluster

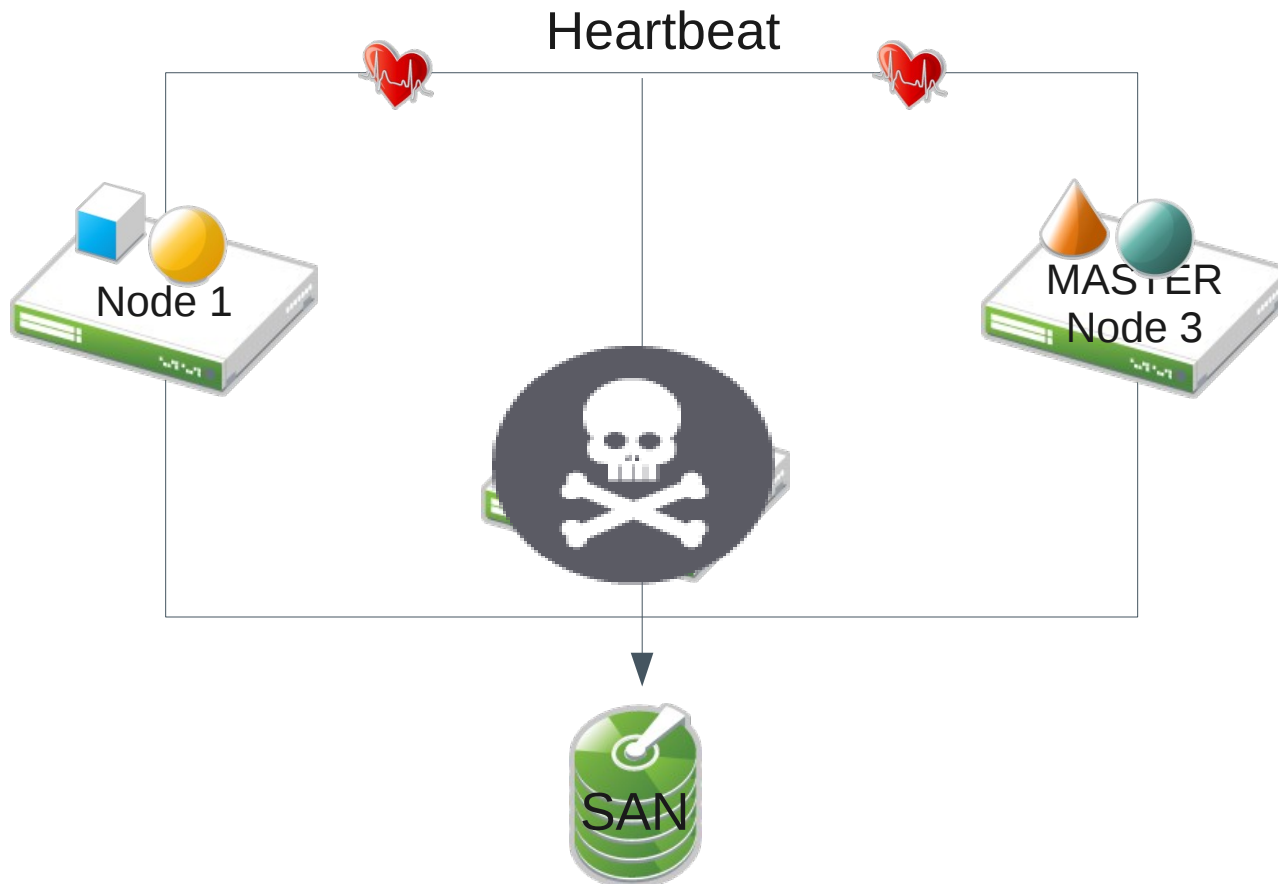




# STONITH the Failed Node Out of the Cluster



# Resources brought up on other nodes the Cluster



# What is required to build a HA cluster using SUSE Linux Enterprise Server?

# SUSE Linux Enterprise High Availability Extension



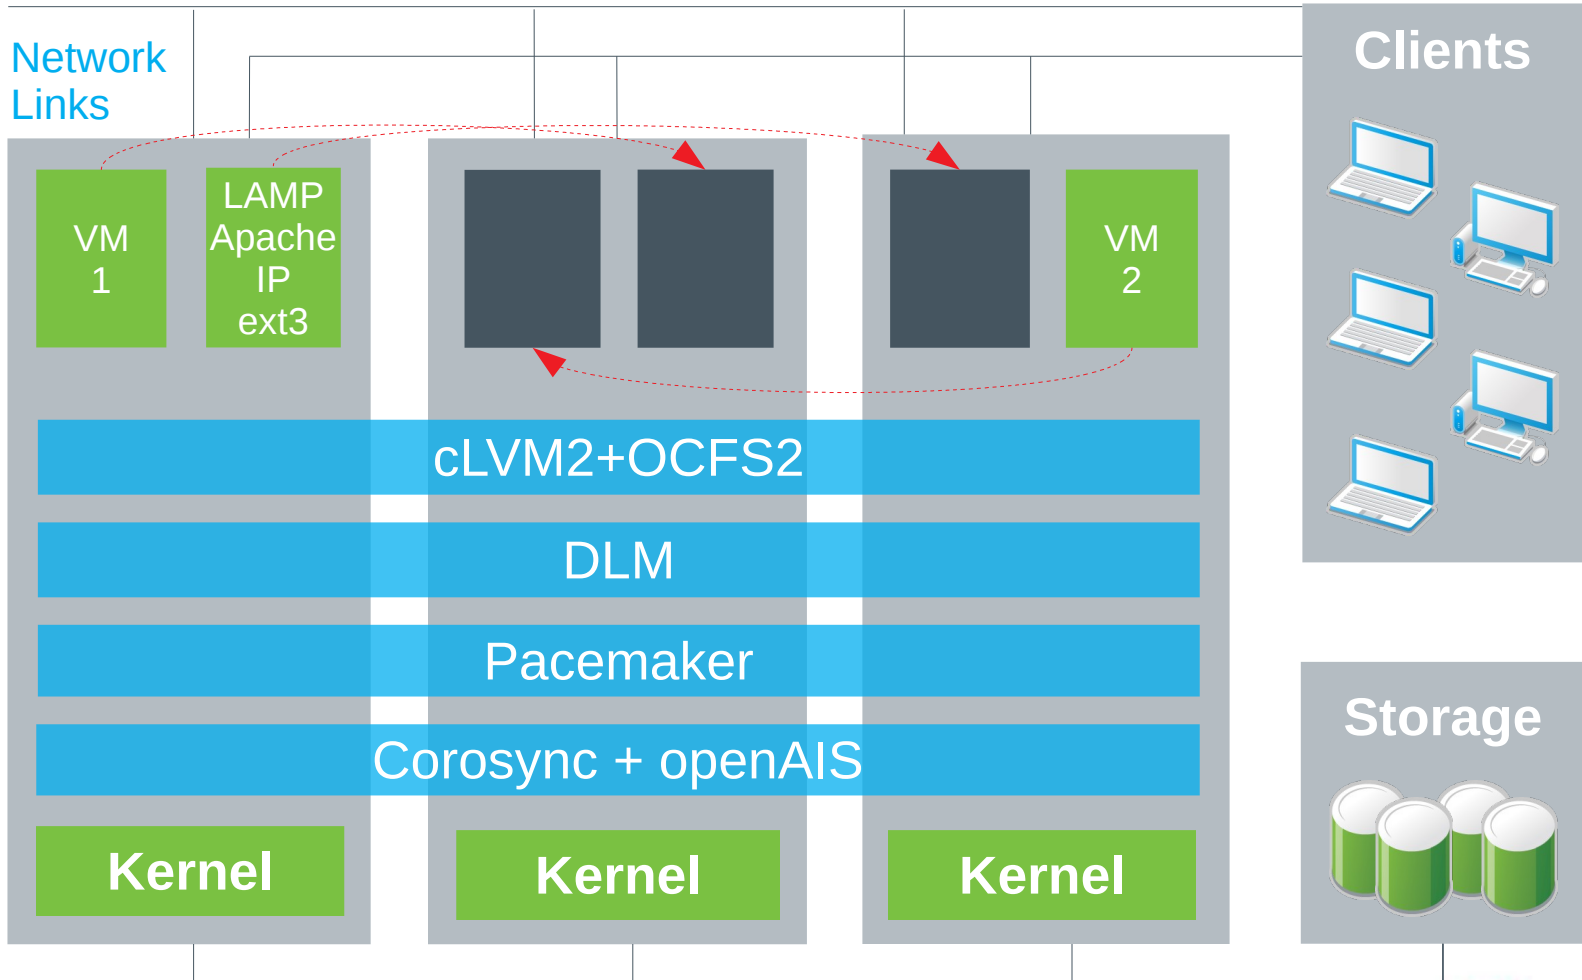
- Most modern and complete open source solution for implementing highly available Linux clusters
- A suite of robust open source technologies that is:
  - Affordable
  - Integrated
  - Virtualization aware
- Used with SUSE Linux Enterprise Server, it helps you:
  - Maintain business continuity
  - Protect data integrity
  - Reduce unplanned downtime for your mission-critical Linux workloads

# Linux High Availability Stack

## The stack includes:

- resource-agents – manage and monitor availability of services
- stonith – IO fencing support (snIPL - simple network IPL for System z)
- corosync and OpenAIS – cluster infrastructure
- Pacemaker – cluster resource manager
- CRM GUI – graphical interface for cluster resource and dependencies editing
- hawk – Web console for cluster monitoring and administration
- CLI – improved command line to interact with the CIB: editing, prepare multiple changes - commit once, syntax validation, etc.

# Cluster Diagram



# Key Use Cases

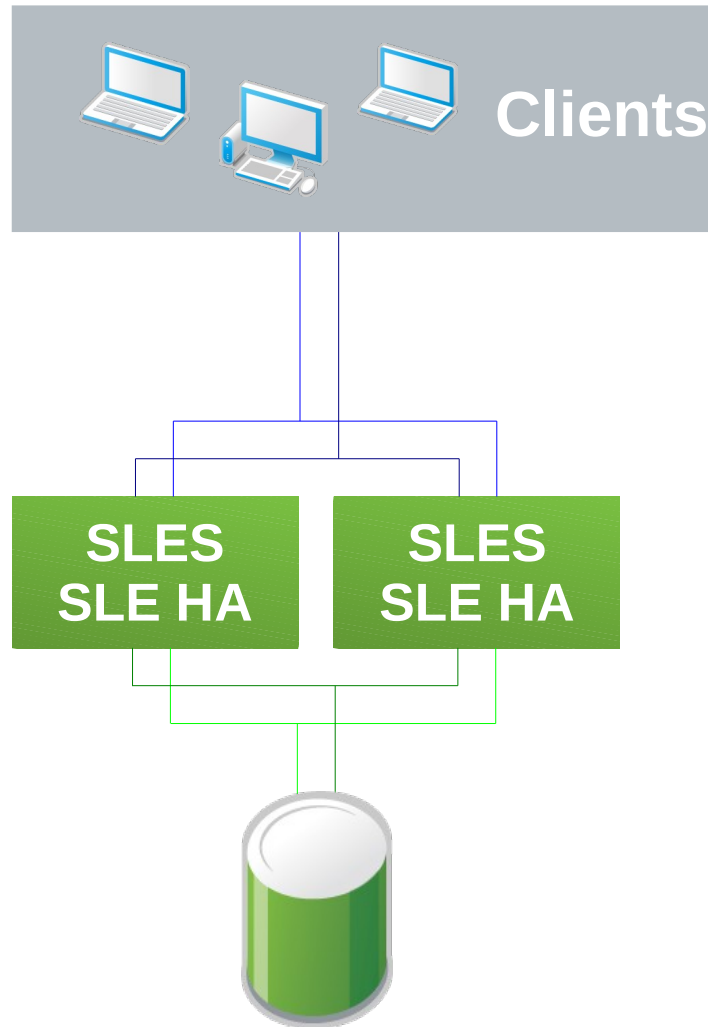
- Achieve high availability of mission-critical services
- Active/active services
  - OCFS2, Databases, Samba File Servers
- Active/passive service fail-over
  - Traditional databases, SAP setups, most regular services
- Private Cloud
  - HA, automation and orchestration for managed VMs
- High availability across guests
  - Build HA on top of a non-HA cloud
- Remote clustering
  - Local (GA), Metro (SP1), and Geographical (SP2) area clusters

# From Local to GEO

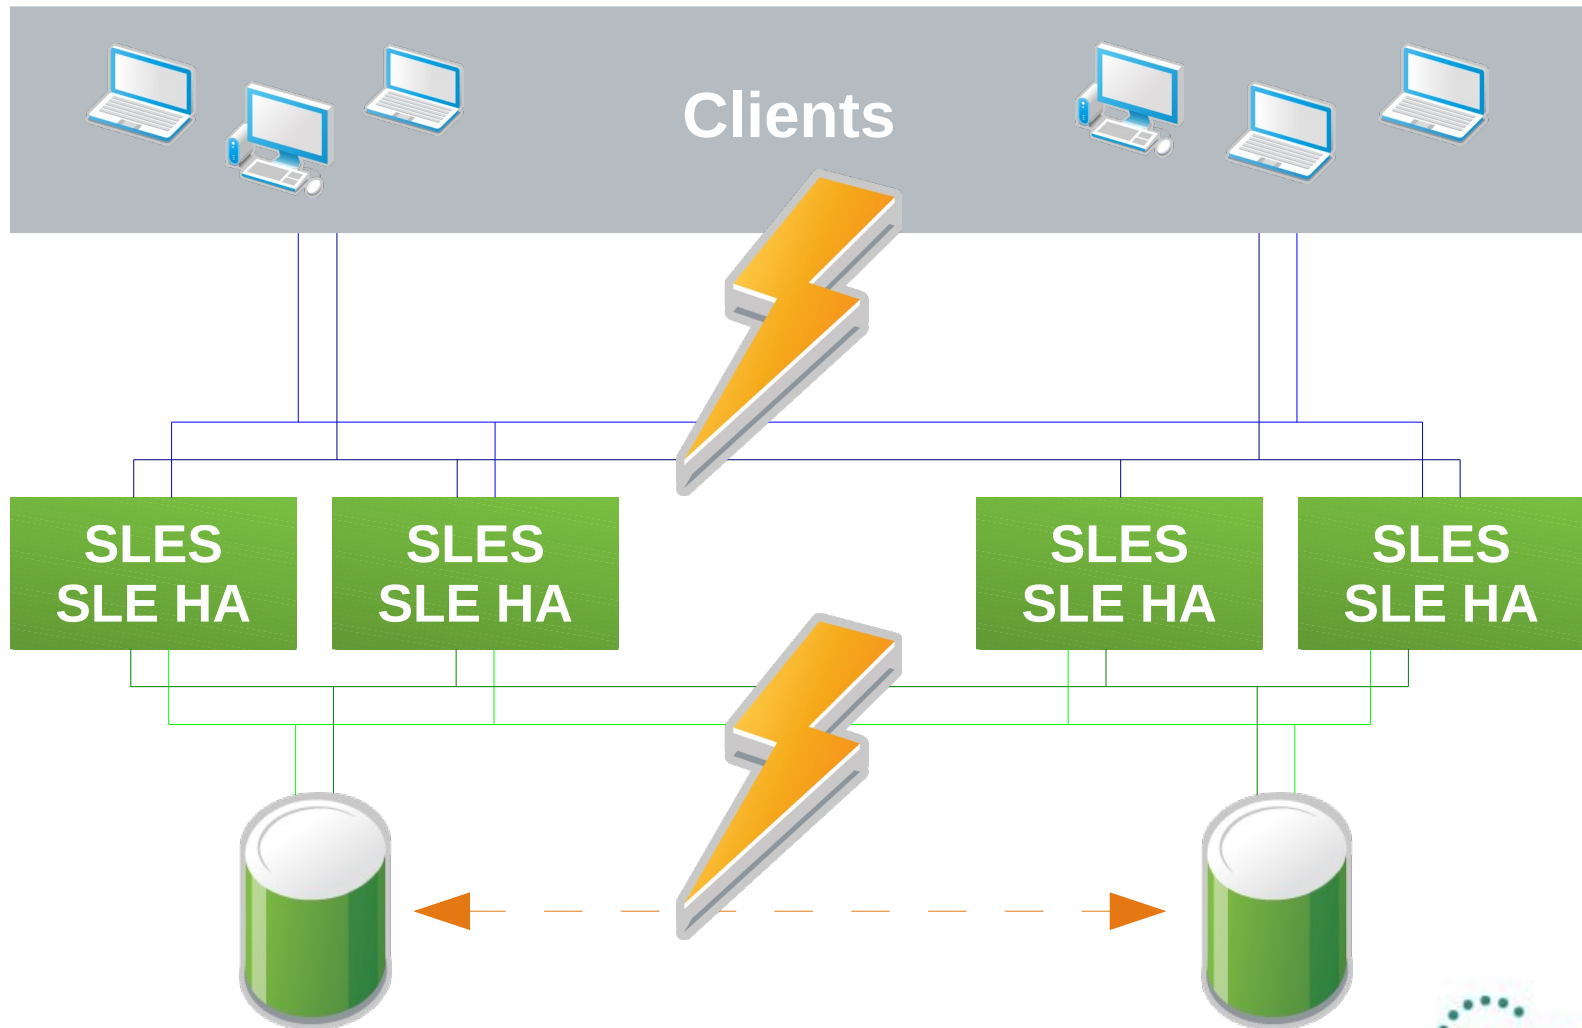
- Local cluster
  - Negligible network latency
  - Typically synchronous concurrent storage access
- Metro area (stretched) cluster
  - Network latency <15ms (~20 miles)
  - Unified / redundant network between sites
  - Usually some form of replication at the storage level
- Geo clustering
  - High network latency, limited bandwidth
  - Asynchronous storage replication



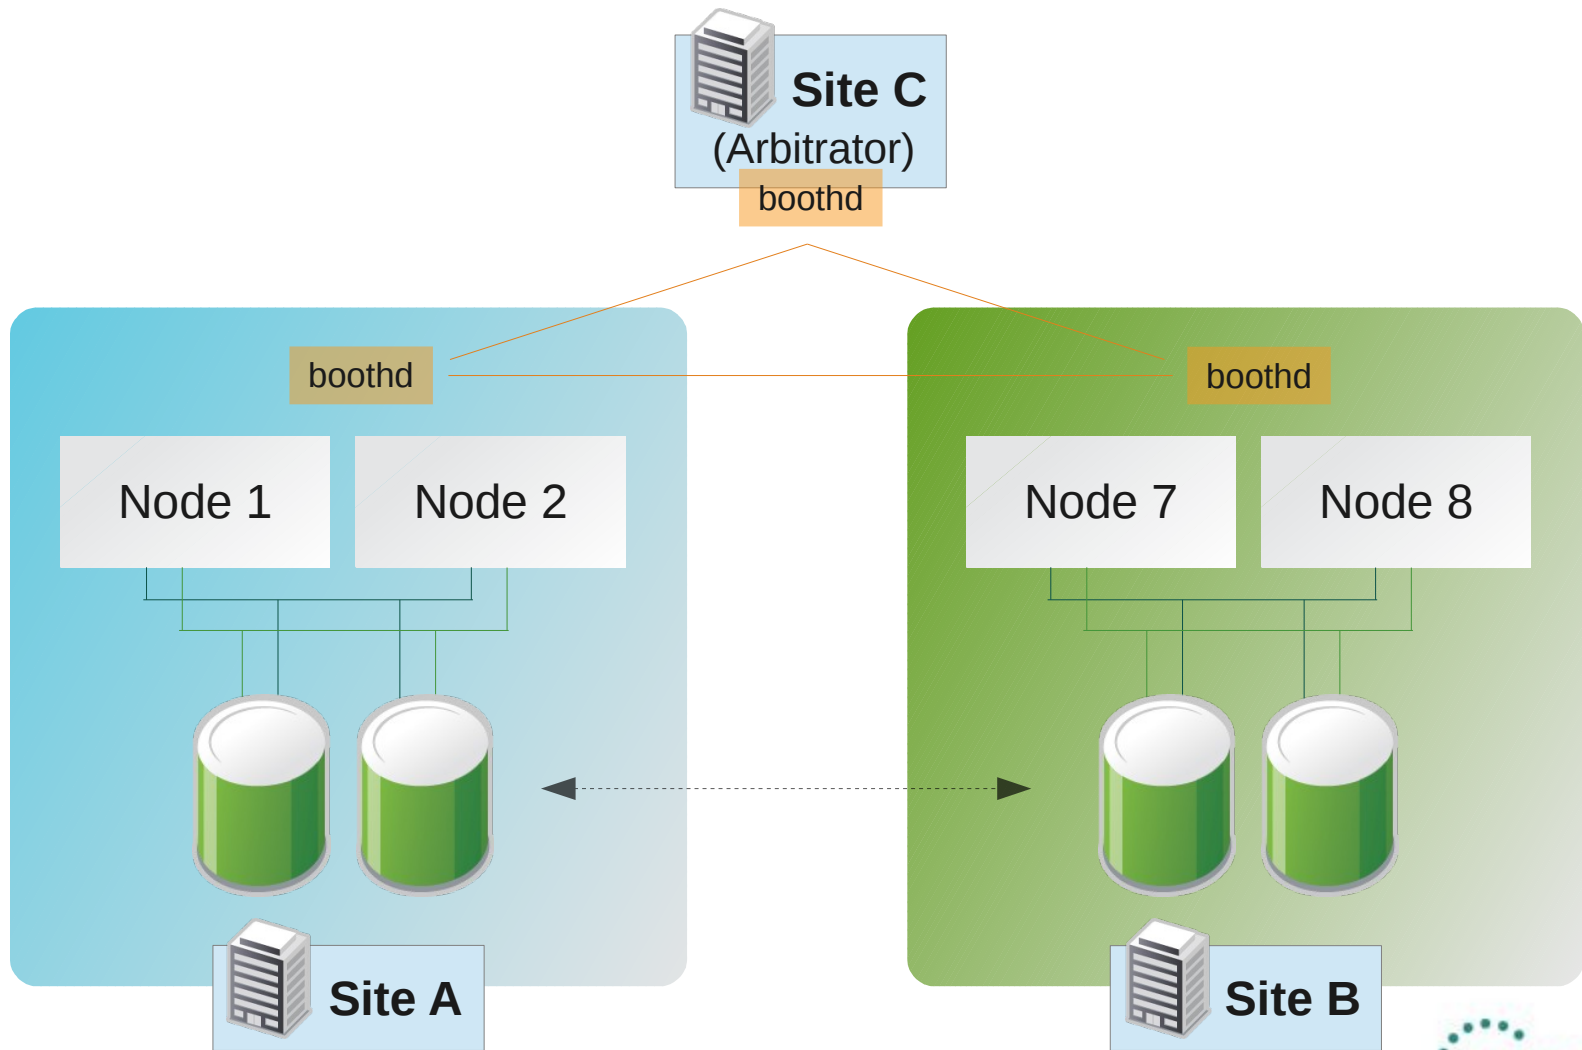
# Local Cluster



# Metro (Stretched) Cluster




# GEO Cluster



# Geo Clustering for SUSE Linux Enterprise High Availability Extension



- Cluster fail-over between different data center locations
  - Provide disaster resilience in case of site failure
  - Each site is a self-contained, autonomous cluster
  - Support both manual and automatic switch-/fail-over
- Extends Metro Cluster capabilities
  - No distance limit between data centers
  - No unified storage / network needed
- Storage replicated as active / passive
  - Leverage Distributed Replicated Block Device (DRBD)
  - Can integrate third-party solutions



# Getting the SLE High Availability Extension

- SUSE® Linux Enterprise High Availability Extension
  - Sold as annual support subscriptions – 1 or 3-year
  - Inherits the support level of the underlying SUSE Linux Enterprise Server subscription
  - Only charged for x86 and x86\_64 – included for free with Itanium, IBM POWER and IBM System z subscriptions
- Geo Clustering for SUSE Linux Enterprise High Availability Extension
  - Sold as annual support subscriptions – 1 or 3-year
  - Additional option for the SUSE Linux Enterprise High Availability Extension
    - *Each system participating in the Geo Cluster needs a subscription for GEOClustering, the High Availability Extension and SUSE Linux Enterprise Server*
  - Inherits the support level of the underlying SUSE Linux Enterprise Server subscription

**Let's build a two node cluster...**

# Session Evaluation

## 12367 - High Availability for Highly Reliable Systems





**Corporate Headquarters**  
Maxfeldstrasse 5  
90409 Nuremberg  
Germany

+49 911 740 53 0 (Worldwide)  
[www.suse.com](http://www.suse.com)

Join us on:  
[www.opensuse.org](http://www.opensuse.org)