



SUSE® Linux Enterprise High Availability

Reliable Availability for a Reliable Platform

Kai Dupke

Senior Product Manager
SUSE Linux Enterprise Server
kdupke@suse.com

Mike Friesenegger

Systems Engineer
mfriesenegger@suse.com

Distribution: pdf any
Date: 2012-08-07
No public document



Topics

SUSE® Linux Enterprise High Availability

The Challenge

Overview

Roadmap & Feature Comparison

Service Pack 2

Demo

Geo Cluster

Architecture

Challenge

Challenge

SUSE® Linux Enterprise High Availability

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?

Overview

Benefits

SUSE® Linux Enterprise High Availability

-  Quickly and easily install, configure and manage clustered Linux servers
-  Ensure continuous access to your mission-critical systems and data
-  Transparent to Virtualization – nodes can be virtual or physical
-  Meet your Service Level Agreements
-  Increase service availability

Overview

SUSE® Linux Enterprise High Availability

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



Key Use Cases

SUSE® Linux Enterprise High Availability

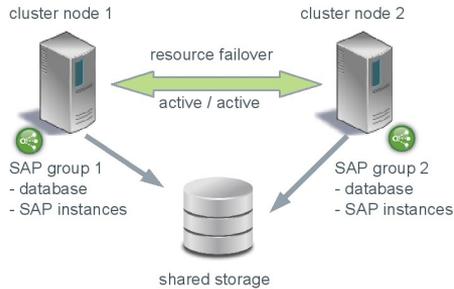
- High availability for mission-critical services
- Active/active services
 - OCFS2, Databases, Samba File Servers
- Active/passive service fail-over
 - Traditional databases, SAP setups, regular services
- Private Cloud
 - HA, automation and orchestration for managed VMs
- High availability across guests
 - Fine granular monitoring and HA on top of virtualization
- Remote clustering
 - Local, Metro, and Geographical area clusters



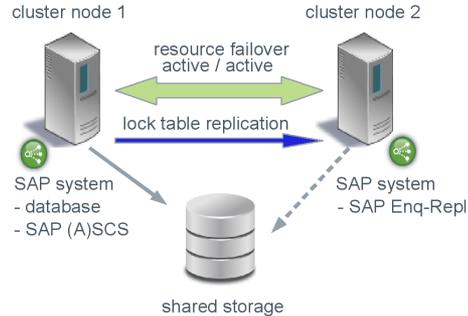
Key Use Cases SAP

SUSE® Linux Enterprise High Availability

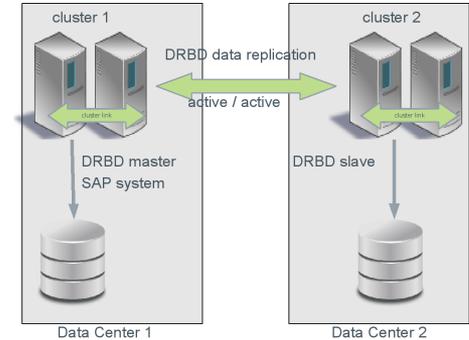
Simple Stack HA



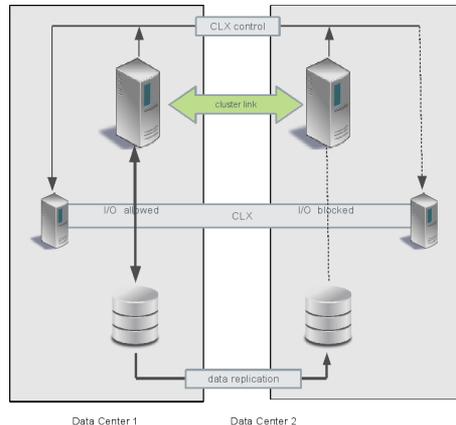
Enqueue Replication



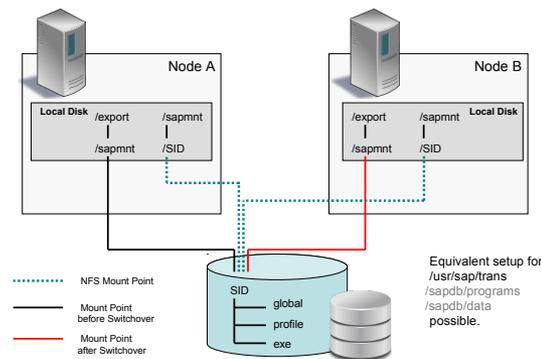
DRBD Data Replication



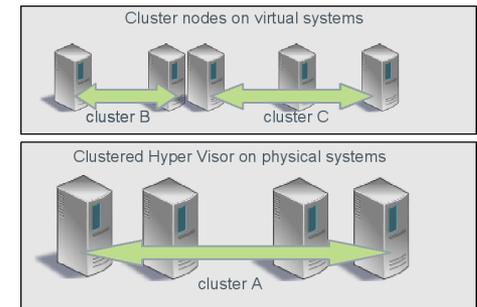
HP CLX SAN Based Mirror



NFS and SAP in one Cluster



HA in Virtualized Environments



Roadmap

Roadmap

SUSE® Linux Enterprise High Availability



SLE HA 11 SP1

- Metro Area Cluster
- Samba Cluster
- Web GUI
- Cluster Test Drive
- Node Recovery

SLE HA 11 SP2

- Geo Cluster
- History Explorer
- Setup Tools
- Improved Web Frontend

SLE HA 11 SP3

- Remote Monitoring

SLE HA 12

- Web Console
- Cloud integration
- Preloaded Clusters

Comparison

Competition

SUSE® Linux Enterprise High Availability

| Competitive Point | SUSE Linux Enterprise High Availability Extension | Red Hat | Symantec VCS |
|--------------------------------|--|--|--|
| Requires shared storage | No | Yes | No |
| Open Source based | Yes | Yes | No |
| Geo Extension | Yes | No | Yes |
| Supports virtualization | Hybrid physical, virtual clusters, protects guests and guest apps; supports KVM, Xen, VMware | KVM, apps within guest, clusters physical, virtual servers | VMware ESX server, protects apps in guests |
| OS integrated tools | Yes | Yes | No |
| Free tools and resource agents | Yes | No (extra for Load Balancer, Clustered Samba, and SAP Resource Agent) | No (extra charged) |
| Platform Support | x86, x86_64, Itanium, IBM POWER, IBM System z | Only on x86 and x86_64 | x86, x86_64 |
| Major Version Upgrade | Yes | No | No |
| Rolling Update | Yes | No | No |
| Node Recovery included | Yes | No | No |
| Cost | \$\$ | \$\$\$ | \$\$\$\$ |

Service Pack 2

Service Pack 2 – New Features

SUSE® Linux Enterprise High Availability

- **Easy Installation & Set-Up**
 - Templates and Wizards
 - Cluster Bootstrap & Join
- **Improved Supportability**
 - History Explorer
 - Log File Query Tools
- **Efficient Management**
 - Access Control Lists
 - Enhanced Web Console
- **Additional Capabilities**
 - Joining of Clustered SAMBA to Active Directory
 - Load Balancer Connection Tracking and Replication
 - Multiple SBD devices for reliable storage-based fencing
 - ReaR support for SUSE boot media



Service Pack 2 – easy setup

SUSE® Linux Enterprise High Availability

- Bootstrapping a cluster is really easy:
 - node1 # sleha-init -i bond0 -t ocfs2 -p /dev/sdb
 - nodeN # sleha-join -c 192.168.2.1
- Configuring a cluster file system or web server
 - Connect to the hawk web console
 - Start the wizard for OCFS2 or web server

Service Pack 2 – Wizards

SUSE® Linux Enterprise High Availability

Cluster Setup Wizard User: hacluster [Log Out](#)

Choose Configuration

This wizard will guide you through the steps necessary to add one of the following configurations to the cluster.

- Web Server
- OCFS2 Filesystem
- OCFS2 (Additional)
- SAP SimpleStack Instance**
- SAP SimpleStack+ Instance
- SAP Database Instance
- SAP ASCS Instance
- SAP Central Instance

Configure a SAP instance including: 1) virtual IP addresses for each of the SAP instance services - ASCS, DB and CI, 2) a RAID 1 host based mirror, 3) a cluster manager LVM volume group and LVM volumes on the RAID 1 host based mirror, 4) filesystems on shared storage for sapmnt, /sapbd and /usr/sap, 5) SAPinstance for - ASCS, a Database ,a Central Instance.

Cluster Setup Wizard User: hacluster [Log Out](#)

SAP SimpleStack Instance: SAP Database Instance

This is the SAP Database instance.

Resource ID:

SID:

DBTYPE:

Constraints User: hacluster [Log Out](#)

- Location
- Colocation
colo-vm INFINITY →
- Order
order-vm INFINITY →
- Ticket
dep-ticket-a site-a
dep-ticket-b site-b

Service Pack 2 – Cluster Simulator

SUSE® Linux Enterprise High Availability

The screenshot displays the 'Cluster Status' interface. At the top left, there is a 'Cluster Status' header with a user icon and the text 'User: hacluster [Log Out](#)'. The main area is a grid of resource status boxes. The first row shows 'hex-0: Online', 'hex-7: Online', and 'hex-9: Online'. Below this is a section for 'Clone Set: base-clone' containing various resources like 'dlm:0: Started', 'o2cb:0: Started', 'clvm:0: Started', 'cmirrord:0: Started', 'vg1:0: Started', and 'ocfs2-1:0: Started'. The second row shows 'fencing-sbd: Started' and 'vm-07: Started'. The third row shows 'vm-08: Started', 'vm-10: Started', 'vm-12: Started', 'vm-14: Started', 'vm-16: Started', 'vm-18: Started', 'vm-20: Started', and 'vm-22: Started'. The fourth row shows 'vm-09: Started', 'vm-11: Started', 'vm-13: Started', 'vm-15: Started', 'vm-17: Started', 'vm-19: Started', 'vm-21: Started', and 'vm-30: Started'. The fifth row shows 'vm-05: Started', 'vm-06: Started', 'vm-31: Started', and 'vm-32: Started'. A dialog box titled 'Inject Operation' is open in the center, with fields for 'Resource: vm-08', 'Operation: start', 'Interval: (ms)', 'Node: hex-0', and 'Result: Not Configured'. It has 'OK' and 'Cancel' buttons. In the bottom right, there is a 'Simulator (initial state)' window with an 'Injected State' text area, a 'Run >' button, and buttons for '+ Node', '+ Op', and '-'. It also has 'Reset' and 'Close' buttons. The bottom left corner of the image has the text 'Copyright © 2009-2012 Novell, Inc.'

Service Pack 2 – Multiple SBD fencing

SUSE® Linux Enterprise High Availability

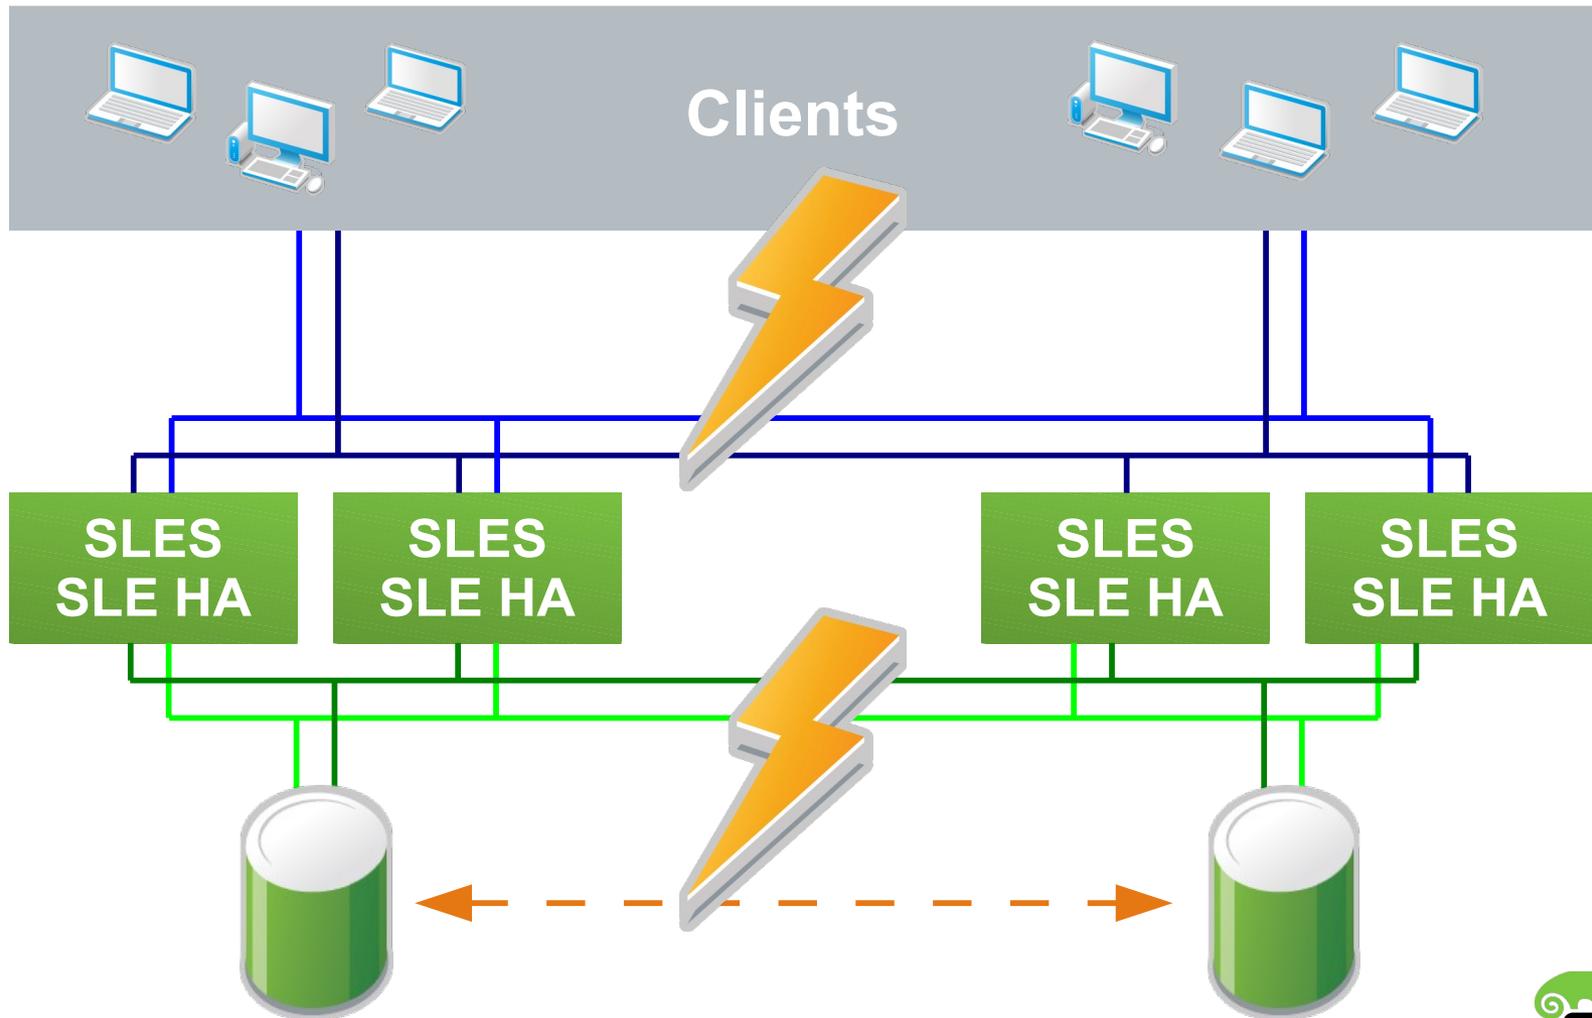
- Shared Block Device (SBD) fencing is recommended by SUSE
 - SBD fencing is highly reliable
 - Independent of management board (firmware, settings, etc.)
 - Equal setup in physical and virtual environments, reducing variance in deployments
- Multiple SBD fencing
 - Supports redundancy in fencing channels
 - Enhanced reliability leveraging independent storage systems
- Integrated with the hardware watchdog devices

Demo

From Local Cluster to Geo Cluster

Local & Stretched Cluster

SUSE® Linux Enterprise High Availability



Geo Cluster – From Local to Geo

SUSE® Linux Enterprise High Availability

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

Geo Cluster – Overview

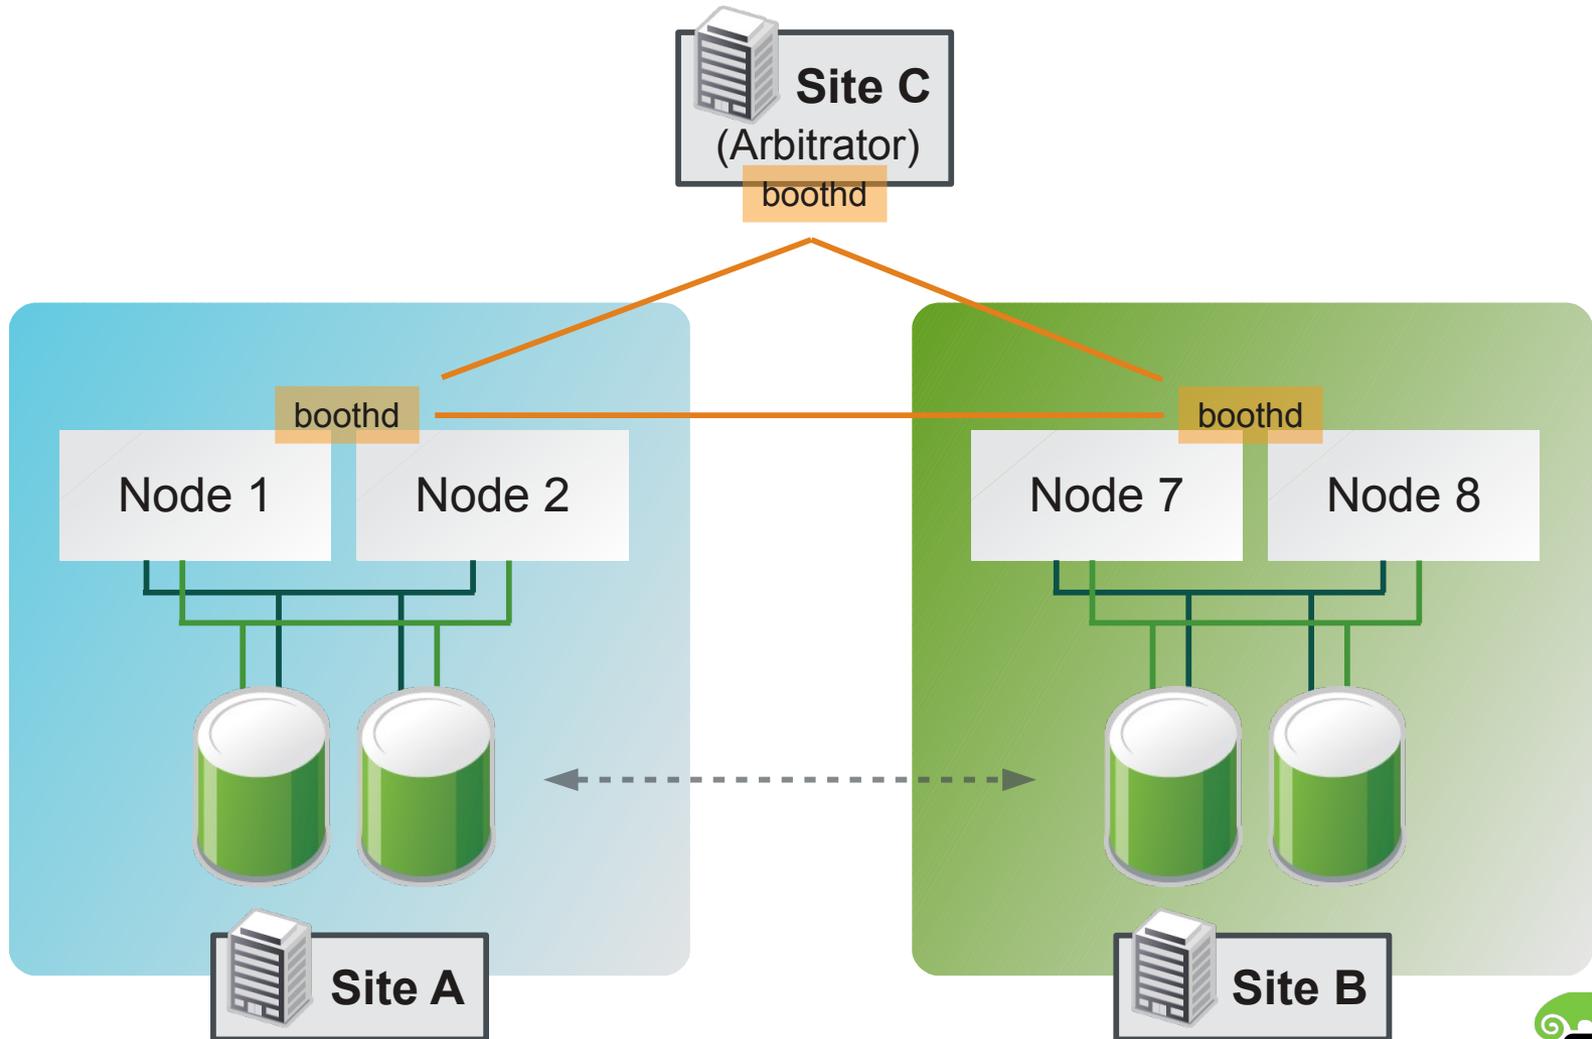
SUSE® Linux Enterprise High Availability

- 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 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 via scripts



Geo Cluster – Setup

SUSE® Linux Enterprise High Availability



Geo Cluster – Delivery

SUSE® Linux Enterprise High Availability

- Additional option for the SUSE Linux Enterprise High Availability Extension
 - Each system participating in the Geo Cluster needs a subscription for the GEO Clustering, the High Availability Extension, and the SUSE Linux Enterprise Server
 - The High Availability Extension is part a subscription for IBM System Z
- Support inherited from base server subscription



Visit us at

booth 324



Thank you.

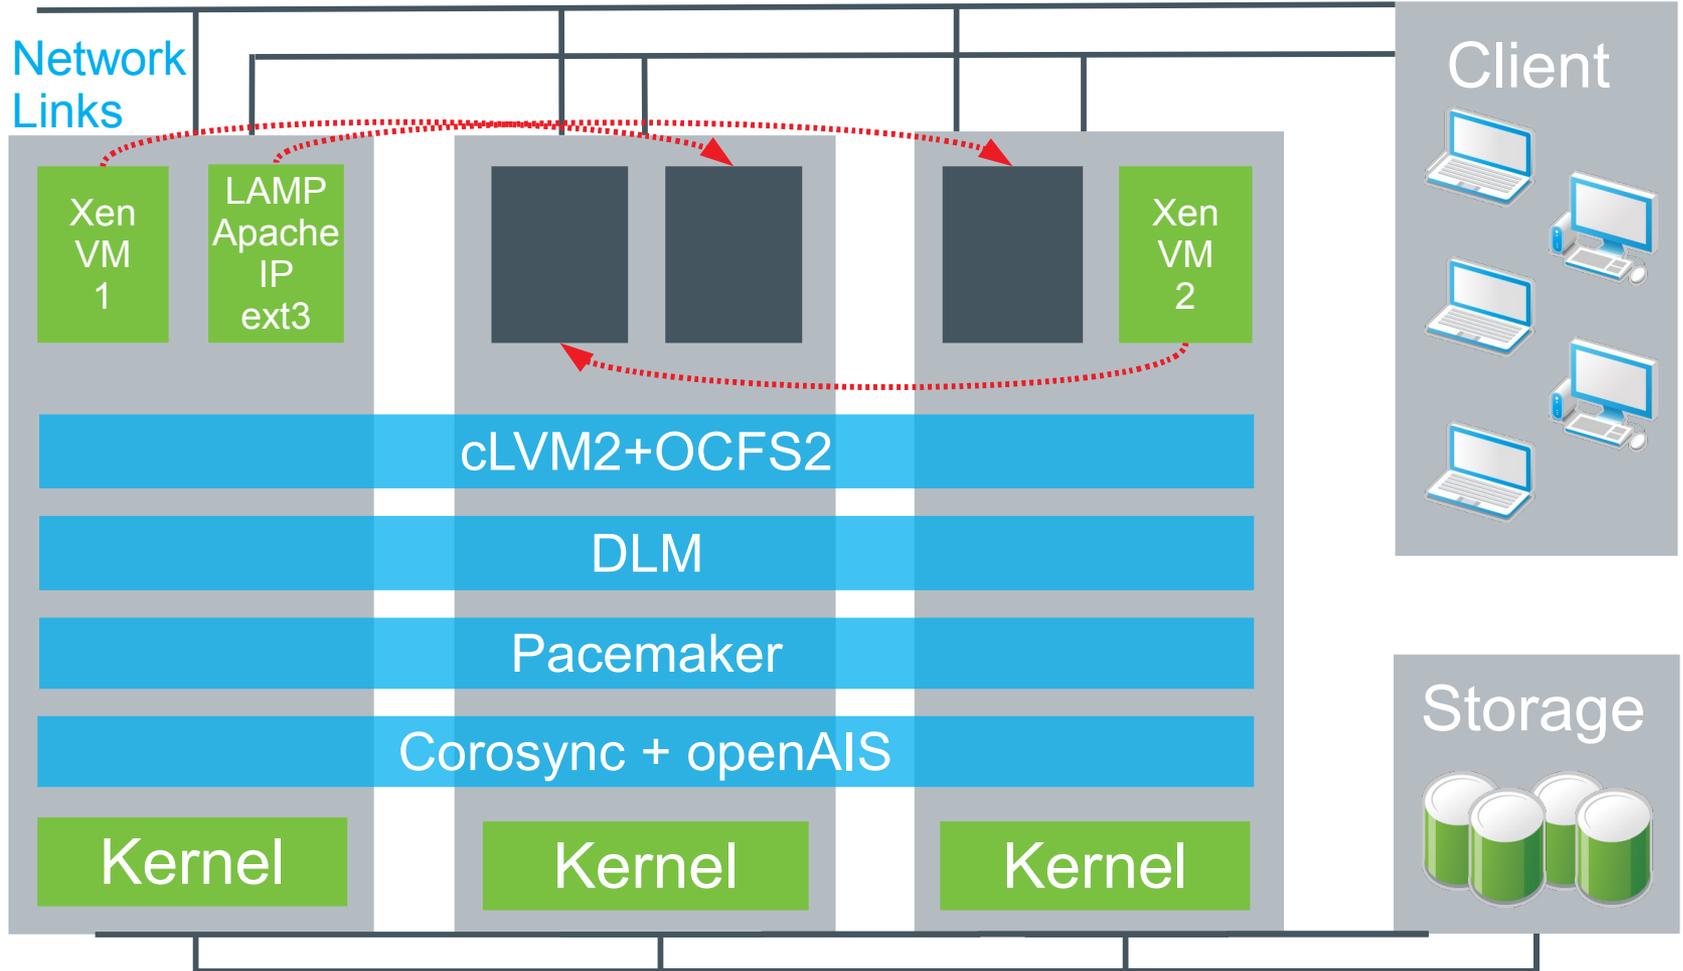


Appendix

Architecture

Cluster Example

SUSE® Linux Enterprise High Availability



Linux High Availability Stack

SUSE® Linux Enterprise High Availability

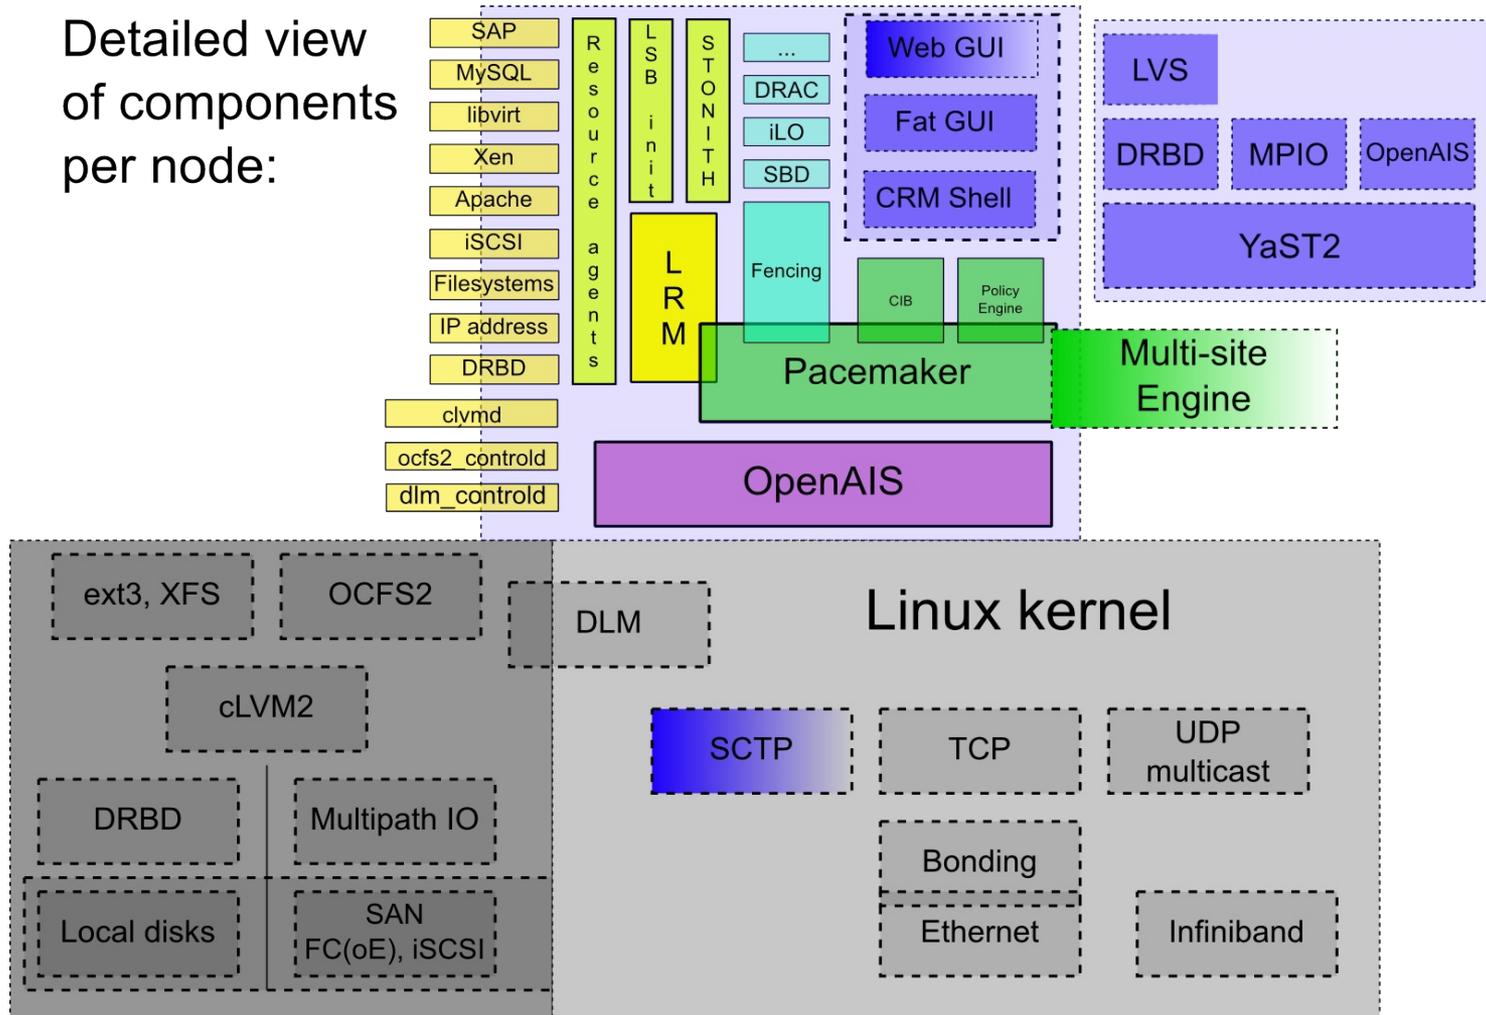
- The stack includes:
 - resource-agents – manage and monitor availability of services
 - stonith – IO fencing support (also Xen and VMware VMs)
 - corosync and OpenAIS – cluster infrastructure
 - Pacemaker – cluster resource manager
 - CRM GUI – graphical interface for cluster resource and dependencies 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.



Detailed Architecture

SUSE® Linux Enterprise High Availability

Detailed view of components per node:



Learn more

www.suse.com/products/highavailability

Thank you.





Unpublished Work of SUSE. All Rights Reserved.

This work is an unpublished work and contains confidential, proprietary and trade secret information of SUSE. Access to this work is restricted to SUSE employees who have a need to know to perform tasks within the scope of their assignments. No part of this work may be practiced, performed, copied, distributed, revised, modified, translated, abridged, condensed, expanded, collected, or adapted without the prior written consent of SUSE. Any use or exploitation of this work without authorization could subject the perpetrator to criminal and civil liability.

General Disclaimer

This document is not to be construed as a promise by any participating company to develop, deliver, or market a product. It is not a commitment to deliver any material, code, or functionality, and should not be relied upon in making purchasing decisions. SUSE makes no representations or warranties with respect to the contents of this document, and specifically disclaims any express or implied warranties of merchantability or fitness for any particular purpose. The development, release, and timing of features or functionality described for SUSE products remains at the sole discretion of SUSE. Further, SUSE reserves the right to revise this document and to make changes to its content, at any time, without obligation to notify any person or entity of such revisions or changes. All SUSE marks referenced in this presentation are trademarks or registered trademarks of Novell, Inc. in the United States and other countries. All third-party trademarks are the property of their respective owners.

