

# SAP Certified Cluster On SUSE® Linux Enterprise Server For SAP Applications



Architecture



Best Practices



Components

**Fabian Herschel, SUSE**

SAP Architect LinuxLab

Fabian.herschel@suse.com

**Markus Gürtler, SUSE**

Architect & Technical ISV Manager

mguertler@suse.com



# Agenda



Overview: SAP and High Availability



SAP “Enqueue Replication” Scenario



SAP Certified Cluster Architecture



SAP in Virtualized Environments



Outlook: SAP HANA in the cluster



# Overview: SAP and High Availability

# What is SAP NetWeaver?

Application server platform for SAP Business Suite

Three application server flavors available:

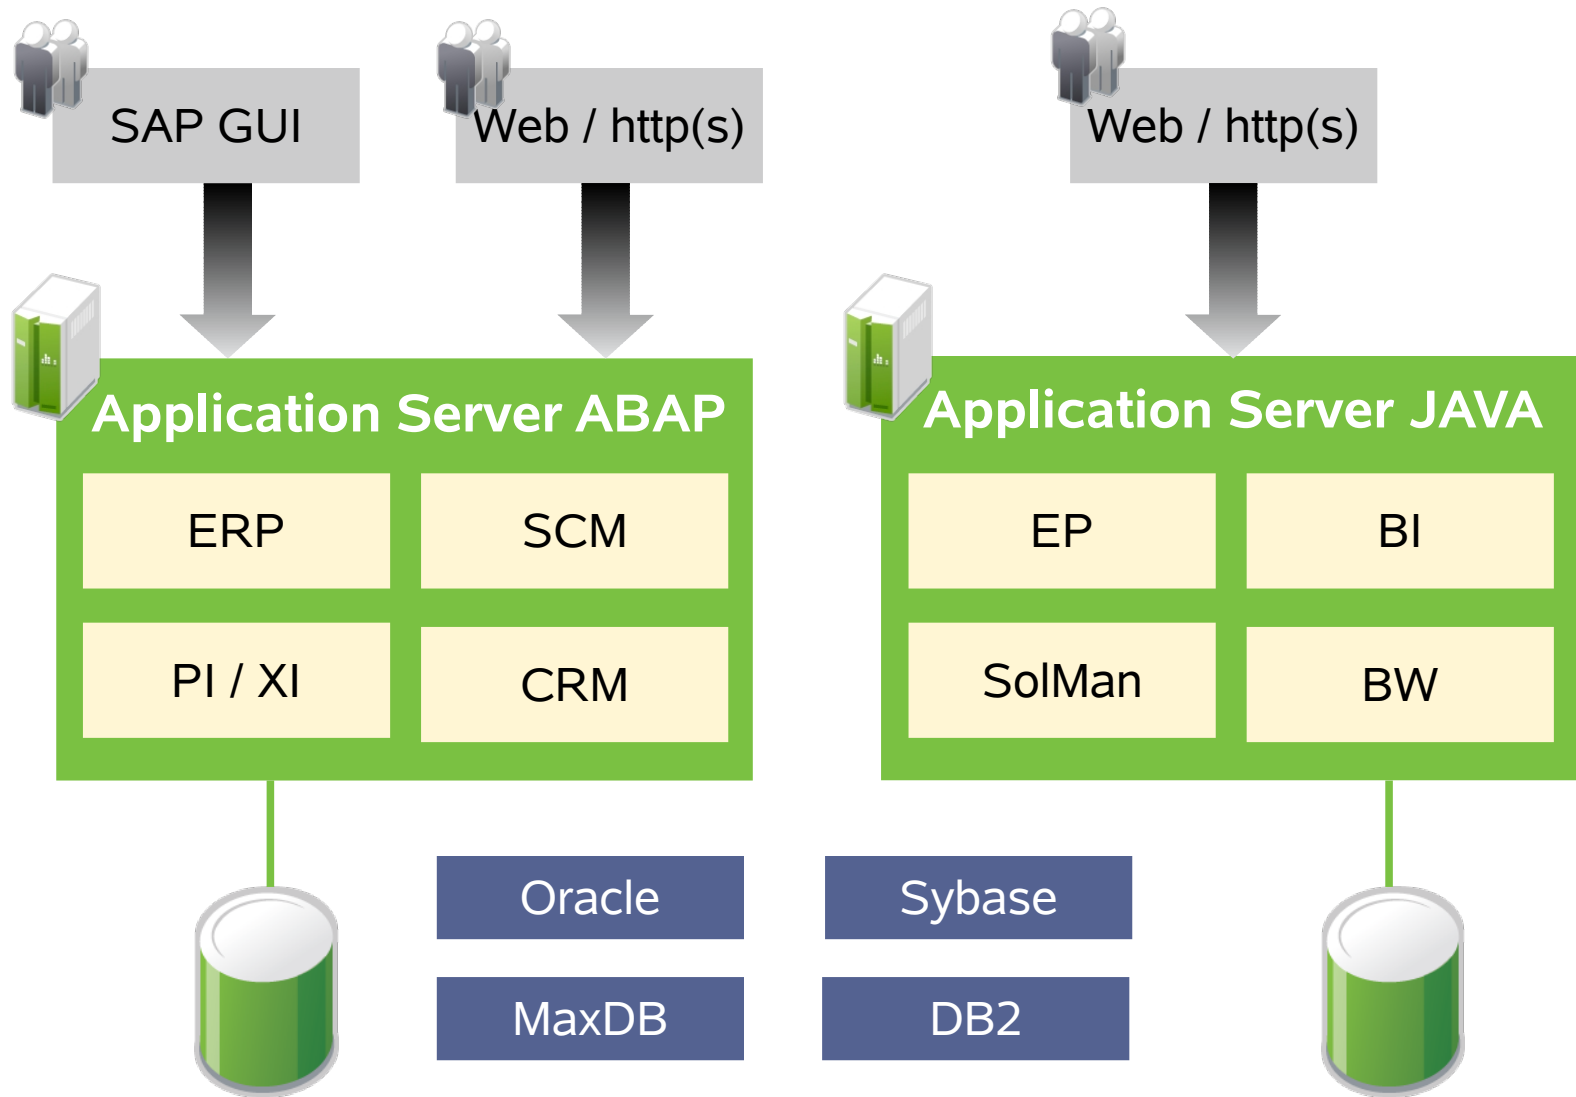
**ABAP:** SAP's business programming language

**JAVA:** JAVA based SAP applications

**Dual-stack:** Combines ABAP and JAVA stack

Users use clients (web browser or SAP GUI) to connect to the SAP application servers

# SAP Application Server Overview



# Why High Availability for SAP?



## Minimized downtimes

- in case of hardware and software failures
- in case of disasters
- for administrative tasks like patching of SAP and OS



## High requirements on SLAs



SAP relies on 3<sup>rd</sup> party vendors for High Availability

Availability	Outage times	
90%	5	weeks
98%	1	weeks
99%	3.7	days
99.8%	18	hours
99.9%	9	hours
99.99%	53	minutes
99.999%	5	minutes
99.99% and 99.999% requires an HA cluster		



# What Does SUSE Offer?



## **SAP certified solution**

Industry proven and SAP certified High Availability solution



## **“Best Practice” guides**

“Best Practice” guides describe detailed the most relevant SAP High Availability scenarios



## **Partner Ecosystem for HA integration**

Strong partner ecosystem for HA integration projects



## **Continually improved**

Our SAP HA solutions are continually improved and extended by new use-cases in a strong collaboration with SAP

# High Availability Training

SAP offers a training course for SUSE HA

- Duration: 3 days
- Content: SLES for SAP + High Availability
- More details at <http://training.sap.com>





# High Availability Software Components Shipped With SLES For SAP

## **SUSE Linux Enterprise High Availability Extension**

The most modern and complete open source solution for implementing highly available Linux clusters

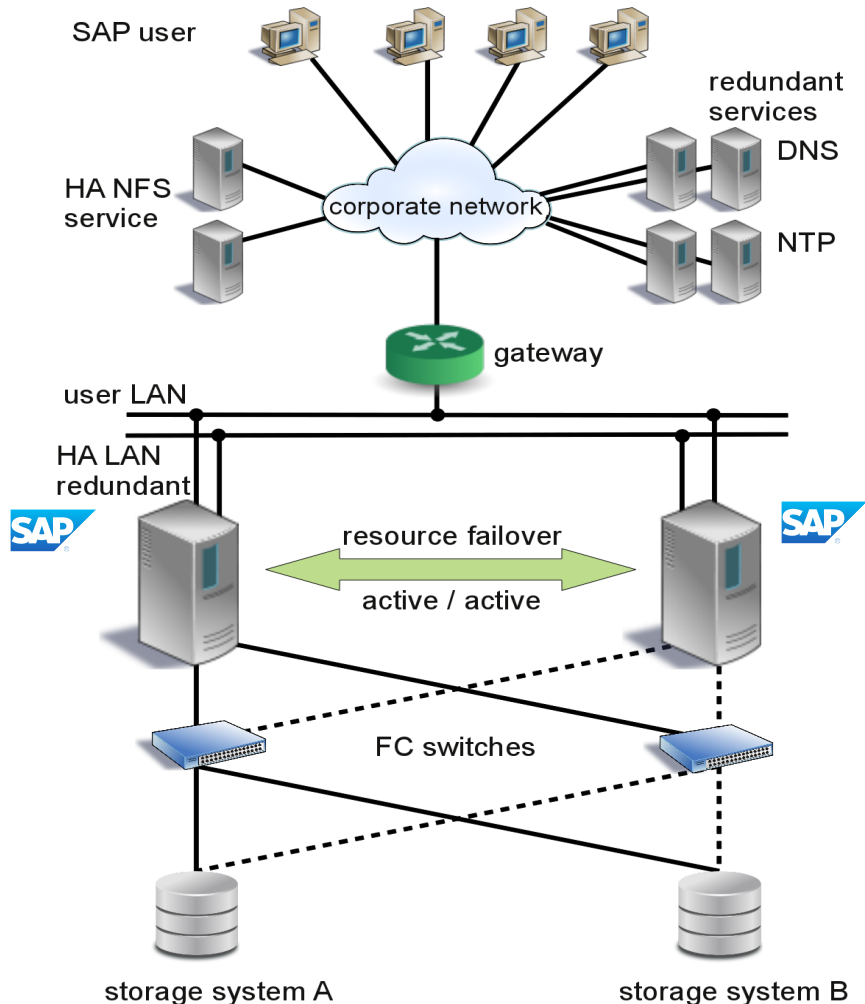
## **SAPInstance resource agent**

Supports SAP NetWeaver stack with kernel  $\geq 6.40$

## **SAPDatabase resource agent**


Supports Sybase, Oracle, MaxDB, DB2

# SAP HA Basis Architecture



- ✓ No single point of failure
- ✓ Split-brain avoidance
- ✓ full redundant and high available
- ✓ Active / Active
- ✓ NFS for shared SAP directories
- ✓ Fast SAN for SAP database


# Cluster Failover



Left node failed

Right node takes over

# Split Brain

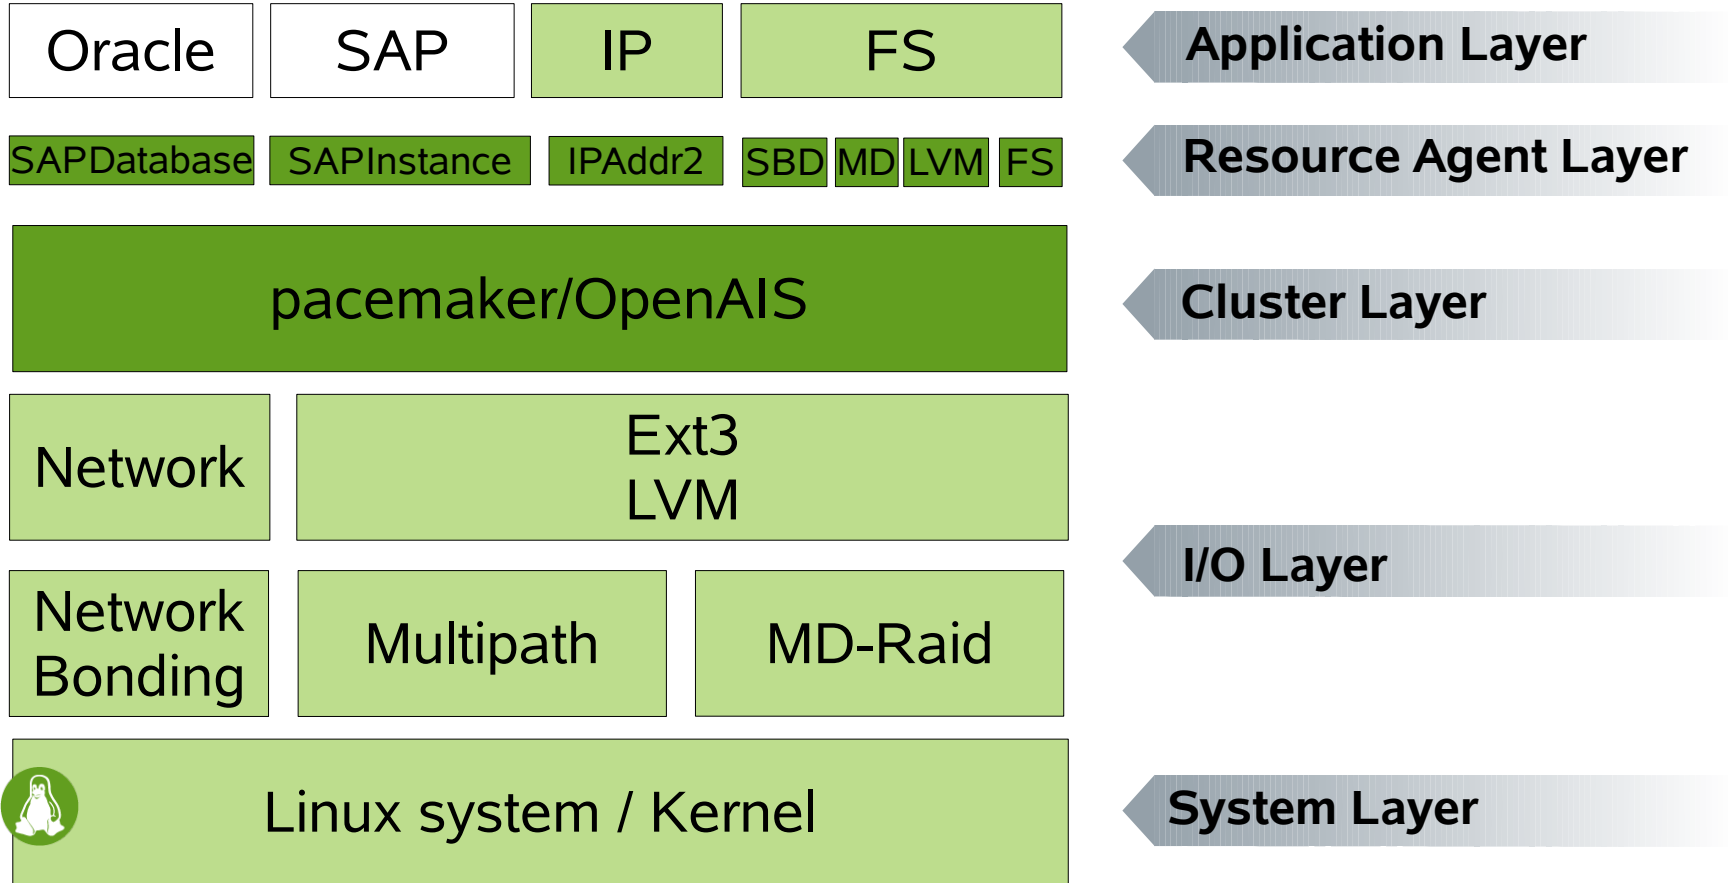


Split brain detection

using fencing

SBD devices -  
Storage-based death

# HA Stack for SAP



# High Availability “Best Practice” Guides


Jointly developed by SUSE, customers and partners

Addressing “real world” use-cases




# High Availability – Use Cases


## Simple Stack




## Enqueue Replication



## DRBD Data Sync



## HA in Virtual Environments



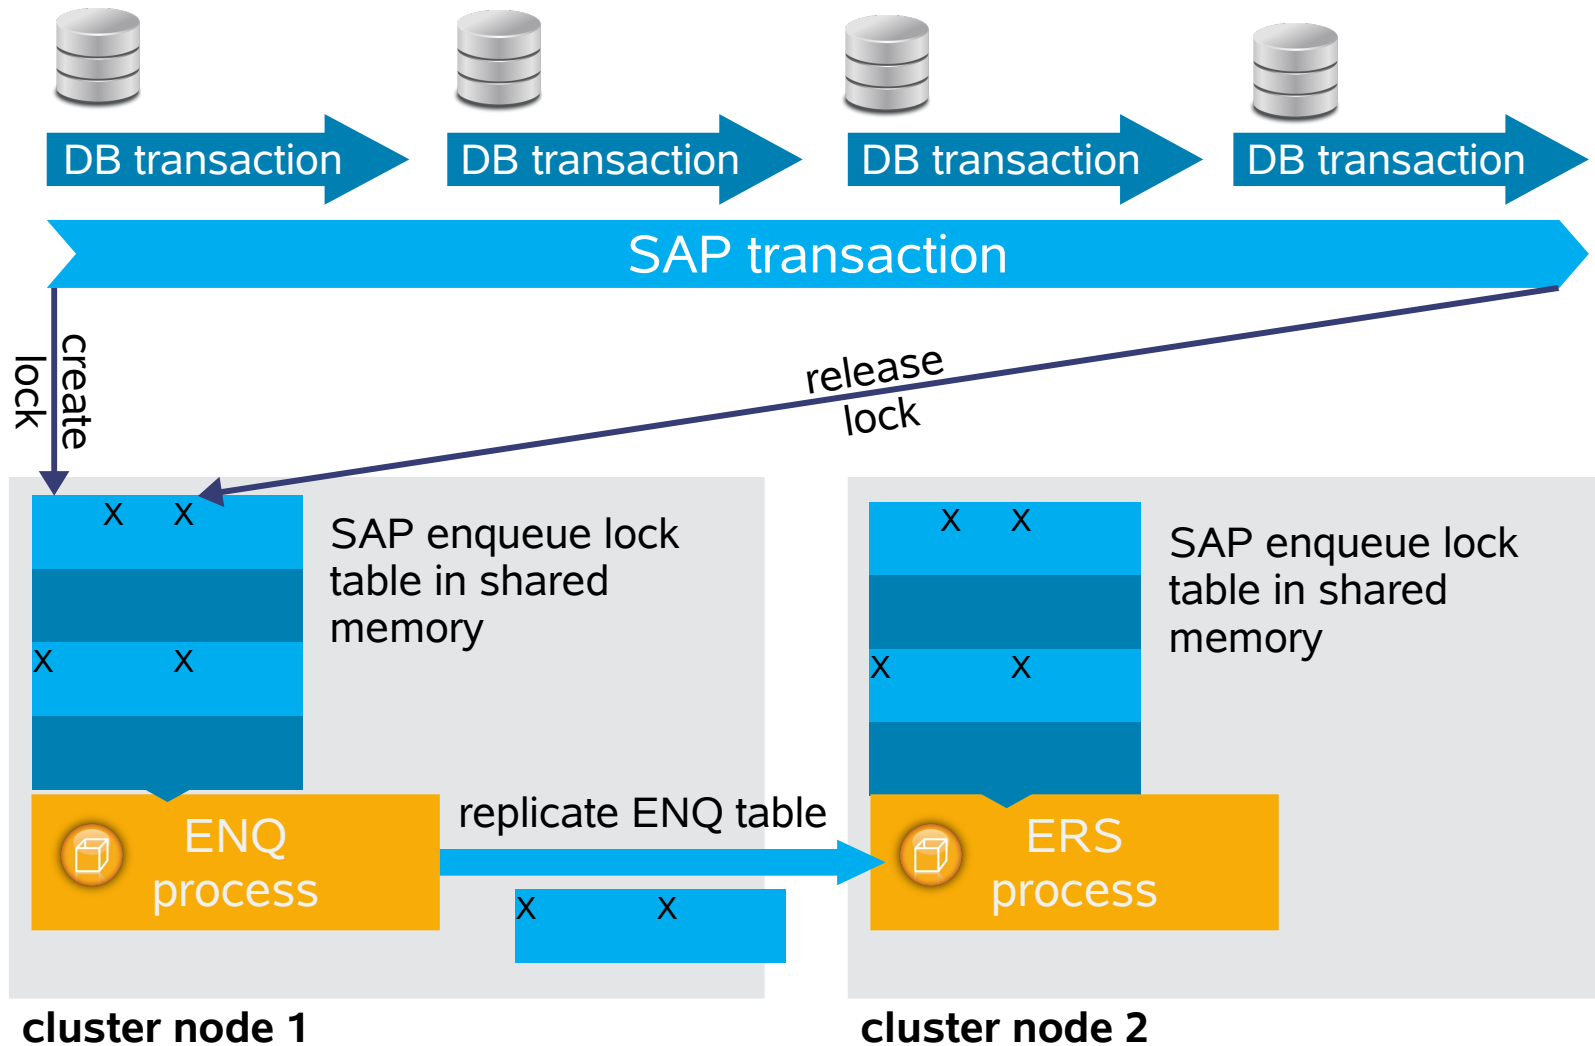





# SAP Enqueue Replication Scenario

# SAP Enqueue-Replication


How does it work?



# Enqueue Replication – Use Case



# Enqueue Replication - Take Over




The cluster registers that Node1 is failed

With help of a STONITH action we ensure that Node1 is really off

Node2 takes all resources, and the Enq.-Repl. Instance get promoted to the (A)SCS instance

The locking table within the shared memory gets transferred and released

# Enqueue Replication - Node Re-Join



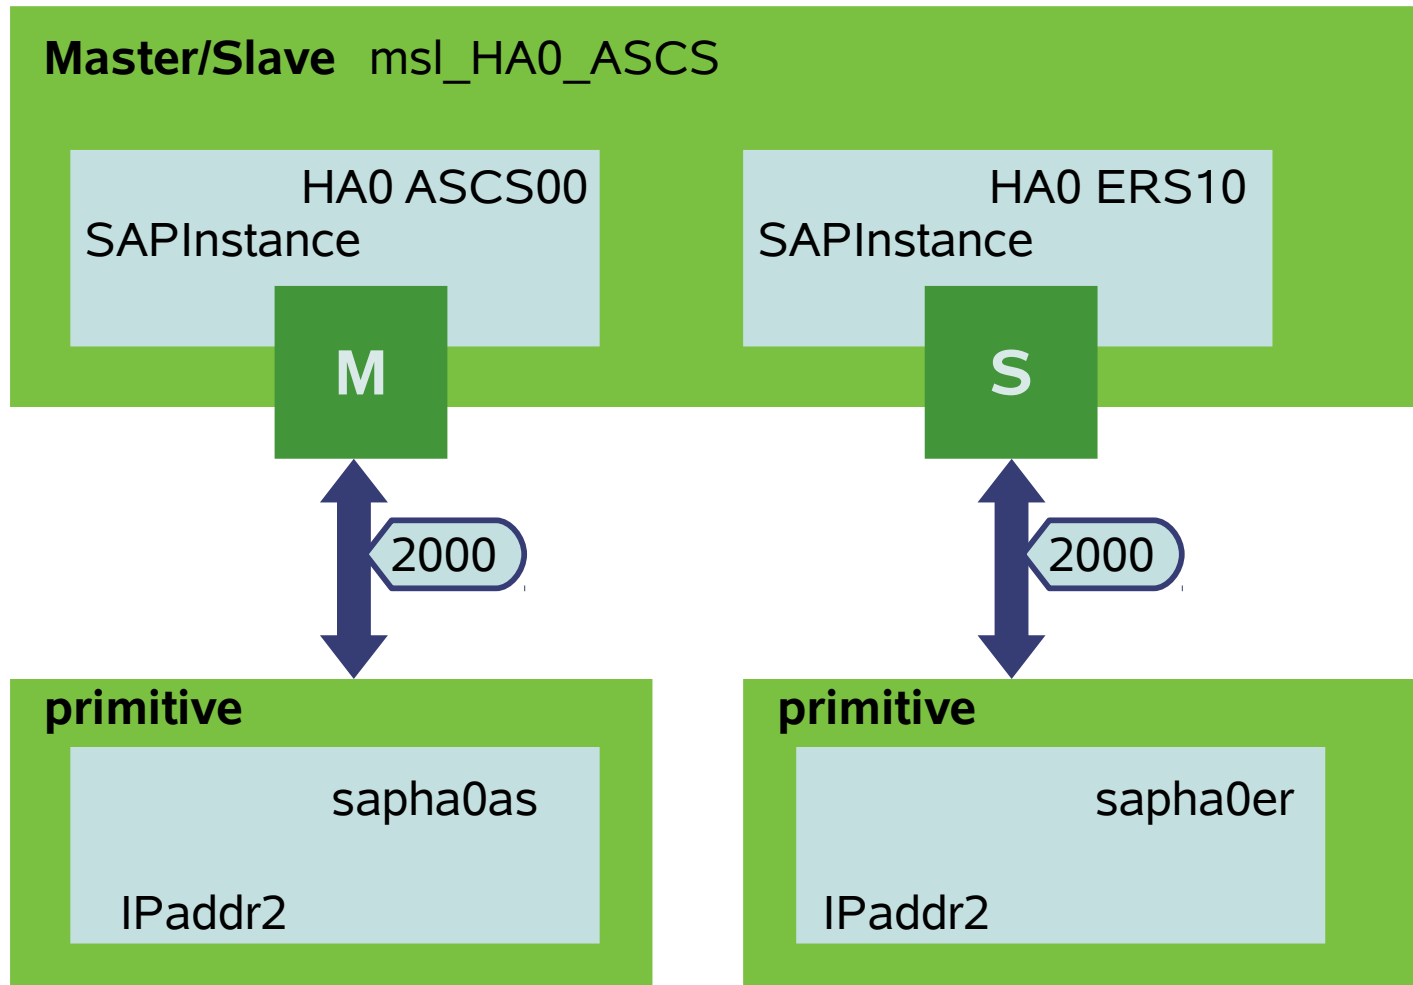
Node 1 comes back and joins the running cluster

The cluster starts the ERS resource on this “new” node

The SAP Enqueue – Replication synchronizes all lock entries

# Cluster Concept - Constraints

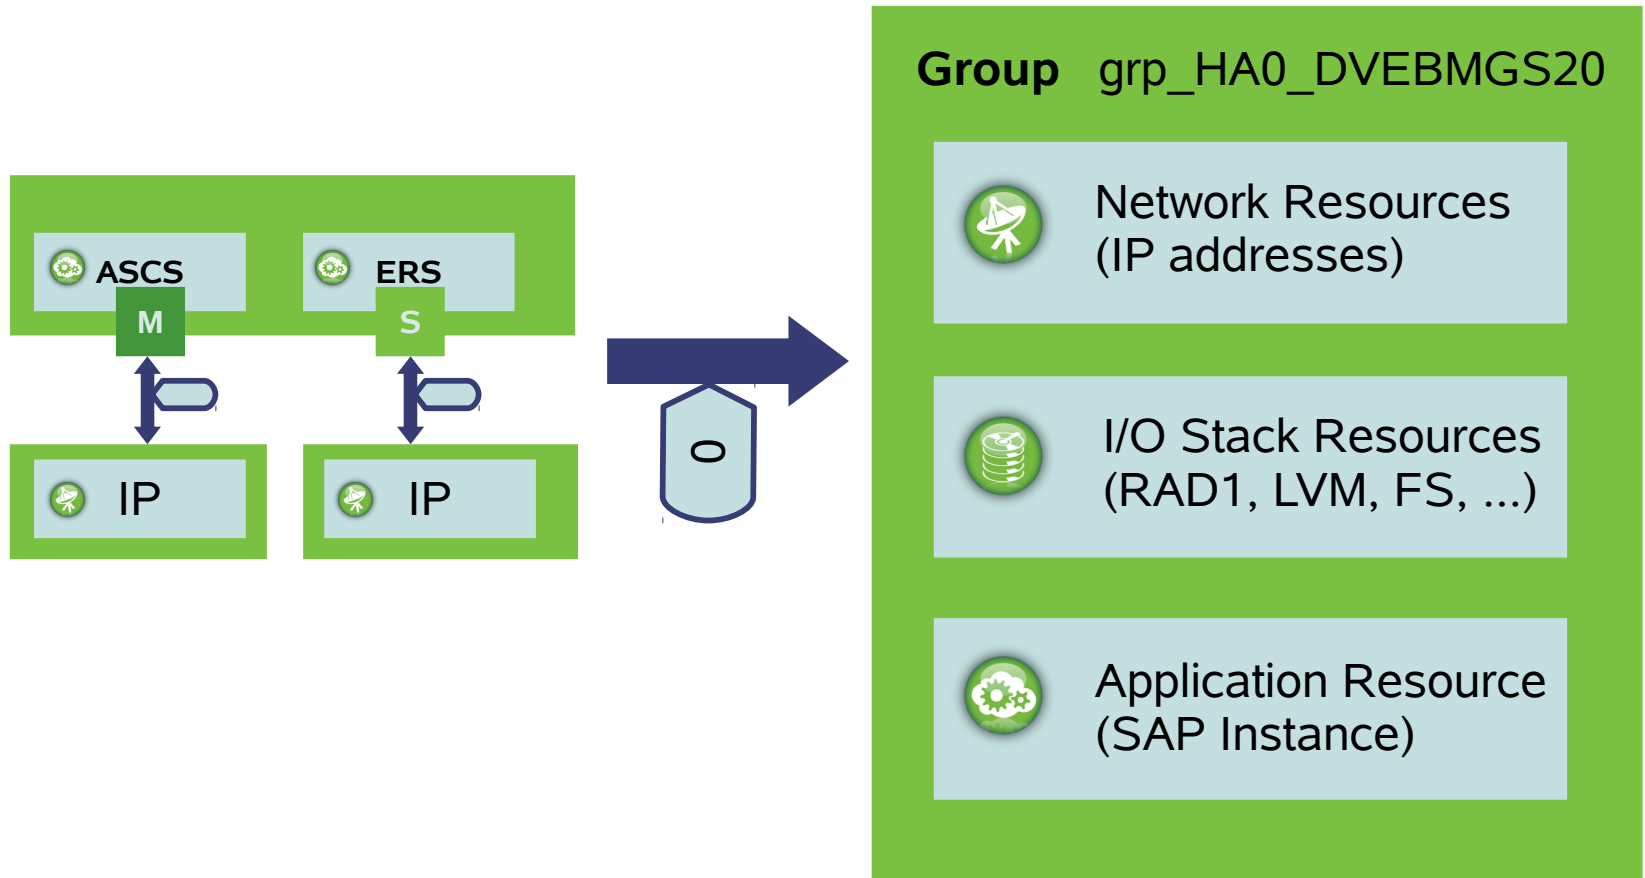
Collocation Constraints for SAP Instances and IP Addresses



Bind IP addresses to M/S status but do not stop, if instance is down

# Cluster Concept - Constraints

Order Constraints for SAP Instances and Groups




Each SAP application server needs the central services (ASCS)



# Cluster Concept - Constraints

Order Constraints for SAP Instances and SAP Database

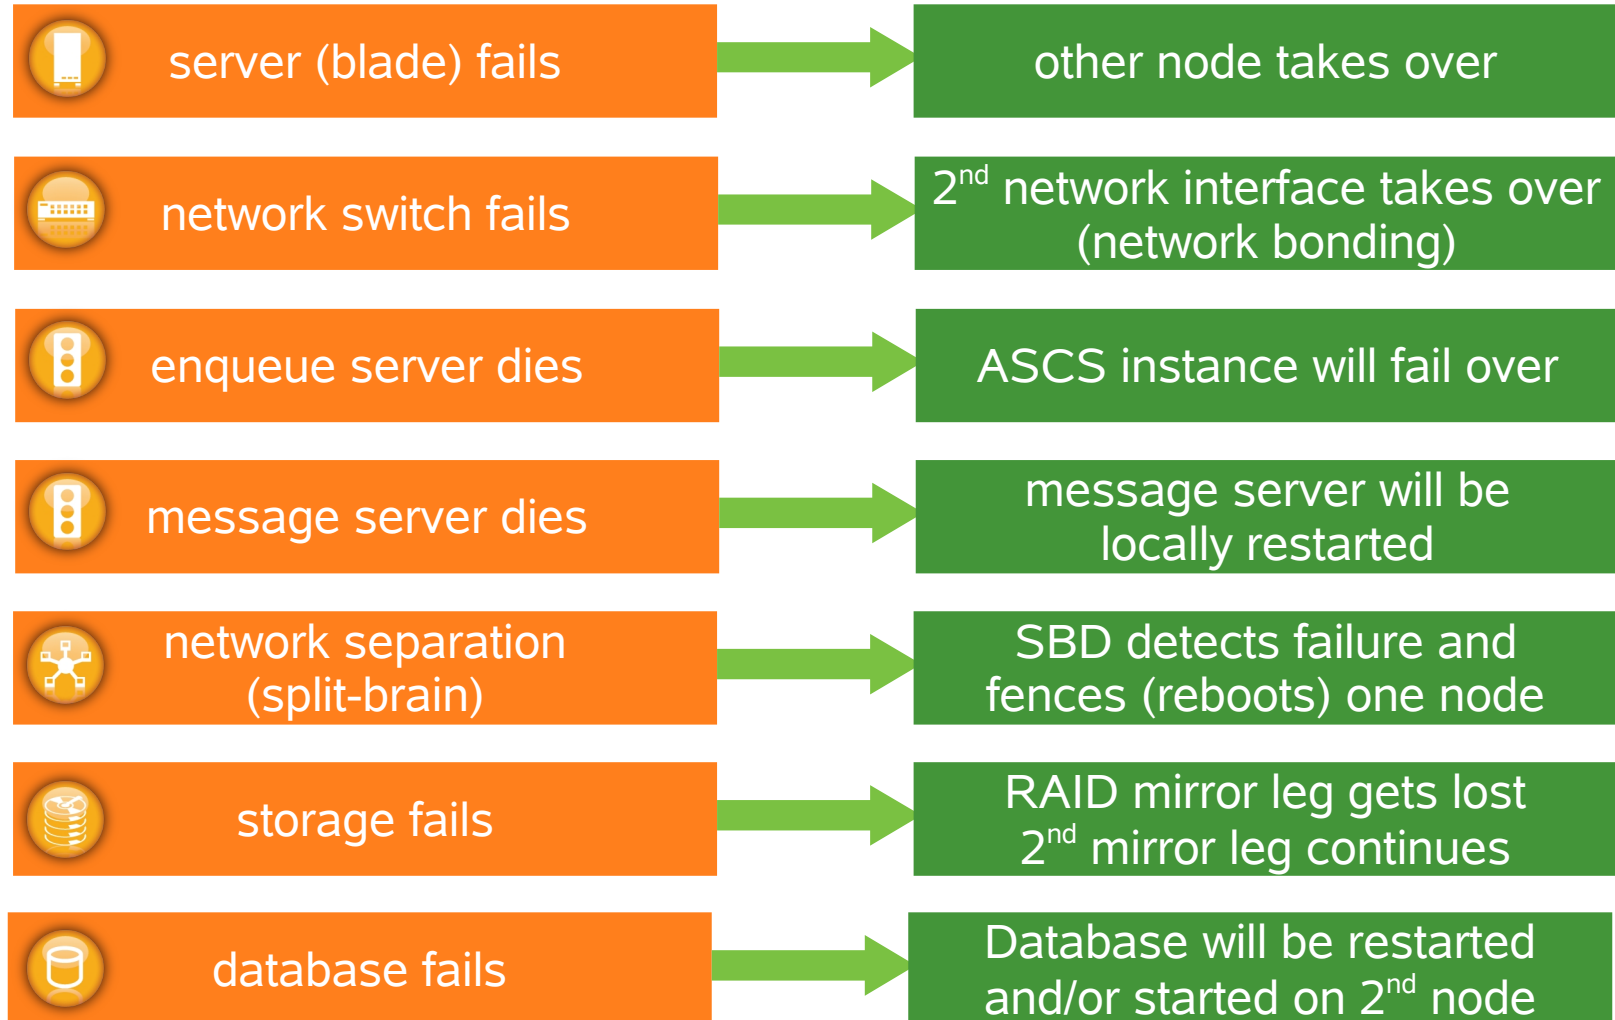


Advisory order constraints for all SAP application servers against Database and central services instance (ASCS)

Optional advisory order constraint between Database and central services instance (ASCS)

# Failures and Solutions

Some examples of “What happens when”



# The Dos and Don'ts

Things you should consider



Keep cluster configuration simple



Use SBD for node fencing (STONITH)



Define and perform tests for all failure scenarios



Follow our best practices

# The Dos and Don'ts

Things you should avoid



Build Cluster cluster without node fencing (STONITH)



Go live without tests planned and done



Go live without proper operations manual



Cluster resource (like SBD and STONITH) timings shorter than SAN timings



# SAP Certified Cluster Architecture

# SAP NetWeaver High Availability Cluster 7.30 Certification



## **Start/Stop infrastructure**

certifies the **Start/Stop infrastructure** within HA-Setups (SAP NetWeaver Management Agents, Cluster-API).



## **Reference SAP HA setup scenario**

Defines a **reference SAP HA setup scenario** that must be used by all HA vendors



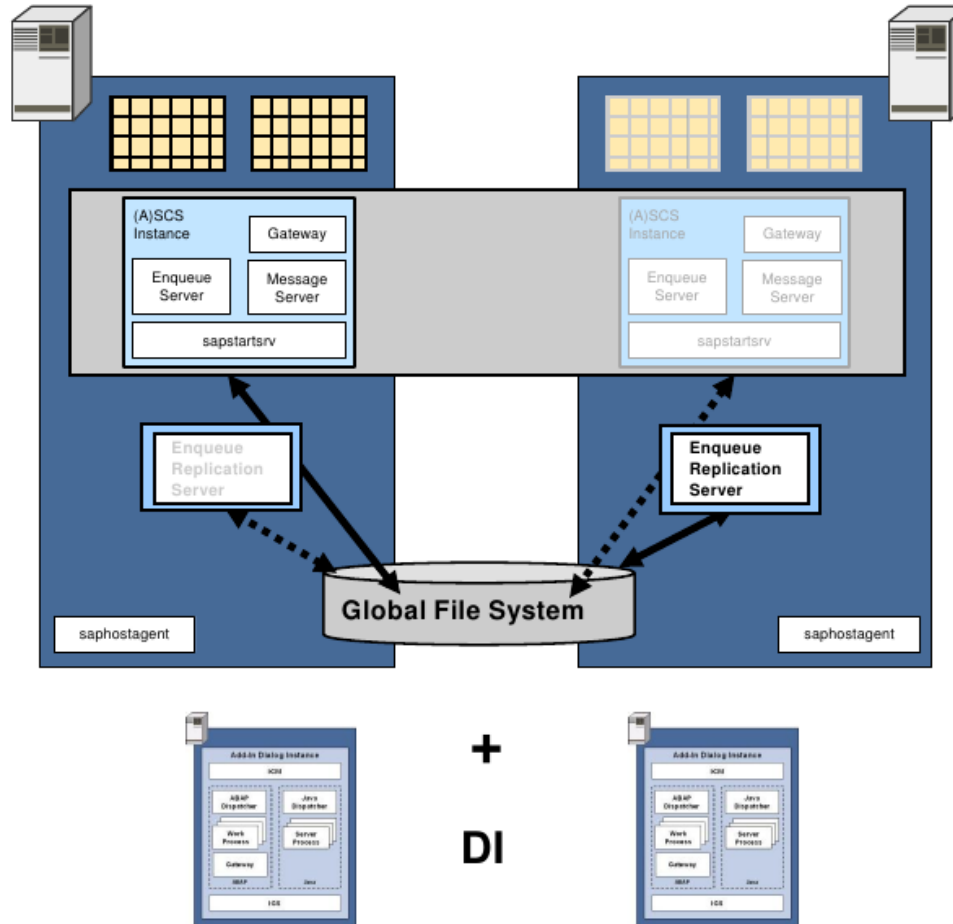
## **HA setups for SAP NetWeaver 7.x**

Unifies **HA setups for SAP NetWeaver 7.x**, based on 7.20 DCK availability for ABAP as well as for Java Application Servers.

<http://scn.sap.com/docs/DOC-26718>



# SAP System in a SAP recommended HA Setup



<http://scn.sap.com/docs/DOC-25453>



## How to Connect SAPSTARSRV and Cluster Frameworks using the Components saphascriptco.so and SAP\_Vendor\_Cluster\_Connector



### Applies to:

All SAP products controlled by sapstartsrv. For more information, visit the [Application Management homepage](#).

### Summary

Running SAP in a high availability cluster environment needs additional interfaces to communicate between the SAP program SAPSTARTSRV and the high availability cluster. At least it is essential to inform the cluster, if a SAP instance is started or shutdown. This article describes the reference interface implementation sap\_suse\_cluster\_connector.

**Author:** Fabian Herschel

**Company:** SUSE Linux Products GmbH

**Created on:** April 3, 2012

**Version:** 1.0

### Author Bio

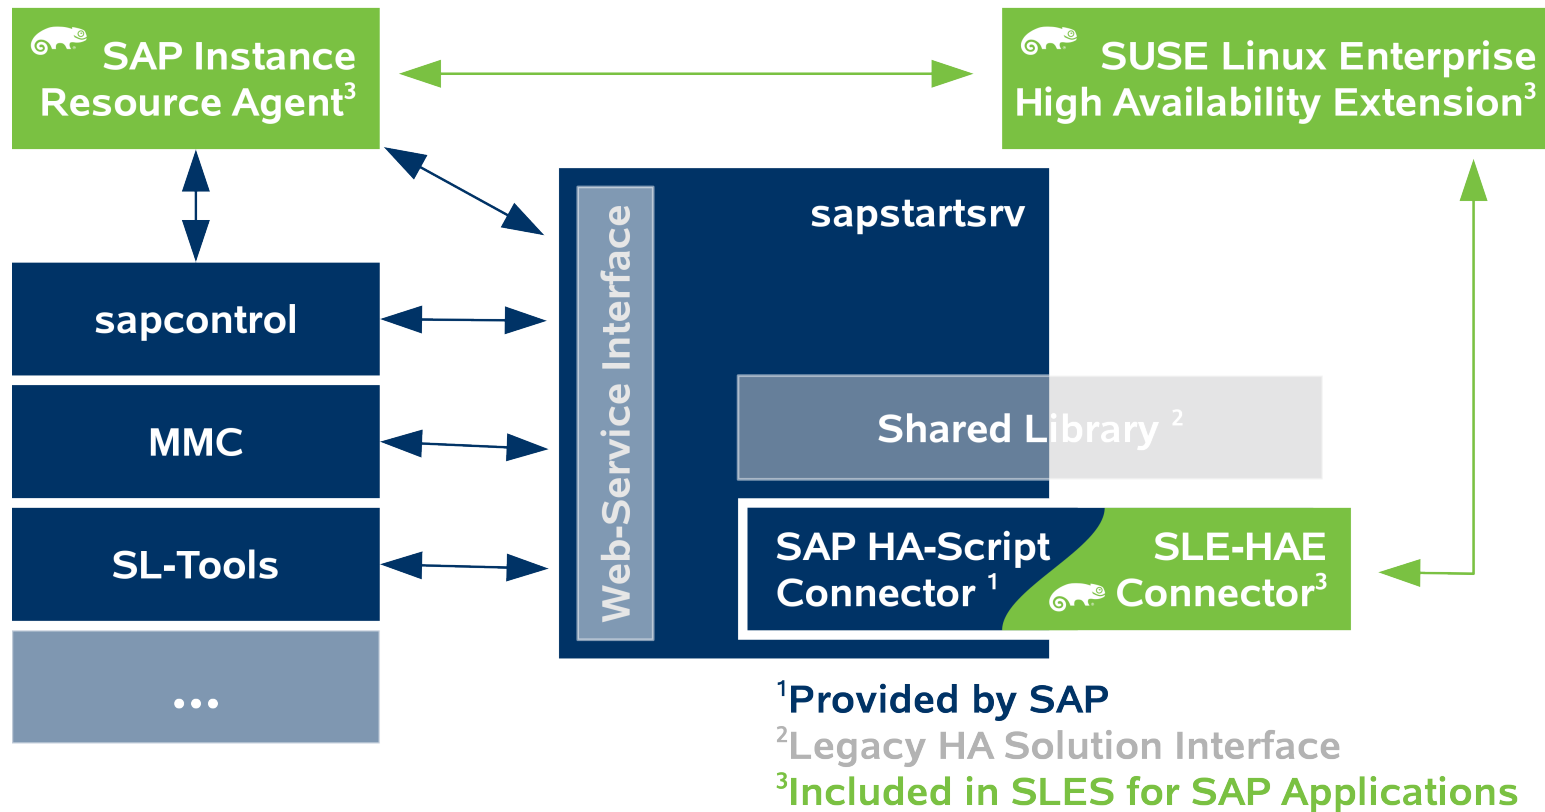


Fabian Herschel is SAP architect and employee of SUSE Linux Products GmbH in Nuremberg, Germany. He is currently working directly at the SAP LinuxLab in Rot and is leading, together with his colleague Markus Gürtler, the creation of the architectures of HA solutions for SAP products running on SUSE Linux Enterprise Server and the SUSE Linux Enterprise High Availability Extension.

<http://scn.sap.com/docs/DOC-28875>


# SAP HA Cluster Interface

Interfaces to integrate our HA solution in SAP



<http://scn.sap.com/docs/DOC-25453>

# SAP Partner Information Center

**The Best-Run Businesses Run SAP**

United States   Store   Newsletter   My SAP Profile  
1-866-609-6150   Call Me Now   Chat Now   Contact Us

● Section   ● SAP.com

Industries

Solutions >

Services >

Technology >

Partners ▾

Co-Innovation Lab (COIL)

About the SAP Ecosystem >

Global & Local Partner Directories

Integration & Certification >

SAP Pinnacle Awards

Partner Spotlight

Partnering with SAP

Referral Program >

SAP Community Network >

## PARTNER INFORMATION CENTER: SOLUTION DETAILS

Partner

Solutions

**Suse Linux Products GmbH**

**SUSE Linux Enterprise High Availability 11**

- **Certification Category:** High Availability Cluster
- **Qualifications:** SAP Certified - Integration with SAP NetWeaver

SAP Interface	SAP Release Levels	Solution Interface Software	Certified Functions	Certification Date
NW-HA-CLU 7.30 - NetWeaver High Availability Cluster 7.30	SAP NetWeaver 7.3	Interface for SUSE Linux Enterprise High 11		25 Sep 2012

<http://www.sap.com/partners/directories/searchpartner.epx>  
→ Search for *Company Name* **SUSE**

# The Certificate



## Certificate

SAP INTEGRATION CERTIFICATION

- ✓ SLES HA 11 on architectures x86\_64 and power64
- ✓ Certification valid 3 years till Sep 2015 / Aug 2016
- ✓ For SAP NetWeaver 7.30
- ✓ Certifies the SAP SUSE cluster integration



So we are...

**SAP<sup>®</sup>** Certified  
Integration with SAP NetWeaver<sup>®</sup>



SAP AG hereby confirms that the interface software  
for the product

SUSE Linux Enterprise High Availability  
11

of the company SUSE LINUX Products GmbH

has been certified for integration with SAP NetWeaver via the  
SAP integration scenario NW-HA-CLU 730.

This certificate confirms the existence of product features in accordance  
with SAP certification procedures. It does not guarantee that the product  
is error-free.

The certification test is documented in report no. 20725237  
and expires September 25, 2015.

Vendor Hardware: x86\_64 platform

Vendor Operating System: SLES 11

SAP Test System: SAP NetWeaver 7.30

Used Integration Tools: none

This configuration meets the requirements for connecting SUSE Linux Enterprise  
High Availability 11 to  
SAP NetWeaver.

Certified Functions:

Integration of sapstartsrv in vendor product

Start/Stop SAP instance via SAPControl

Manual failover test via vendor product

Walldorf, September 25, 2012

Mr. Jürgen Bierlein, SAP AG

SAP, R3, and SAP NetWeaver are registered trademarks of SAP AG Germany.  
All other names are registered or unregistered trademarks of the individual firms.  
<http://www.sap.com/loc>

**SAP<sup>®</sup>** Certified  
Integration with SAP NetWeaver<sup>®</sup>

# SAP Certified – HAWK Cluster Status

Enqueue Replication with External Database (SLE HAE 11 SP2)

The screenshot displays the SAP Cluster Status web interface. The browser address bar shows the URL: <https://ls3198v7.wdf.sap.corp:7630/main/status>. The page title is "Cluster Status".

The interface shows a list of resources and their status:

- cl2n01: Online
- cl2n02: Online
- Master/Slave Set: msl\_sap\_enqrepl\_HA0
- rsc\_sap\_HA0\_ASCS00:0: Master
- rsc\_sap\_HA0\_ASCS00:1: Slave
- stonith-sbd: Started
- rsc\_ip\_HA0\_sapha0as: Started
- rsc\_ip\_HA0\_sapha0ci: Started
- rsc\_fs\_HA0\_dvebmgs01: Started
- rsc\_sap\_HA0\_DVEBMGS01: Started
- rsc\_ip\_HA0\_sapha0d2: Started
- rsc\_fs\_HA0\_d02: Started
- rsc\_sap\_HA0\_D02: Started

The resources are organized into columns. The first column contains the status of the cluster nodes and the Master/Slave Set. The second column contains the status of the resources. The third column is labeled "Inactive Resources".

**SAP®** Certified  
Integration with SAP NetWeaver®

# Sample Configuration

Primitive: rsc\_sapinst\_HA0\_ASCS00\_sapha0as

**Resource Agent** SAPInstance

## Parameters

Parameter	Value
InstanceName	HA0_ASCS00_sapha0as
ERS_InstanceName	HA0_ERS10_sapha0er
AUTOMATIC_RECOVER	true

## Operations

Type	Interval	Timeout	Start delay	Role	On fail
monitor	120s	60s	120s		
monitor	121s	60s	120s	slave	
monitor	119s	60s	120s	master	
start	0	180s			
stop	0	240s			block






# SAP in Virtualized Environments



# Use Case: Cluster in Cluster



# VMware HA and SLE HAE

SUSE Linux Enterprise  
High Availability Extension

- \* Both SLE HA Nodes running on ESX server 1
- \* ESX Server 3 is powered down

DB  
OS

APP  
SCS  
OS

APP  
OS

APP  
OS

APP  
OS

VMware HA and DRS Cluster

VMware ESX

VMware ESX

(VMware ESX)



# VMware HA and SLE HAE

SUSE Linux Enterprise  
High Availability Extension

- \* VM is migrated to ESX server 2 without
- \* SLE HA cluster interference

DB  
OS

APP  
SCS  
OS

vMotion

APP  
SCS  
OS

APP  
OS

APP  
OS

APP  
OS

VMware HA and DRS Cluster

VMware ESX

VMware ESX

(VMware ESX)



# VMware HA and SLE HAE

SUSE Linux Enterprise  
High Availability Extension

\* SLE HA cluster now runs on different  
ESX servers to have HA against  
Hardware failures...

DB  
OS

APP  
SCS  
OS

APP  
OS

APP  
OS

APP  
OS

VMware HA and DRS Cluster

VMware ESX

VMware ESX

(VMware ESX)



# VMware HA and SLE HAE

SUSE Linux Enterprise  
High Availability Extension

\* SLE HA cluster now runs on different  
ESX servers to have HA against  
Hardware failures...

DB  
OS

APP  
SCS  
OS

APP  
OS

APP  
OS

APP  
OS

VMware HA and DRS Cluster

VMware ESX

VMware ESX

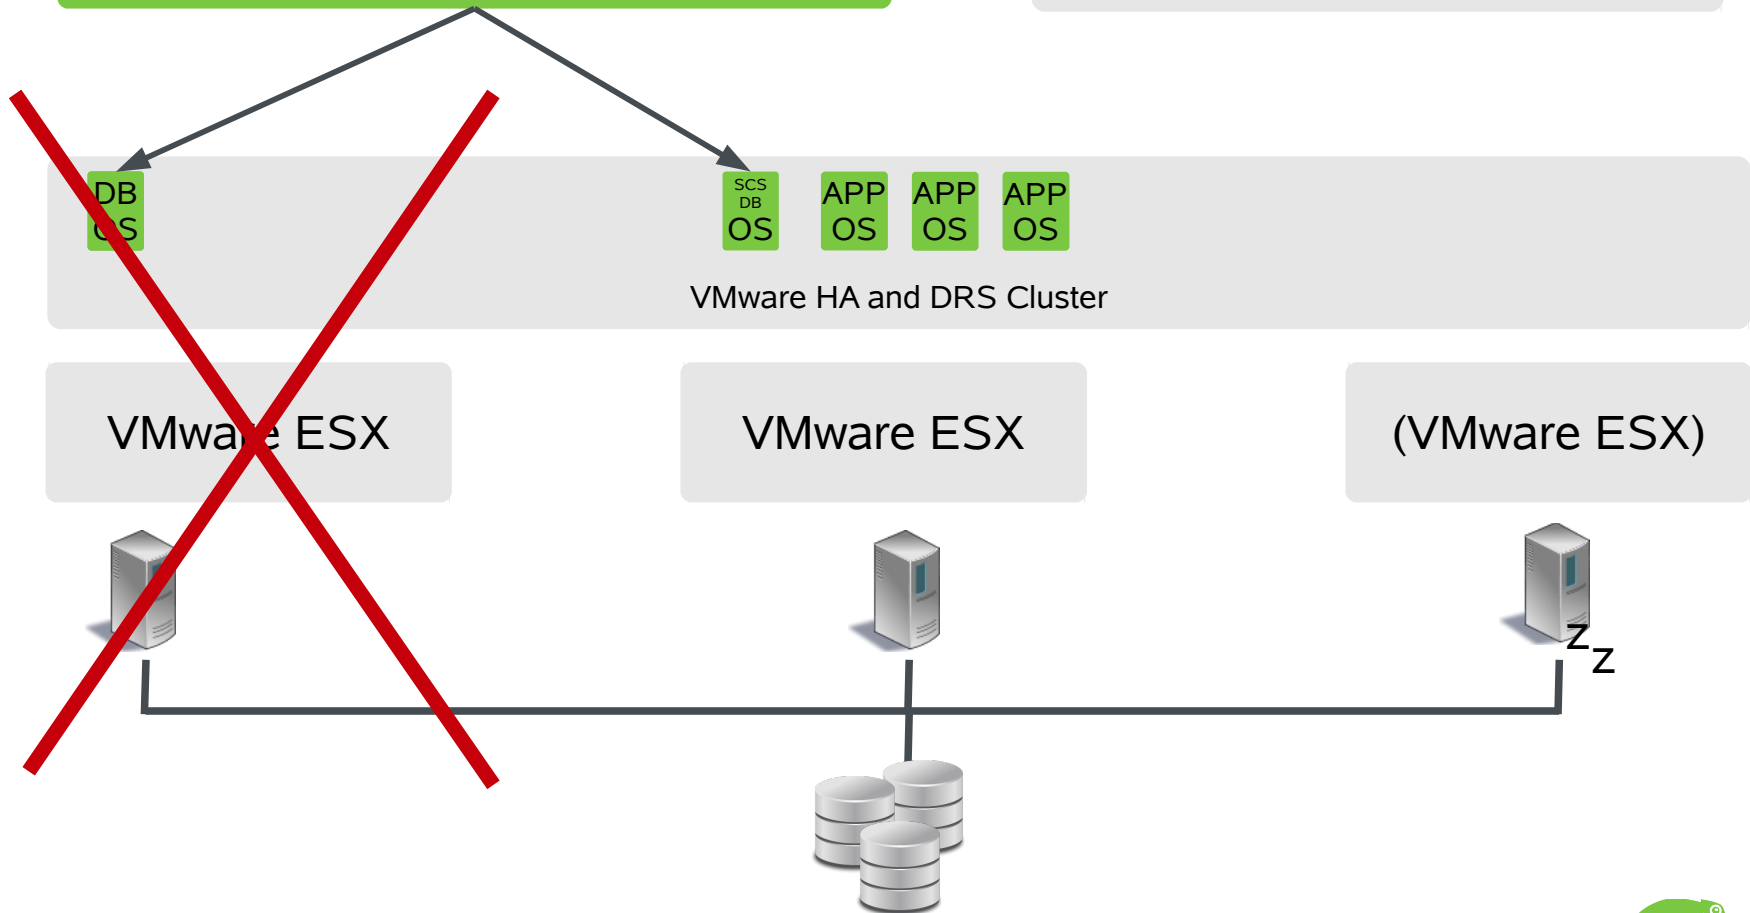
(VMware ESX)



# VMware HA and SLE HA

SUSE Linux Enterprise  
High Availability Extension

\* ... This was just in time, because  
Unfortunately a ESX hardware system fails  
\* SLE HA migrates the Database and  
optionally shutdown an Application Server



# VMware HA and SLE HA

SUSE Linux Enterprise  
High Availability Extension

- \* ESX server 1 is now in hardware Maint.
- \* VMware DPM powers up ESX server 3
- \* Failed Virtual Machines get started by VMware HA

SCS  
DB  
OS

APP  
OS

APP  
OS

APP  
OS

OS

VMware HA and DRS Cluster

(VMware ESX)

VMware ESX

VMware ESX

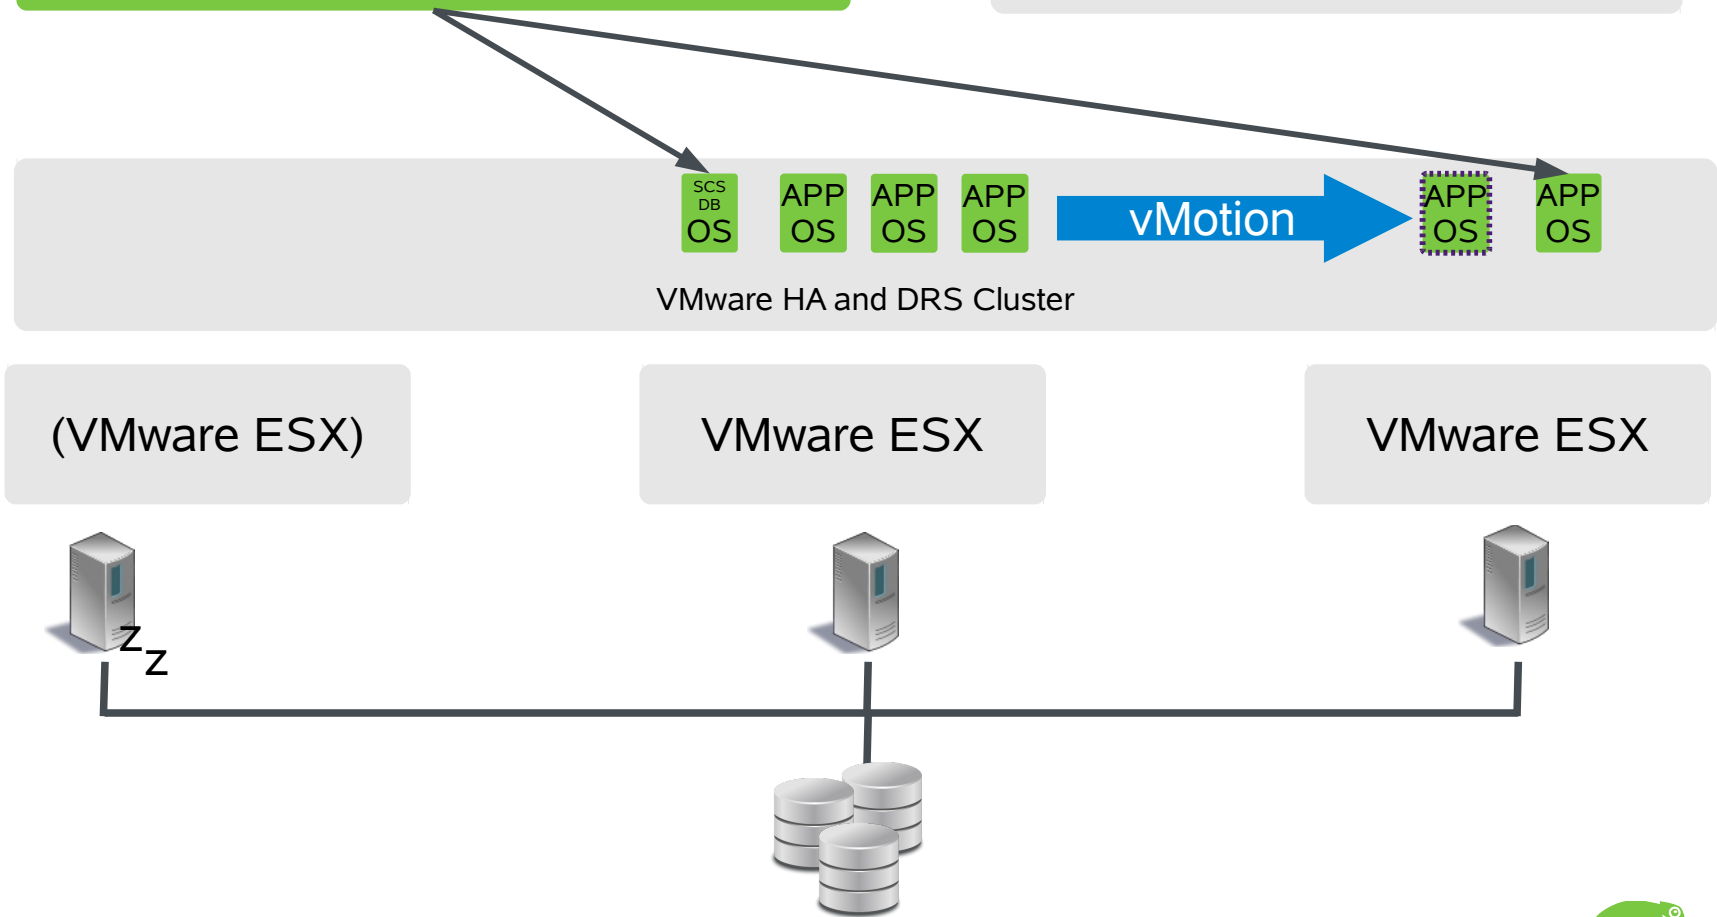




# VMware HA and SLE HA

SUSE Linux Enterprise  
High Availability Extension

\* One of the virtual machines with an SAP application server is migrated to ESX server 3  
\* SLE HA starts the SAP application Server on the cluster node



# VMware HA and SLE HA

SUSE Linux Enterprise  
High Availability Extension

\* Migration is ready with complete  
business continuity

SCS  
DB  
OS

APP  
OS

APP  
OS

APP  
OS

APP  
OS

APP  
OS

VMware HA and DRS Cluster

(VMware ESX)

VMware ESX

VMware ESX






## Outlook: SAP HANA in the Cluster

# Outlook: HANA in a Cluster

## HANA Single Box – Cold Takeover



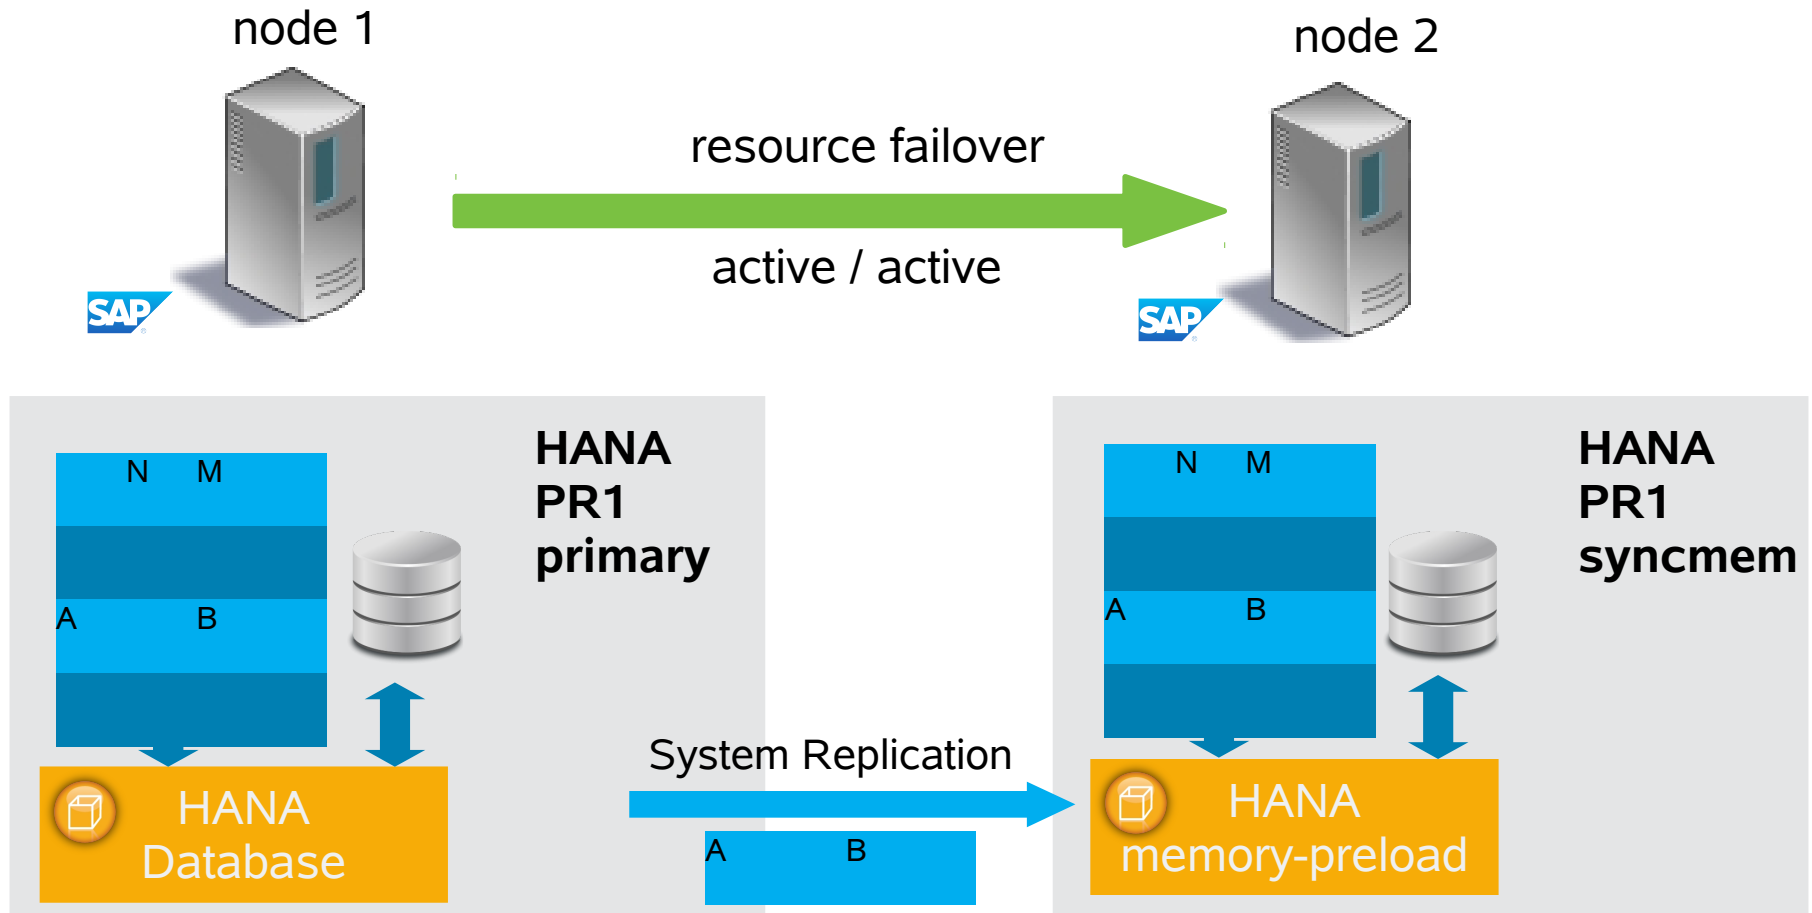
Productive HANA DB on node1

Take over takes quite long, because data need to be loaded into memory

Optional QA HANA DB on second node2

# Outlook: HANA in a Cluster

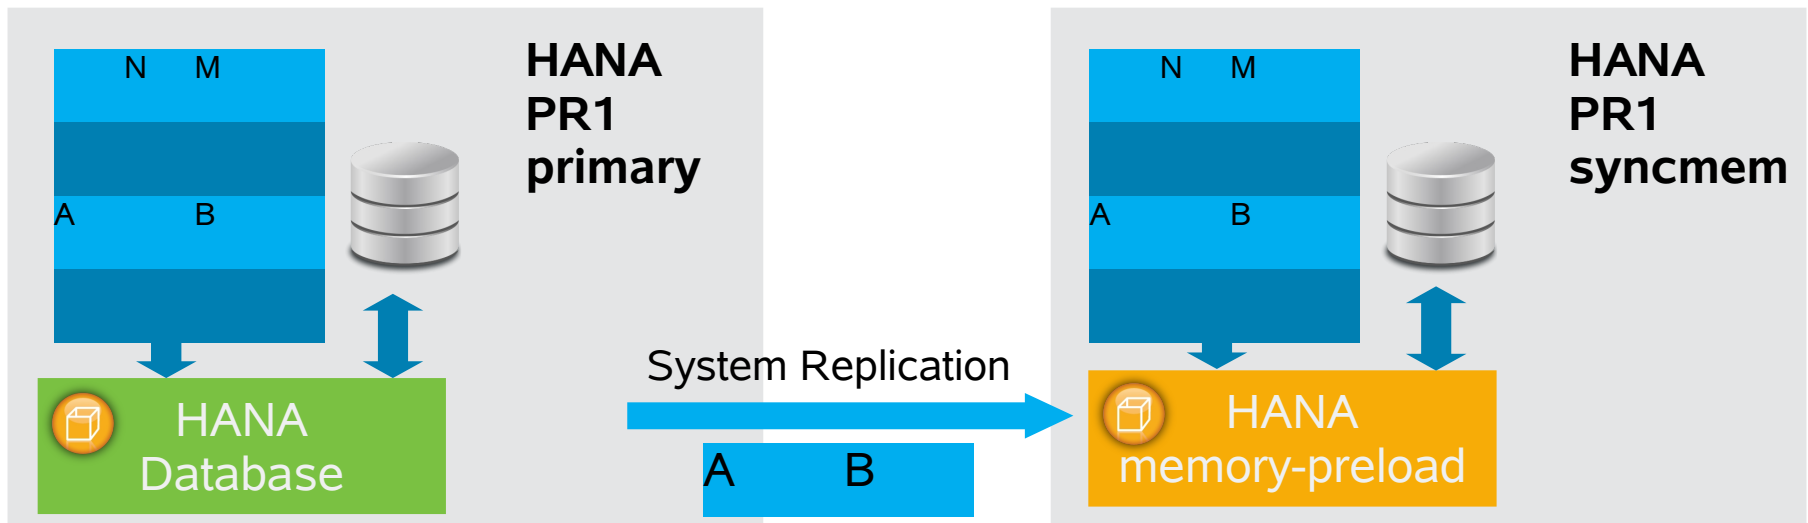

HANA “Single Box – System Replication”



This cluster concept is still under scoping and research.

# Outlook: HANA in a Cluster

HANA “Multi Node – System Replication”



This cluster concept is still under scoping and research.

Find Our Best Practices at  
[www.suse.com/products/sles-for-sap/resource-library/](http://www.suse.com/products/sles-for-sap/resource-library/)

Thank you.







**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)

## **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.

