

# SAP HANA System Replication and SUSE® Linux Enterprise Server for SAP Applications

Architecture, Best Practices and Components

**Fabian Herschel, SUSE**

SAP Senior Architect LinuxLab

Fabian.herschel@suse.com



# Agenda



**Concept**



**Delivery**



**Prerequisites and Parameters**



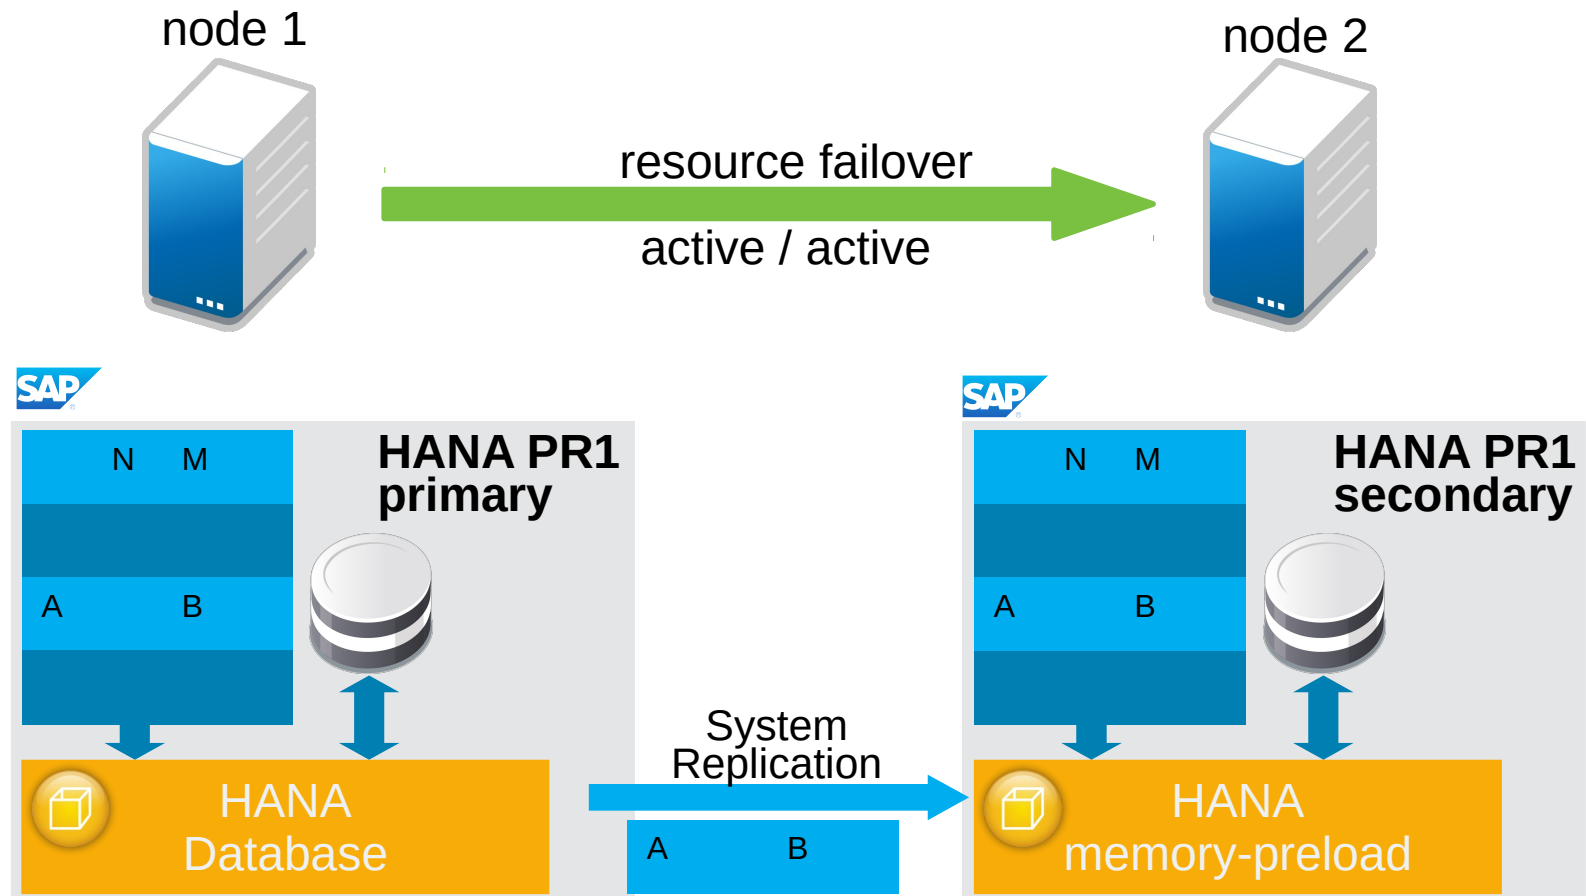
**Outlook**



**Next-Steps**

# HANA in a SUSE® Linux Enterprise High Availability Extension Cluster

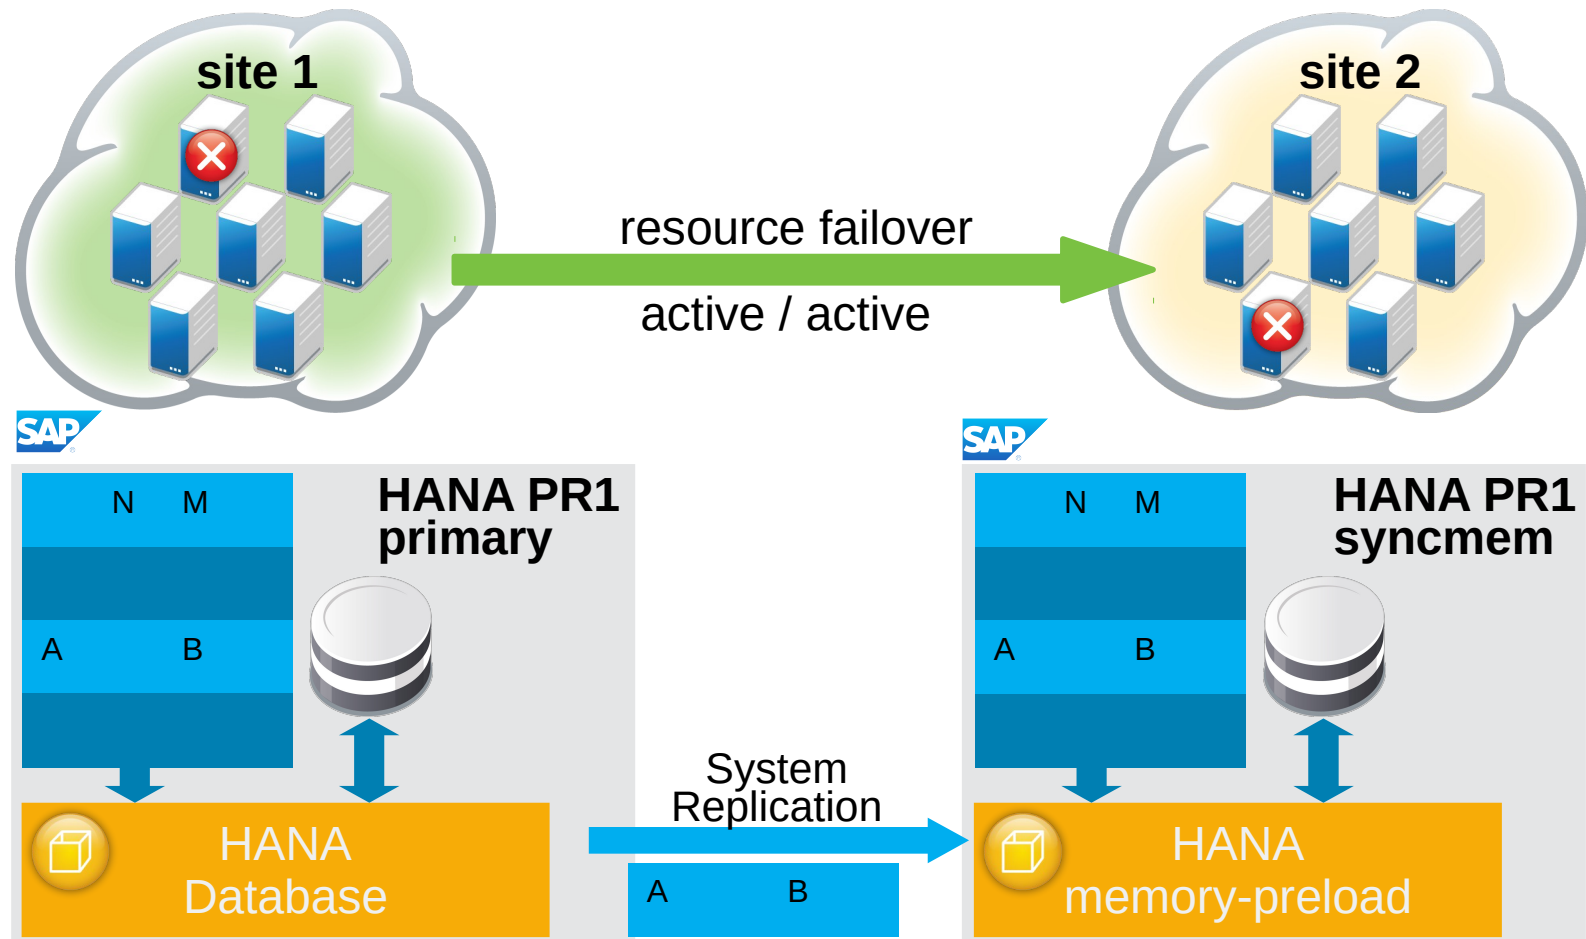
HANA Single Box – System Replication / Scale-UP



SUSE is still testing this cluster concept

# Outlook: HANA in a SUSE Linux Enterprise High Availability Extension Cluster

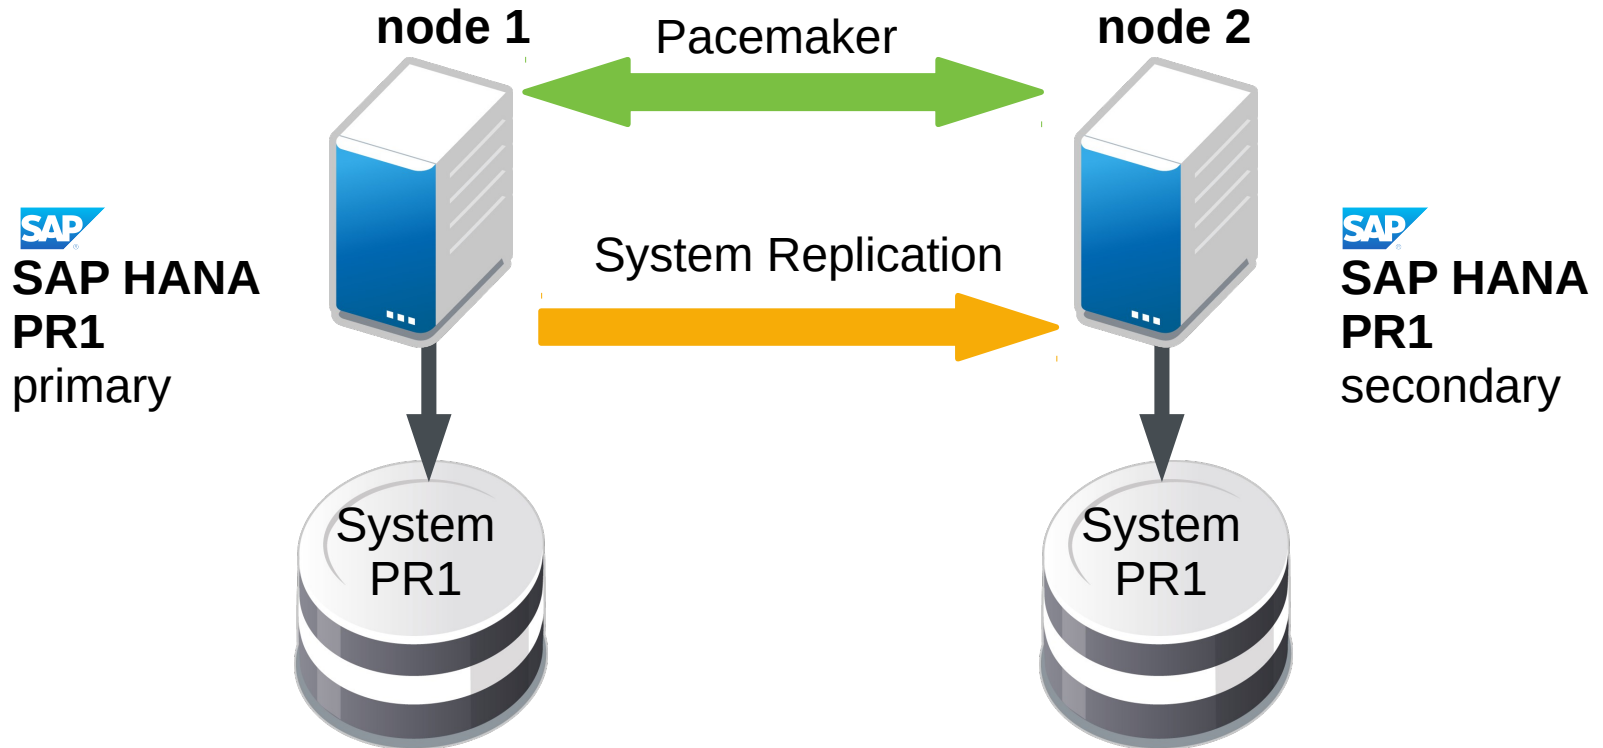
HANA Multi Node – System Replication / Scale-OUT



SUSE is still scoping and researching this cluster concept.

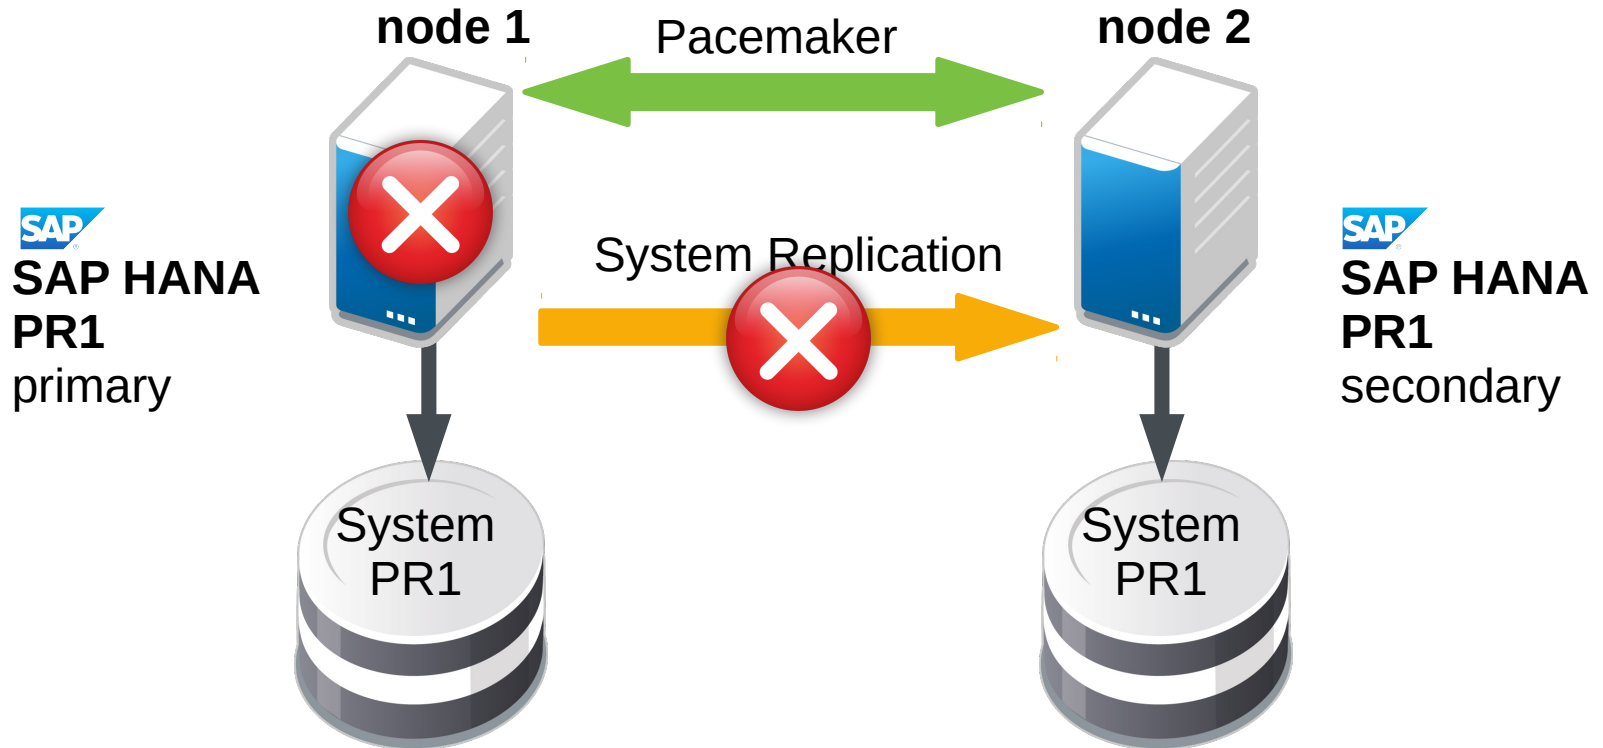
# SAP HANA SR and SUSE Linux Enterprise High Availability Extension Cluster

HANA Single Box



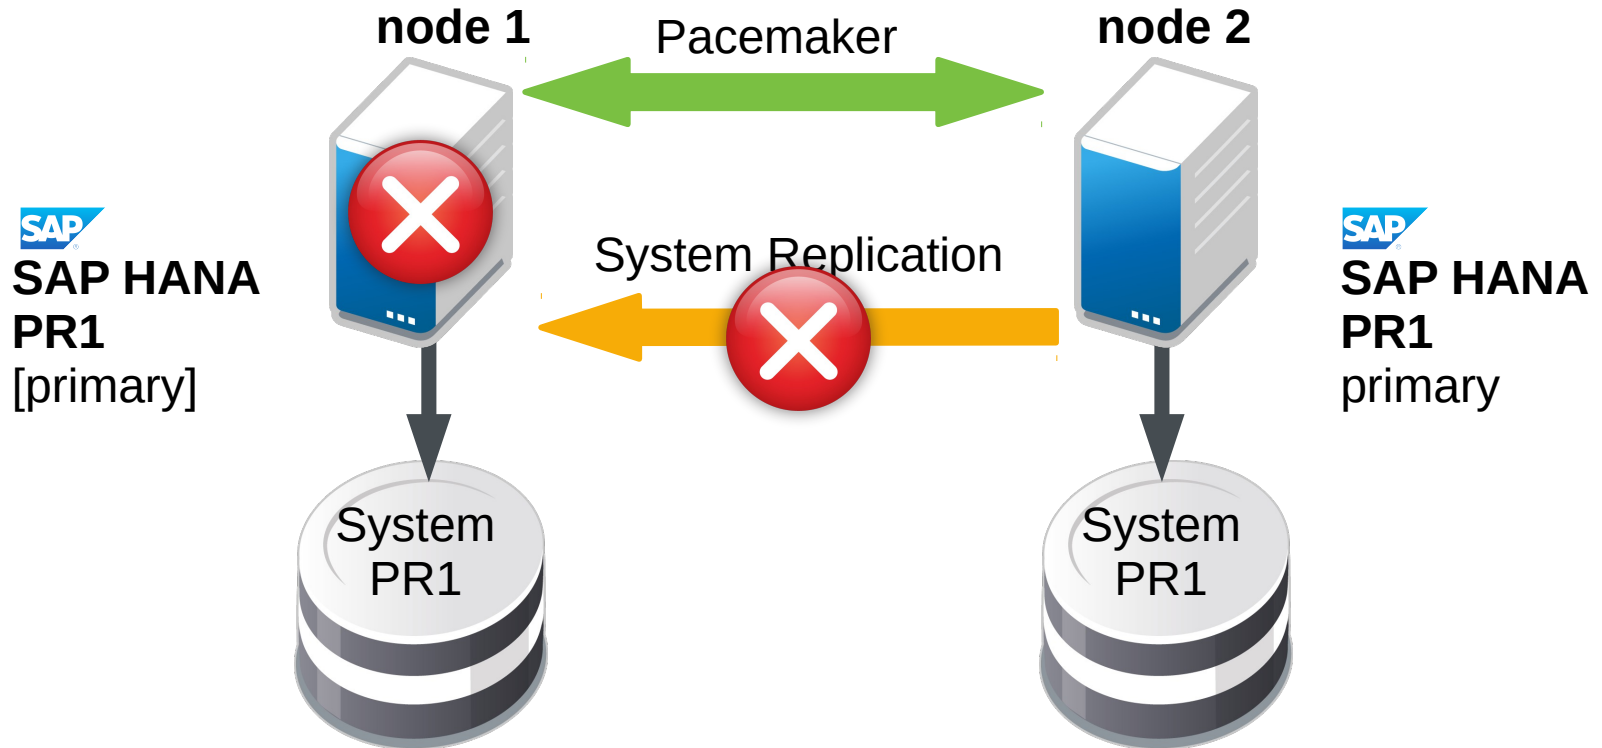
# SAP HANA SR and SUSE Linux Enterprise High Availability Extension Cluster

HANA Single Box



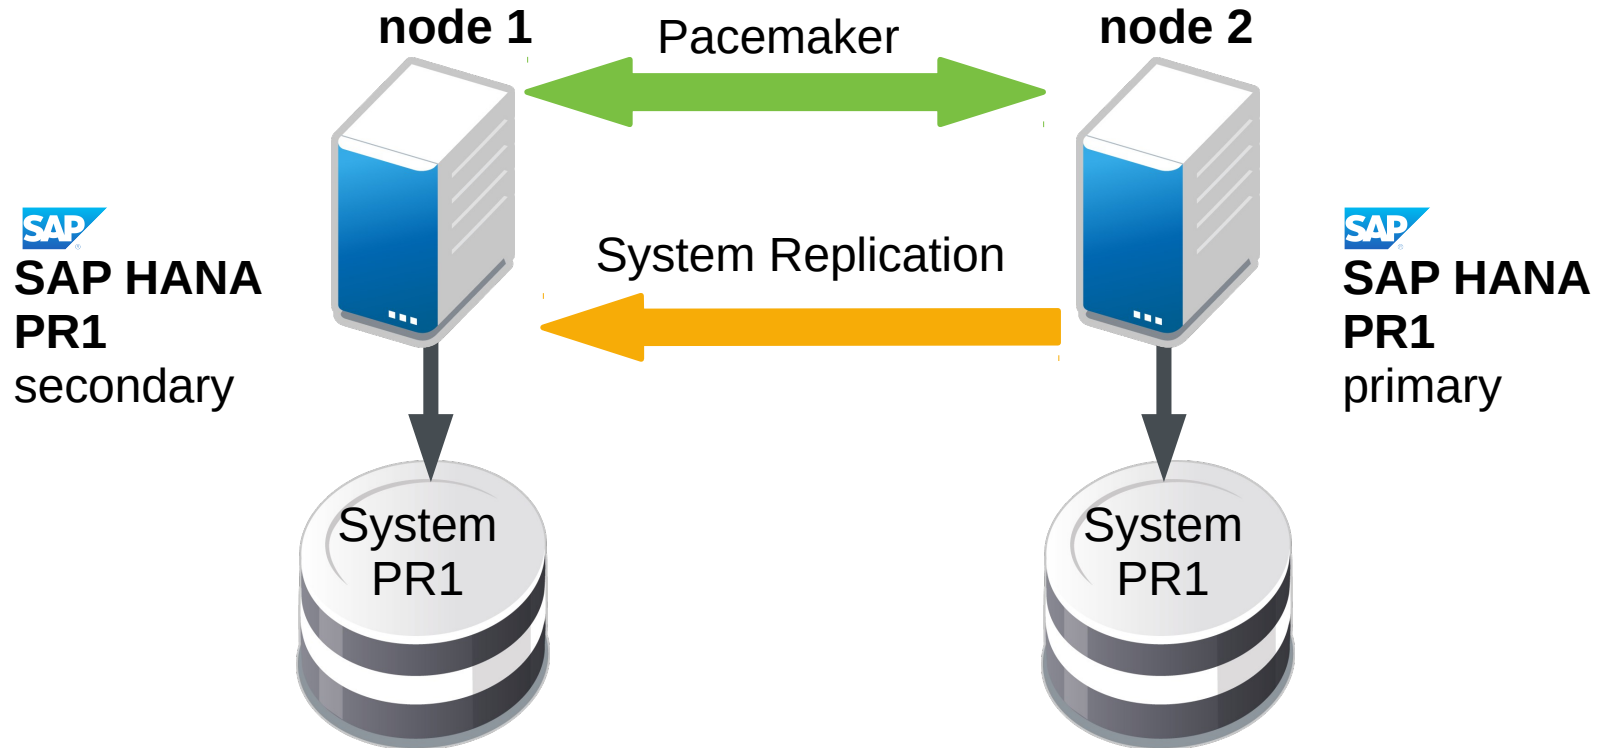
# SAP HANA SR and SUSE Linux Enterprise High Availability Extension Cluster

HANA Single Box



# SAP HANA SR and SUSE Linux Enterprise High Availability Extension Cluster

HANA Single Box

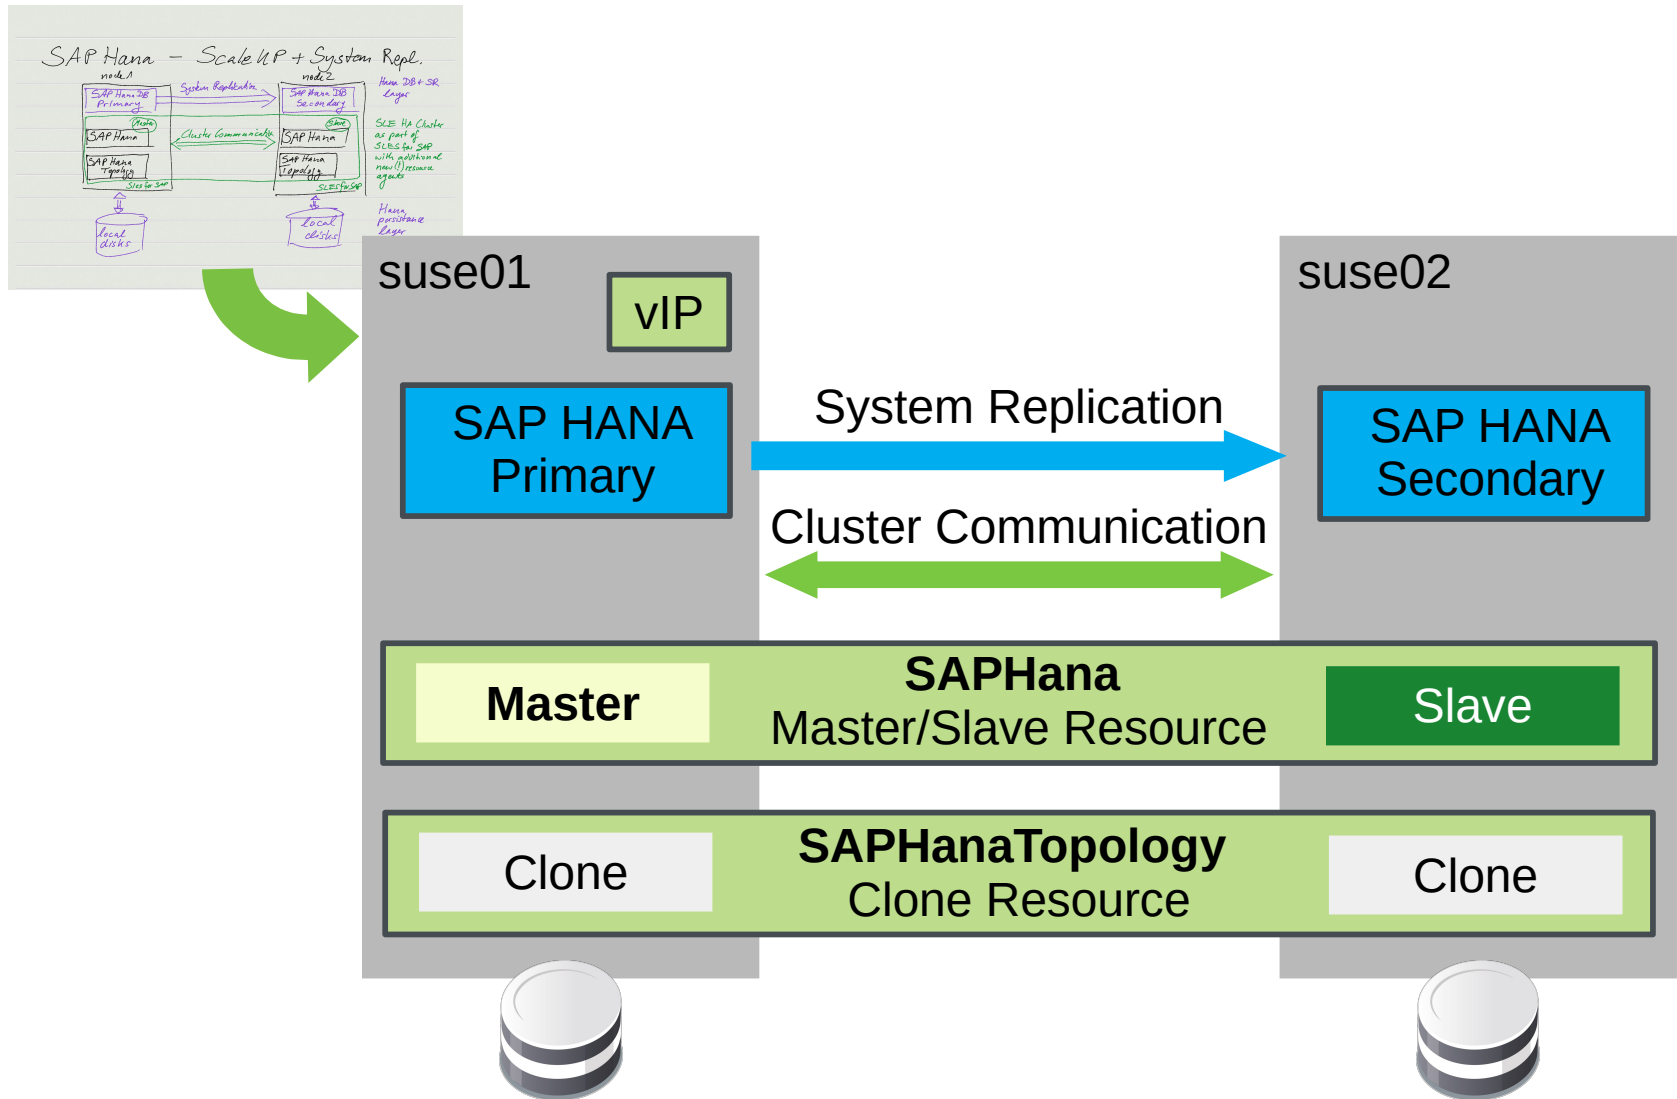


Direction of the system replication will only be changed if the parameter `AUTOMATED_REGISTER` is been changed to "true."

We recommend starting with the default: `AUTOMATED_REGISTER="false"`




# From Concept to Implementation



# What is the Delivery?

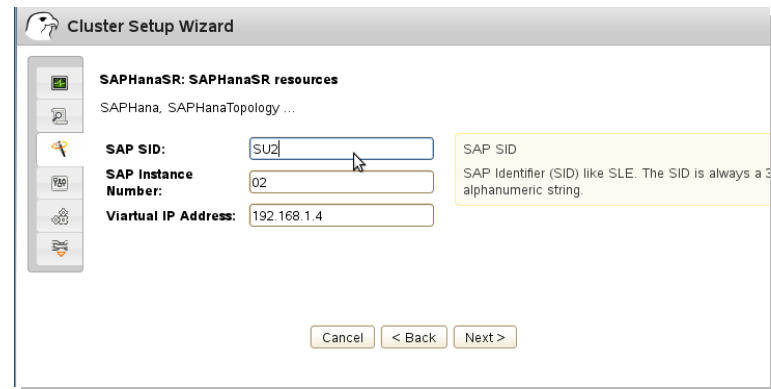
Package **SAPHanaSR** with two resource agents:  
**SAPHanaTopology** and **SAPHana**

## Setup Guide



and

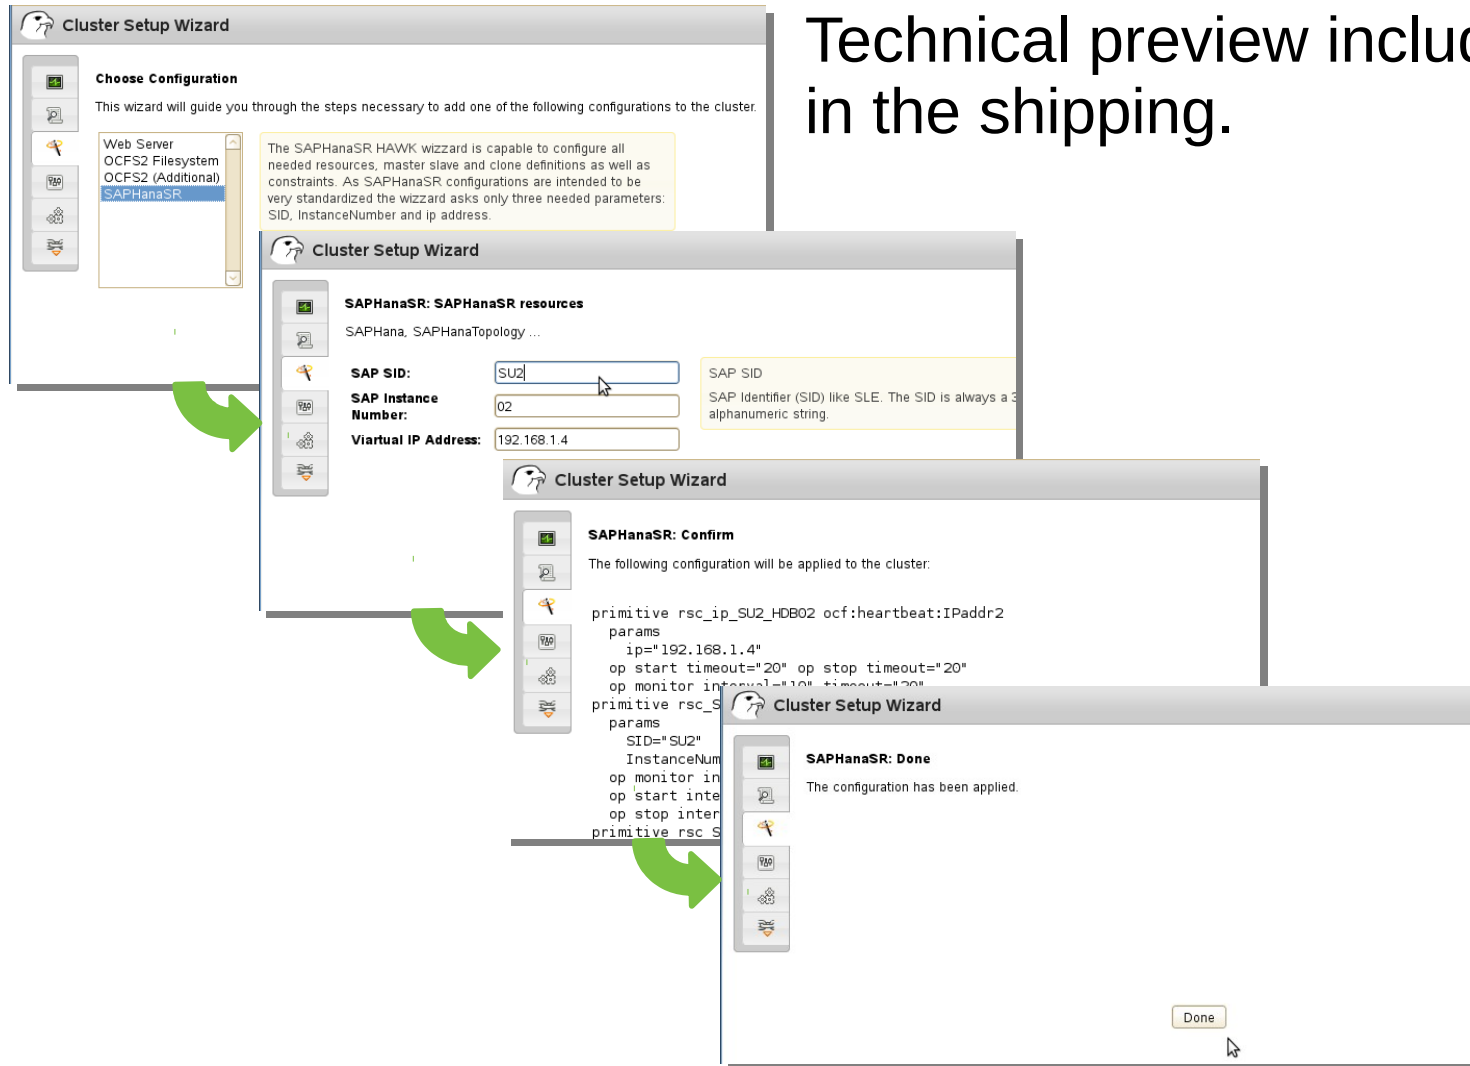
## SAPHanaSR HAWK Wizard



SAPHanaSR HAWK Wizard  
is a technical preview.

# SAPHanaSR HAWK Wizard

Technical preview included  
in the shipping.



# SAPHanaSR in HAWK

The screenshot shows the HAWK Cluster Status web interface running in a Mozilla Firefox browser. The browser's address bar shows the URL "Cluster Status". The page title is "Cluster Status". The user is logged in as "haclient" and can click "Log Out".

The interface displays the status of various resources in a table-like structure. The resources are organized into two main sections: "Master/Slave Set" and "Clone Set".

**Master/Slave Set: msl\_SAPHana\_NDB\_HDB00**

- fscs98: Online
- fscs99: Online

**Clone Set: cln\_SAPHanaTopology\_NDB\_HDB00**

- rsc\_SAPHana\_NDB\_HDB00:0: Slave
- rsc\_SAPHana\_NDB\_HDB00:1: Master
- rsc\_SAPHanaTopology\_NDB\_HDB00:0: Started
- rsc\_SAPHanaTopology\_NDB\_HDB00:1: Started
- stonith\_fscs99: Started
- stonith\_fscs98: Started
- rsc\_ip\_NDB\_HDB00: Started

On the right side of the interface, there is a section titled "Inactive Resources" with three icons representing different resource states.

# Current Allowed Scenarios

- Two-node clusters
- Scale-up (single-box to single-box) HANA system replication
- Both nodes are in the same network segment (layer 2)
- Preferred site takeover active. There is no other SAP HANA system (like QA) on the replicating node that needs to be stopped during takeover (requires additional testing)
- Both SAP HANA instances have the same SAP Identifier (SID) and Instance Number

# Current Allowed Scenarios (Cont.)

- Both cluster nodes in-time sync (ntp)
- Technical users and host names resolved locally
- Both physical and “virtual” SAP host names
- Distance / Latencies. If the cluster nodes are installed in different data centers or data center areas, the environment must match the requirements of the SUSE Linux Enterprise High Availability Extension cluster. This specifically means the network latencies between the nodes and the recommended maximum distance. Please review our product documentation for recommendations.

# Parameters for SAP HANA Topology

Parameter	Description
SID	SAP identifier (SLE)
InstanceNumber	SAP instance number (00)

# Parameters for SAPHana

Parameter	Description
SID	SAP identifier (SLE)
InstanceNumber	SAP instance number (00)
PREFER_SITE_TAKEOVER	Defines, if <i>Secondary</i> should take over, if <i>primary</i> instance (not node) fails otherwise prefer local restart of SAP HANA (true)
AUTOMATED_REGISTER	Defines if a former <i>Primary</i> should be registered as new <i>Secondary</i> after former <i>Secondary</i> completed the takeover and became new <i>Primary</i> (false)
DUPLICATE_PRIMARY_TIMEOUT	Defines the the lpa time gap needed between two primaries to select the “correct” primary.



# The Five Interfaces

**HANA Startframework:** sapstartsrv / sapcontrol / HDB  
(calls, output format “GetProcessList”)

**HANA-Topology:** landscapeHostConfiguration.py  
(rc, output format)

**SR-Topology:** hdbnsutil  
(calls, output format “-sr\_state --sapcontrol=1”)

**SAP Hostagent:** saphostctrl  
(call, output format “ListInstances”)

**SR-Status:** hdbsql (now) / systemReplicationStatus.py  
(future) (now; rc, calls, output format)

# Preview Scale-Out

Every 2.0s: bash ./crm\_mon+attr  
17:17:21 2014

Wed Apr 2

Last updated: Wed Apr 2 17:17:21 2014  
Last change: Wed Apr 2 17:14:27 2014 by hacluster via crmd on lv9042  
Stack: classic openais (with plugin)  
Current DC: lv9048 - partition with quorum  
Version: 1.1.10-65bb87e

6 Nodes configured, 6 expected votes

14 Resources configured

Online: [ lv9041 lv9042 lv9048 lv9049 lv9050 lv9051 ]

Full list of resources:

stonith-sbd (stonith:external/sbd): Started lv9042

rsc\_ip\_LNX\_HDB42 (ocf::heartbeat:IPaddr2): Started lv9049

Master/Slave Set: msl\_SAPHana\_LNX\_HDB42 [rsc\_SAPHana\_LNX\_HDB42]

Masters: [ lv9049 ]

Slaves: [ lv9041 lv9042 lv9048 lv9050 lv9051 ]

Clone Set: cln\_SAPHanaTopology\_LNX\_HDB42 [rsc\_SAPHanaTopology\_LNX\_HDB42]

Started: [ lv9041 lv9042 lv9048 lv9049 lv9050 lv9051 ]

Host \ Attr	clone_state	remoteHost	roles	site	srmode	sync_state	vhost
-------------	-------------	------------	-------	------	--------	------------	-------

lv9041	DEMOTED	lv9051	3:S:master1:master:standby:master	WALLDORF	sync	SOK	lv9041 30
lv9042	WAITING		3:P:master1:slave:worker:standby	ROT	sync	PRIM	lv9042
1396451811							
lv9048	DEMOTED	lv9051	3:S:master3:slave:worker:slave	WALLDORF	sync	SFAIL	lv9048 30
lv9049	PROMOTED	lv9041	3:P:master2:master:standby:master	ROT	sync	PRIM	lv9049
1396451811							
lv9050	DEMOTED	lv9051	3:S:master2:slave:worker:standby	WALLDORF	sync	SFAIL	lv9050 30
lv9051	DEMOTED	lv9041	3:P:master3:slave:worker:slave	ROT	sync	PRIM	lv9051
1396451805							

rsc\_SAPHana\_LNX\_HDB42:4 promotion score on lv9049: 154  
rsc\_SAPHana\_LNX\_HDB42:0 promotion score on lv9041: 100  
rsc\_SAPHana\_LNX\_HDB42:5 promotion score on lv9051: 60  
rsc\_SAPHana\_LNX\_HDB42:3 promotion score on lv9042: -9000  
rsc\_SAPHana\_LNX\_HDB42:1 promotion score on lv9048: -INFINITY  
rsc\_SAPHana\_LNX\_HDB42:2 promotion score on lv9050: -INFINITY

# Next Steps

Feedback from SAP, partners and customers

Integration of the solution in partner products

Upstream open-source project

Scoping, discussing and implementing Scale-Out

Multiple-Synchronisations (chain, star, ...)

Additional scenarios (TST/QAS, multi node, ...)

**You are invited to join  
the community :-)**

**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.

