

SUSE® Linux Enterprise 11

Platform for SAP

Certified Architecture Overview

Fabian Herschel

SUSE Linux GmbH

Lars Pinne

SUSE Linux GmbH



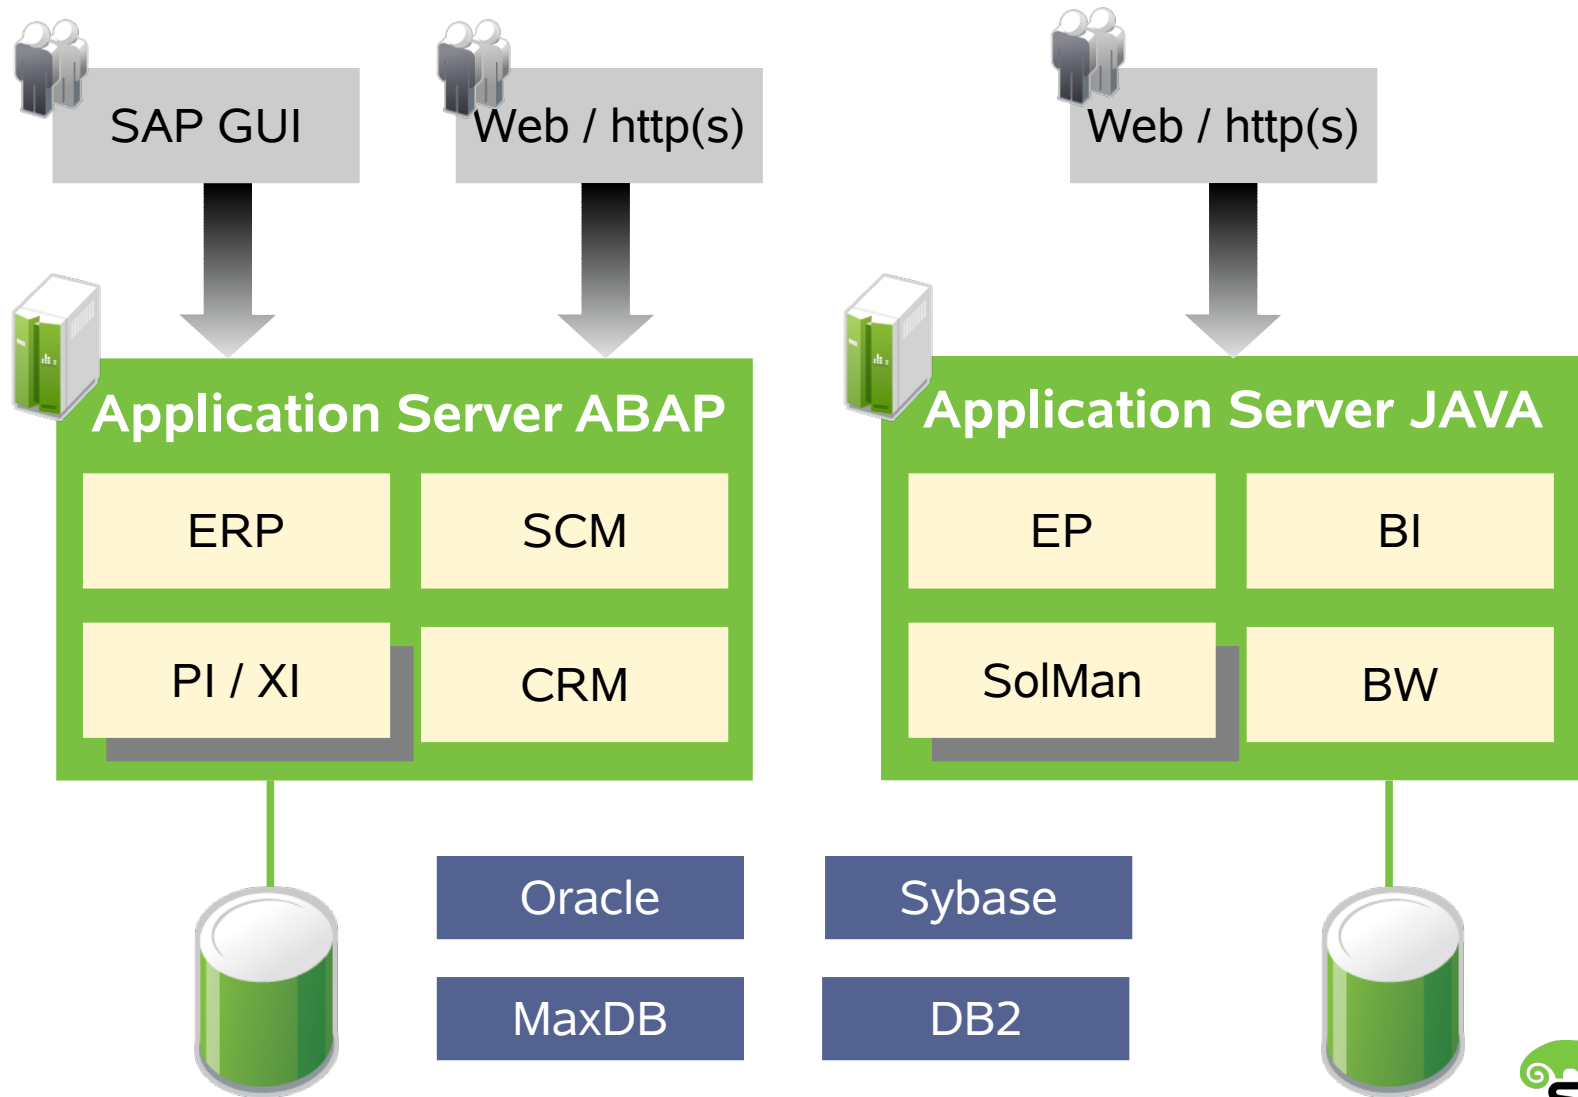
Agenda

- Certified Architecture for SAP
- Storage and Network
- Cluster Manager Pacemaker/Corosync
- Cluster Test and Maintenance
- SAP Dual Stack Enqueue Replication
- Appendix


SAP® Certified
Integration with SAP NetWeaver®

Architecture for SAP

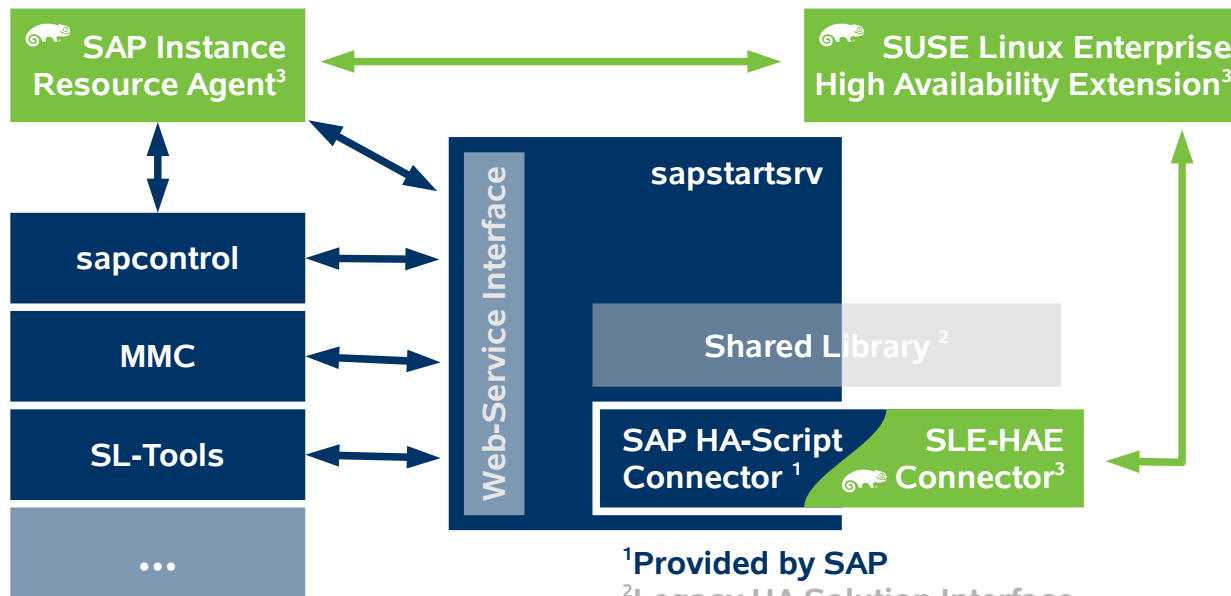
SAP Application Server Overview



Enqueue Replication – Use Case



Architecture: Design and Goals

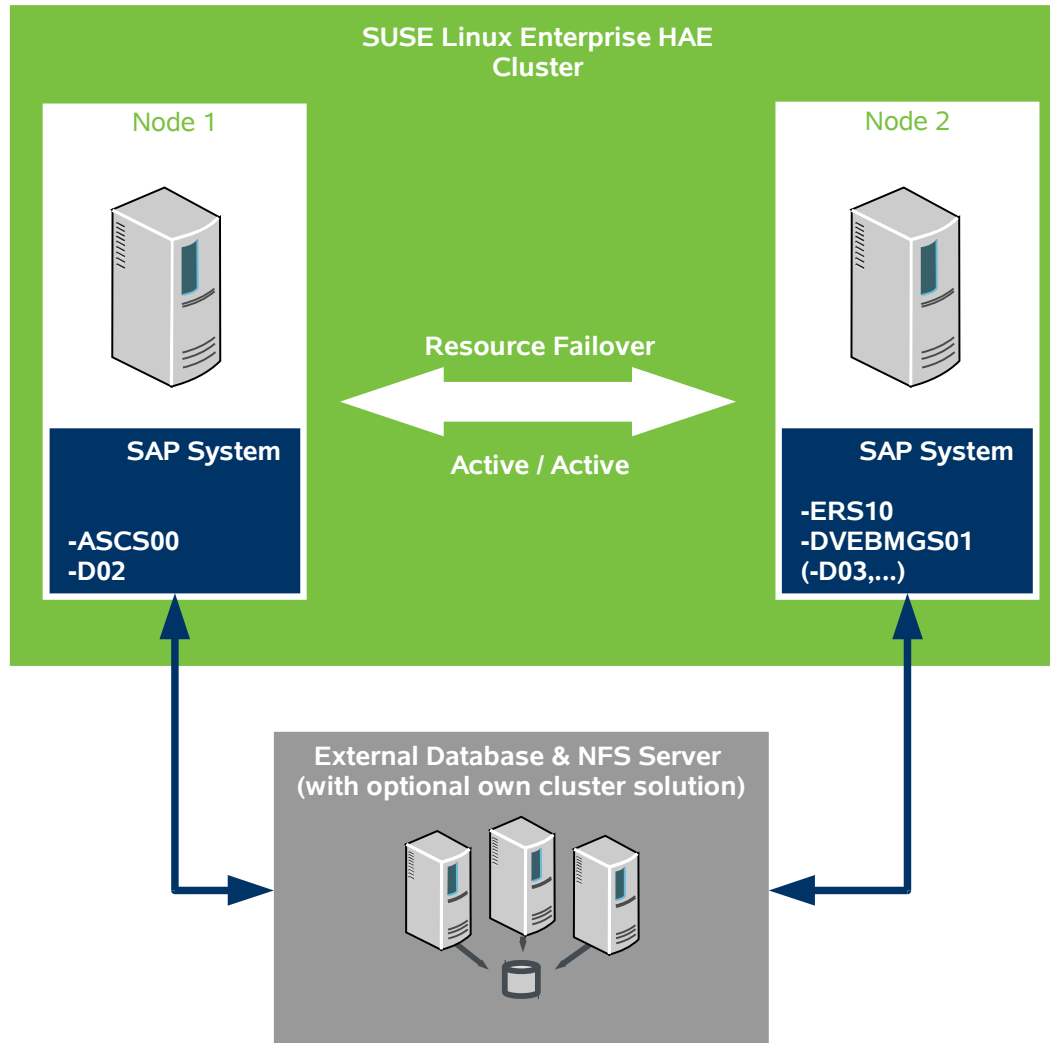


¹Provided by SAP

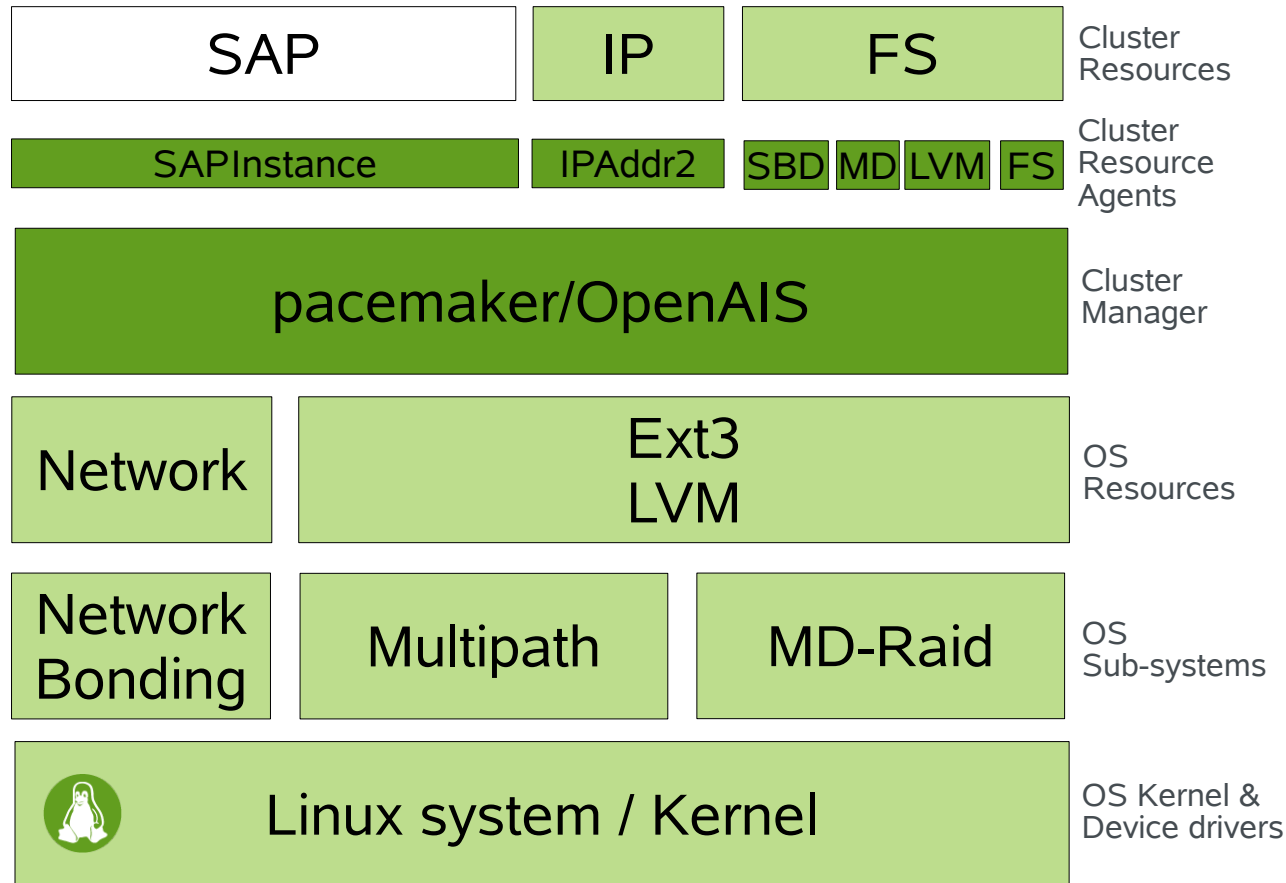
²Legacy HA Solution Interface

³Included in SLES for SAP Applications

Architecture: Overview



HA Stack for SAP



- Cluster manager pacemaker/openAIS from SLE11-HAE
- Resource agents for SAP Instance is part of SLE11-HAE
- Supported databases
 - DB/2
 - MaxDB
 - Sybase ASE
 - Oracle
- Server fencing
 - Redundant SBD

Storage and Network

Filesystems for SAP #1

Enqueue Replication with Integrated Database

- Local filesystems

- /usr/sap

- /usr/sap/HB2

- /usr/sap/HB2/ERS10

- /usr/sap/HB2/ERS11

- Shared storage filesystems

- /usr/sap/HB2/DVEBMSG02

- /sapdb

- /sapdb/HB2/sapdata

- /sapdb/HB2/saplog

- Network-mounted filesystems

- /sapmnt/HB2

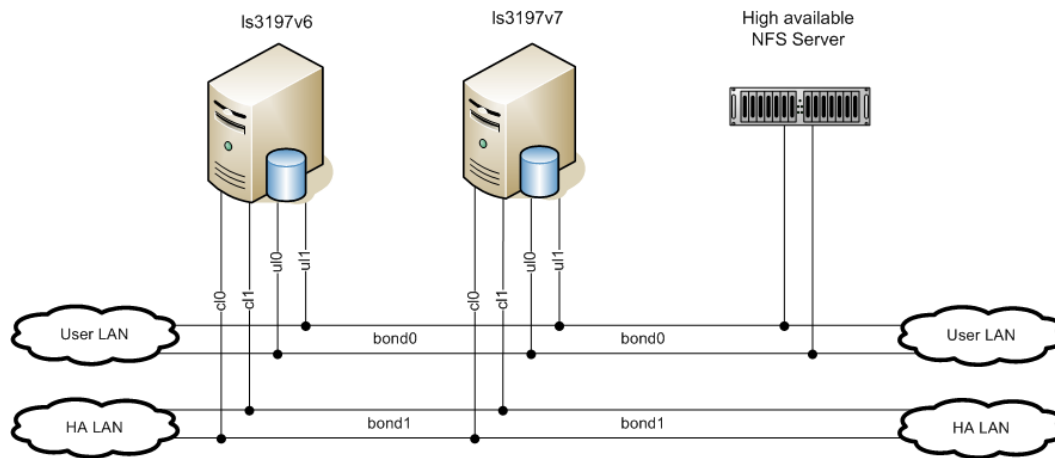
- /usr/sap/HB2/ASC00

- /usr/sap/HB2/SCS01

Note: HB2 is the name of this particular SAP system instance

- Filesystems for ERS instances are local on every node because of the Master-Slave mechanism of the SAPInstance resource agent.
- Filesystems for database tables and transaction logs are on a shared storage for active/passive clustered DB operation.
- Filesystems for ASCS and SCS are on a network filesystem for optional scale-out.

Network



- Separate networks for internal and external communication
- Network bonding to increase network stableness.
- Network infrastructure should be redundant, too.

Cluster Manager Pacemaker/Corosync

Resource Agents for SAP 1/2

- Operations of the SAPInstance resource agent are done by using the startup framework called SAP Management Console or sapstartsrv.
- sapstartsrv knows 4 status colours:
 - GREEN = everything is fine
 - YELLOW = something is wrong, but the service is still working
 - RED = the service does not work
 - GRAY = the service has not been started
- SAPInstance resource agent will interpret GREEN and YELLOW as OK, statuses RED and GRAY are reported as NOT_RUNNING to the cluster.
- Depending on the status the cluster expects from the resource, it will do a restart, failover or just nothing.

Resource Agents for SAP 2/2

- SAPInstance resource agent
 - disp+work
 - msg_server
 - enservice
 - enrepserver
 - jcontrol
 - jstart
 - wdisp
- SAPDatabase resource agent
 - DB/2
 - MaxDB
 - Oracle
 - Sybase

- Supported Releases:
 - SAP WebAS ABAP Release 4.6C - 7.10
 - SAP WebAS Java Release 6.40 - 7.10 (min. 6.40 SP22, 7.00 SP15, 7.10 SP00)
 - SAP WebAS ABAP + Java Add-In Release 6.20 - 7.30
 - DB/2 UDB 9.x, 10.x
 - Oracle 10gR2, 11gR2
 - SAP-DB / MaxDB 7.6, 7.7
 - Sybase ASE 15.7

<http://www.suse.com/products/sles-for-sap/resource-library>

Cluster Resources #1

Enqueue Replication with External Database (SLES11)

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

The cluster resources are listed in a table with three columns: Resource Name, Status, and Inactive Resources. The resources are grouped by their Master/Slave Set: msl_sap_enqrepl_HA0.

Resource Name	Status	Inactive Resources
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	

Cluster Resources #2

Enqueue Replication with Integrated Database (SLES10)

The screenshot displays the Linux HA Management Client interface. The left pane shows a tree of cluster resources. The right pane shows the parameters for the selected resource, `rsc_SAP_HB2_ASCS00_ascshost:0`.

Current Running on [ls3197v6]

Name	Value
InstanceName	HB2_ASCS00_ascshost
AUTOMATIC_RECOVER	true
START_PROFILE	/sapmnt/HB2/profile/START_ASCS00_ascshost
ERS_InstanceName	HB2_ERS10_ers10host
ERS_START_PROFILE	/sapmnt/HB2/profile/START_ERS10_ers10host

Buttons: Add Parameter, Delete Parameter, Apply, Reset

Connected to 127.0.0.1

Two Enqueue Replication Groups

SCS and ASCS with Replicated Enqueue

The screenshot shows the Pacemaker GUI (on cl3n01) with the following configuration:

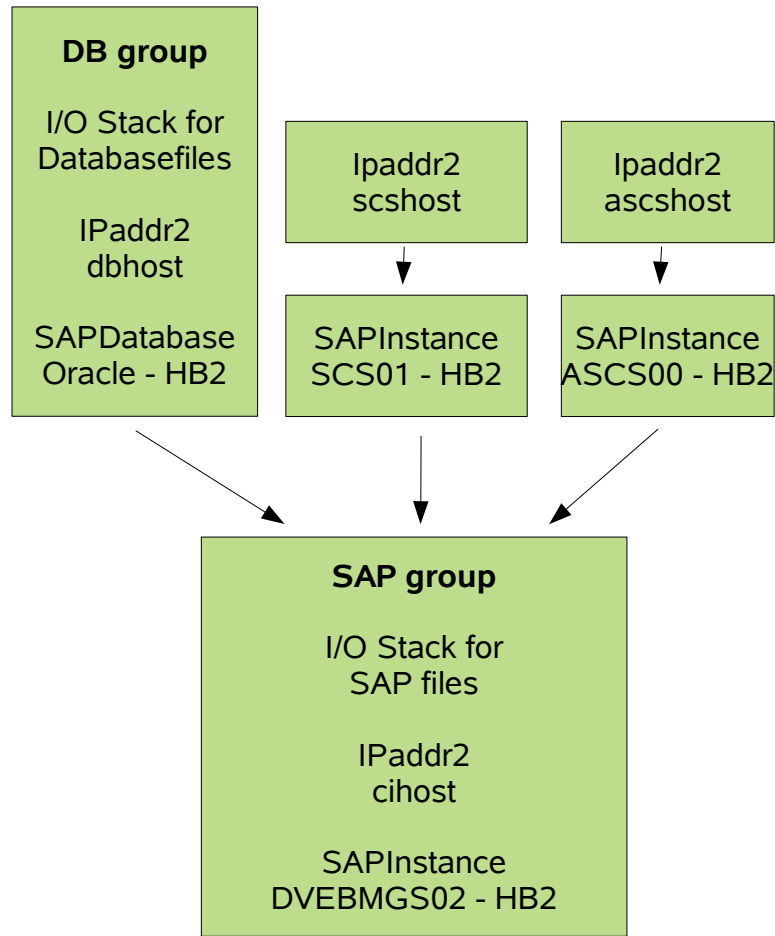
Name	Status	Details
Cluster	have quorum	Openais & Pacemaker
cl3n01	online (dc)	
cl3n02	online	
Resources		
rsc_stonith_sbd	running on [cl3n01]	stonith::external/sbd
grp_sapdb_so1	group	
rsc_ip_SO1_sapso1db	running on [cl3n01]	ocf::heartbeat:IPAddr2
rsc_md_SO1_md0	running on [cl3n01]	ocf::heartbeat:Raid1
rsc_lvm_SO1_sapso1db	running on [cl3n01]	ocf::heartbeat:LVM
rsc_fs_SO1_sapdb	running on [cl3n01]	ocf::heartbeat:Filesystem
rsc_sapdb_SO1	running on [cl3n01]	ocf::heartbeat:SAPDatabase
grp_sapci_so1	group	
rsc_ip_SO1_sapso1ci	running on [cl3n01]	ocf::heartbeat:IPAddr2
rsc_sapinst_SO1_DVEBMGS20_sapso1ci	running on [cl3n01]	ocf::heartbeat:SAPInstance
msl_sapas_SO1	master	
rsc_sapinst_SO1_ASCS00_sapso1as:0	running (Master) on [cl3n02]	ocf::heartbeat:SAPInstance
rsc_sapinst_SO1_ASCS00_sapso1as:1	running (Slave) on [cl3n01]	ocf::heartbeat:SAPInstance
msl_sapcs_SO1	master	
rsc_sapinst_SO1_SCS01_sapso1cs:0	running (Slave) on [cl3n02]	ocf::heartbeat:SAPInstance
rsc_sapinst_SO1_SCS01_sapso1cs:1	running (Master) on [cl3n01]	ocf::heartbeat:SAPInstance
rsc_ip_SO1_sapso1er	running on [cl3n01]	ocf::heartbeat:IPAddr2
rsc_ip_SO1_sapso1e2	running on [cl3n02]	ocf::heartbeat:IPAddr2
rsc_ip_SO1_sapso1cs	running on [cl3n01]	ocf::heartbeat:IPAddr2
rsc_ip_SO1_sapso1as	running on [cl3n02]	ocf::heartbeat:IPAddr2

Connected to hacluster@127.0.0.1 (Simple Mode)

- SLE11-HA Cluster add-on for SLES11
- Resource agent for SAP Instance is part of SLE11-HA
- Enqueue and Replicated Enqueue as Master/Slave resource
- SCS and ASCS with Replicated Enqueue, independent from each other
- Similar to SAP certified setup

Constraints: Orders #2

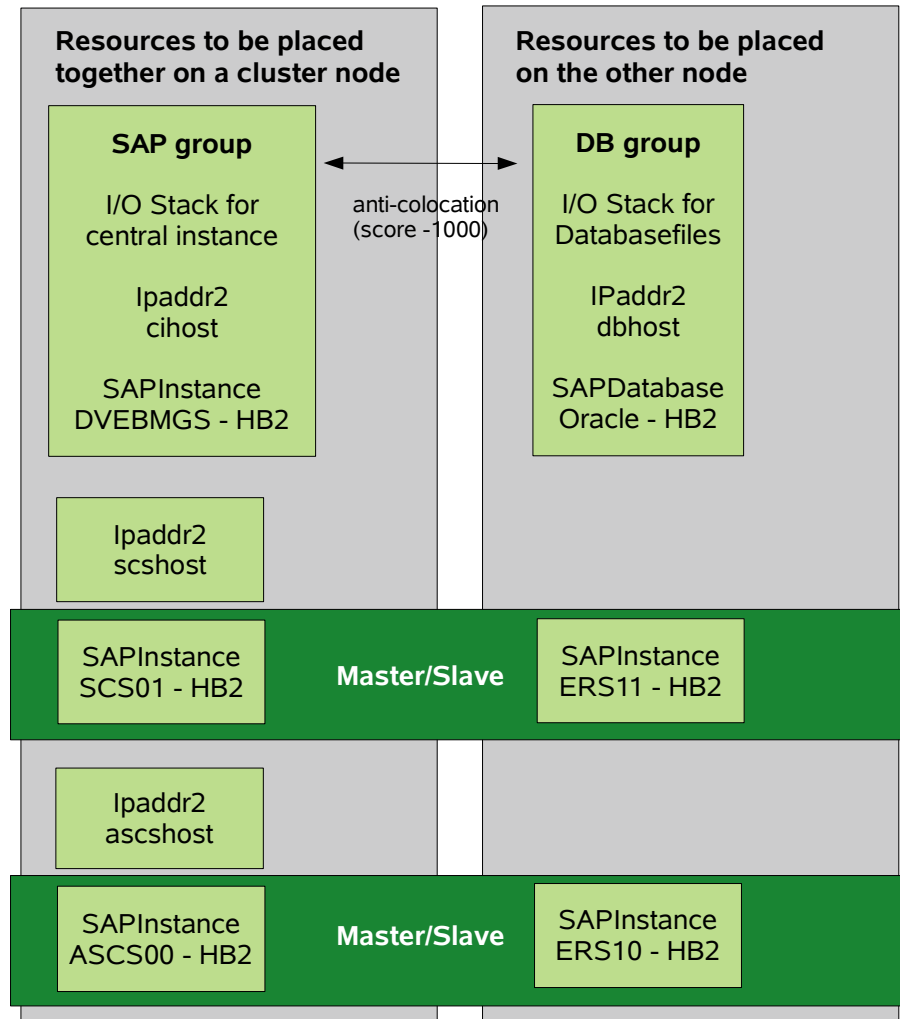
Enqueue Replication with Integrated Database



- The virtual service address is started before the matching SAP instance
- The database and the service instances SCS and ASCS are started before the central instance
- Beside the normal start/stop sequence an order constraint does also define, what to do if a services fails and has to be restarted: All successive services have to be restarted to. This is the reason, why we do not define an order constraint between the database and the service instances SCS and ASCS.

Constraints: Co-location #2

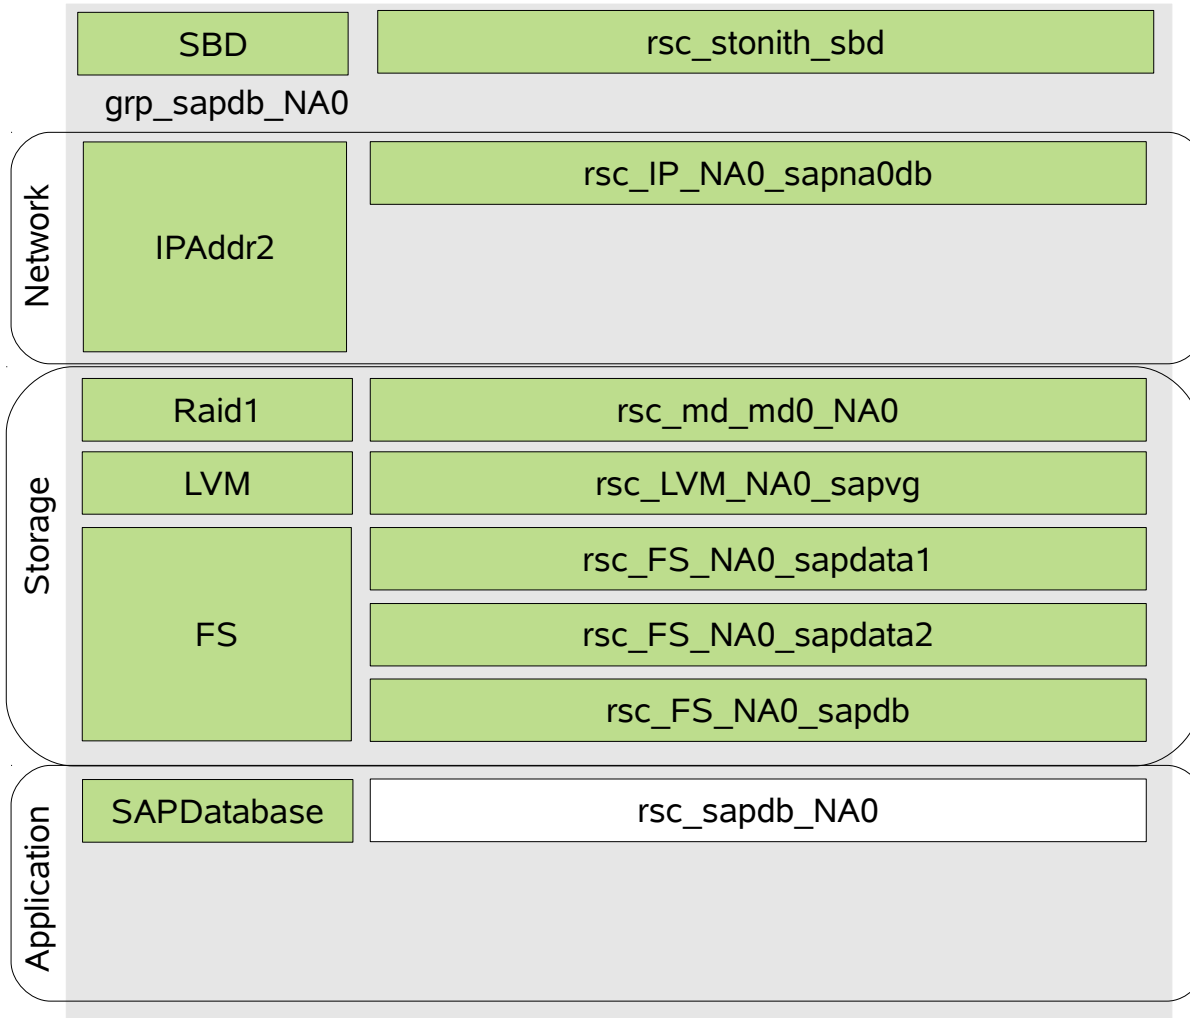
Enqueue Replication with Integrated Database



- Normal scored anti-colocation between the Database and the central instance (DVEBMGS). This should help to balance the load. This anti-colocation is not a must, your operational concepts decide, whether this constraint should be a anti-colocation or a colocation.
- The colocations between the virtual IP addresses and the corresponding SAP instance.

Cluster Resources #3

Simple Stack with SBD for External Database



Cluster Resources #3

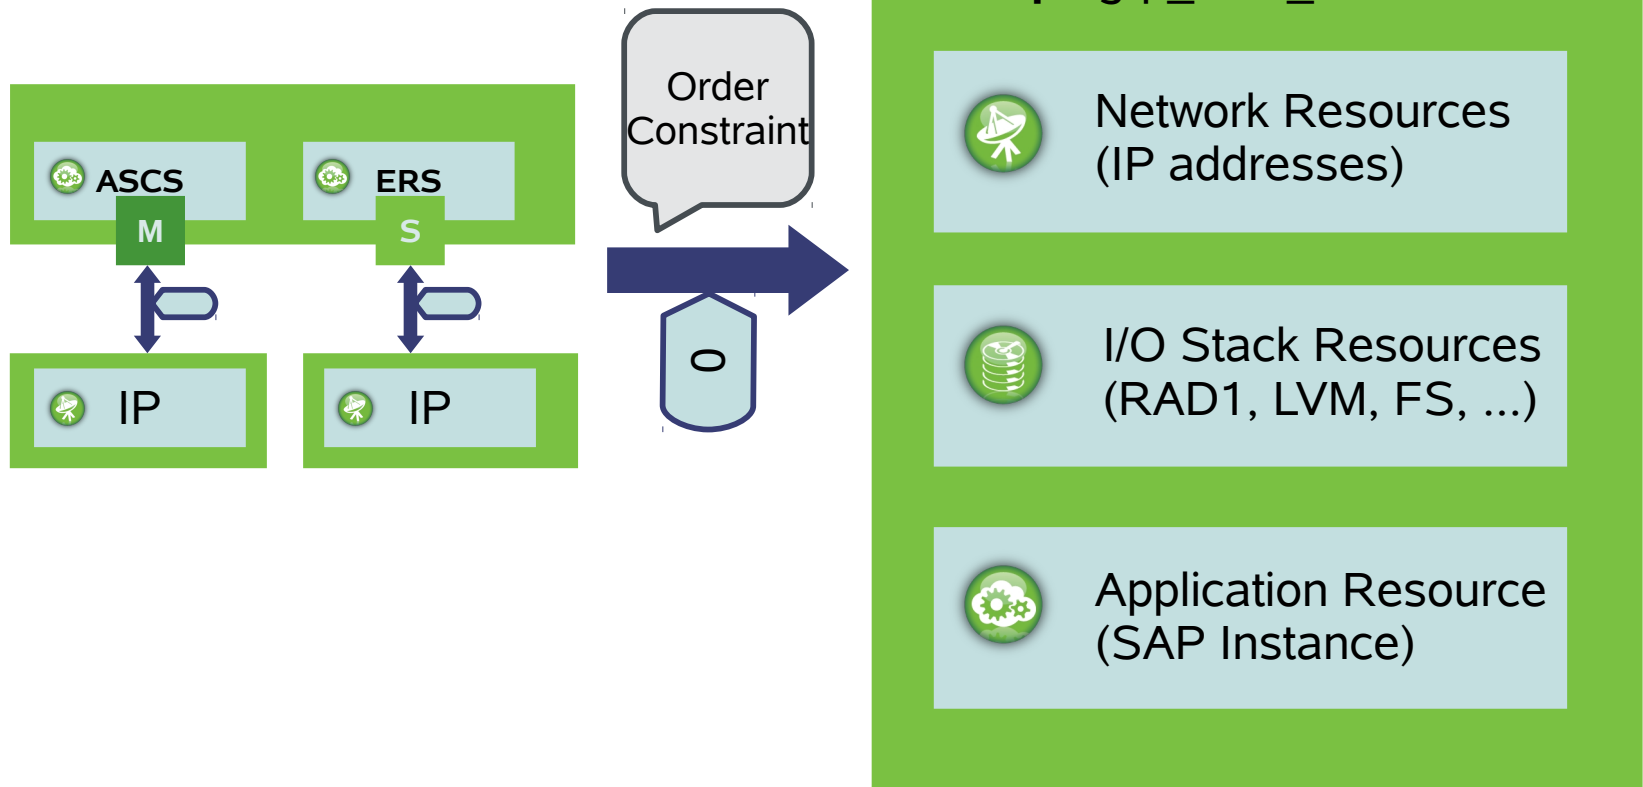
Simple Stack with SBD for External Database

- SBD as primitive resource.
- One common resource group for all database resources.
- Location constraints are needed to ensure that the database is only started if the node fulfills all operational needs.
- Co-location constraints are implicitly defined by adding all resources to one common group.
- Order constraints are implicitly defined by adding all resources to one common group.

- Cluster Resources:
 - SBD: Server fencing
 - IPAddr2: Virtual IP address
 - Raid1: MD Raid1 arrays
 - LVM: LVM Volume Groups
 - Filesystem: Ext3 filesystems
 - SAPDatabase: MaxDB database

Cluster Concept - Constraints

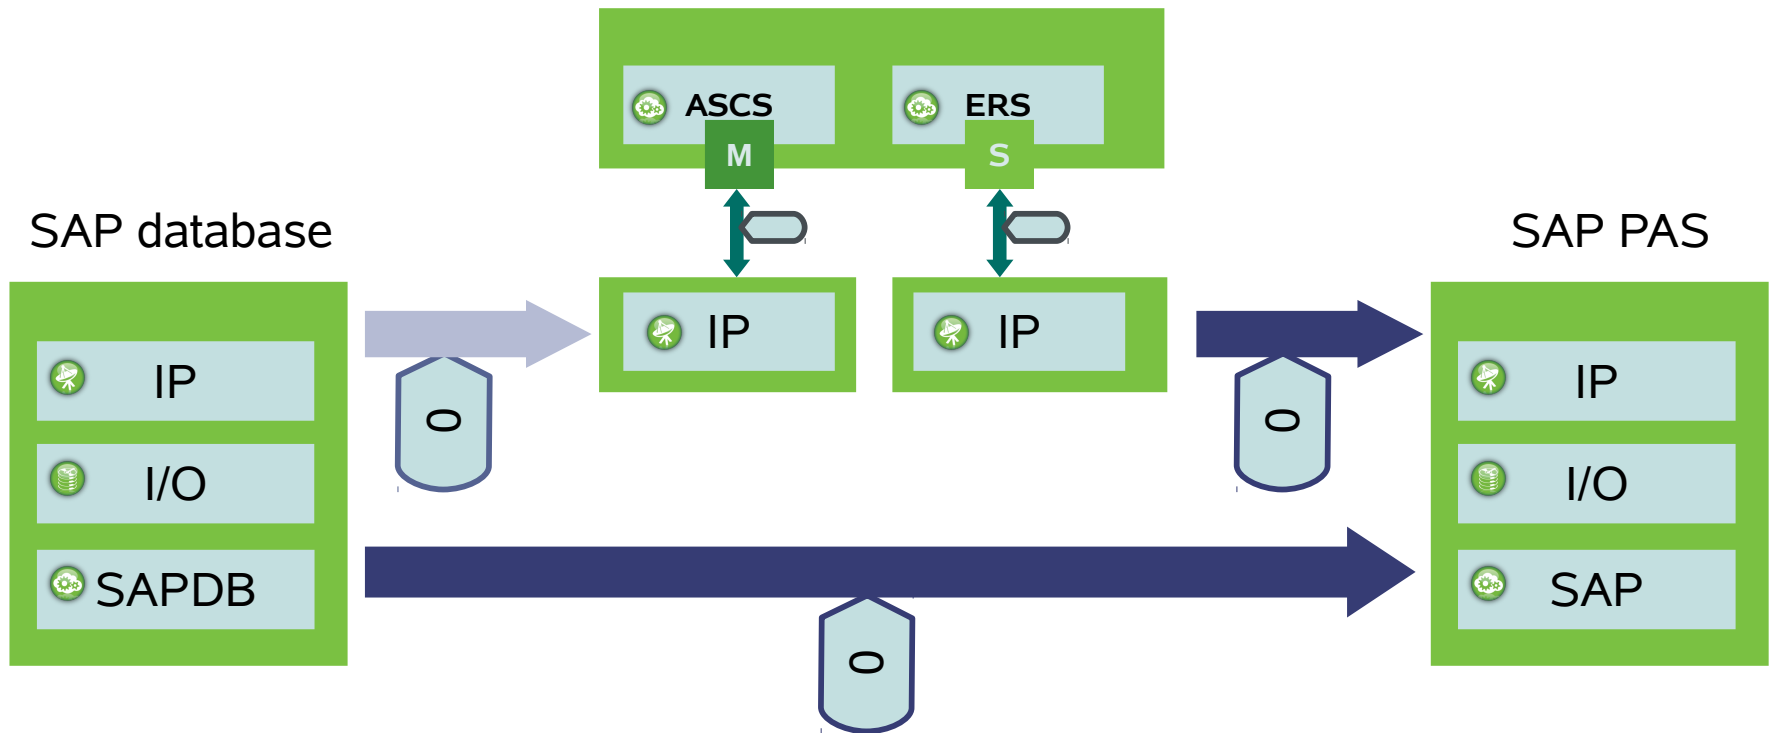
Order Constraints for SAP Instances and Groups



Each SAP application server needs the central services (ASCS)

Cluster Concept - Constraints

Order Constraints for SAP Instances and SAP Database

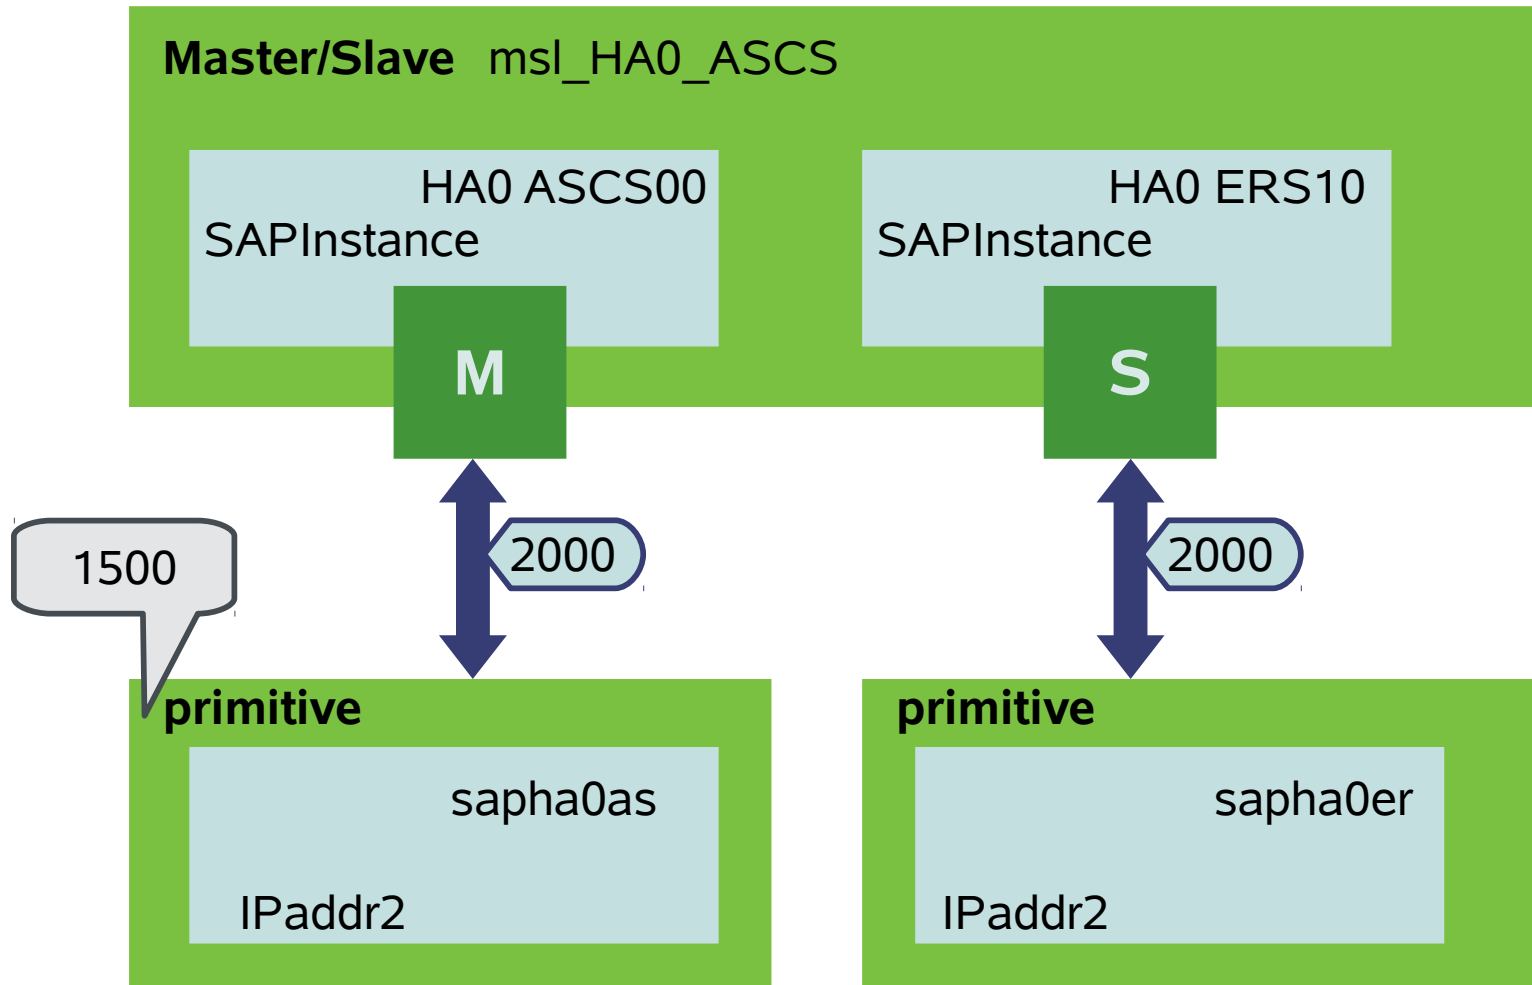


Advisory order constraints for all SAP application servers against Database and central services instance (ASCS)

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

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 Test and Maintenance

Cluster Test #2

Enqueue Replication with Integrated Database

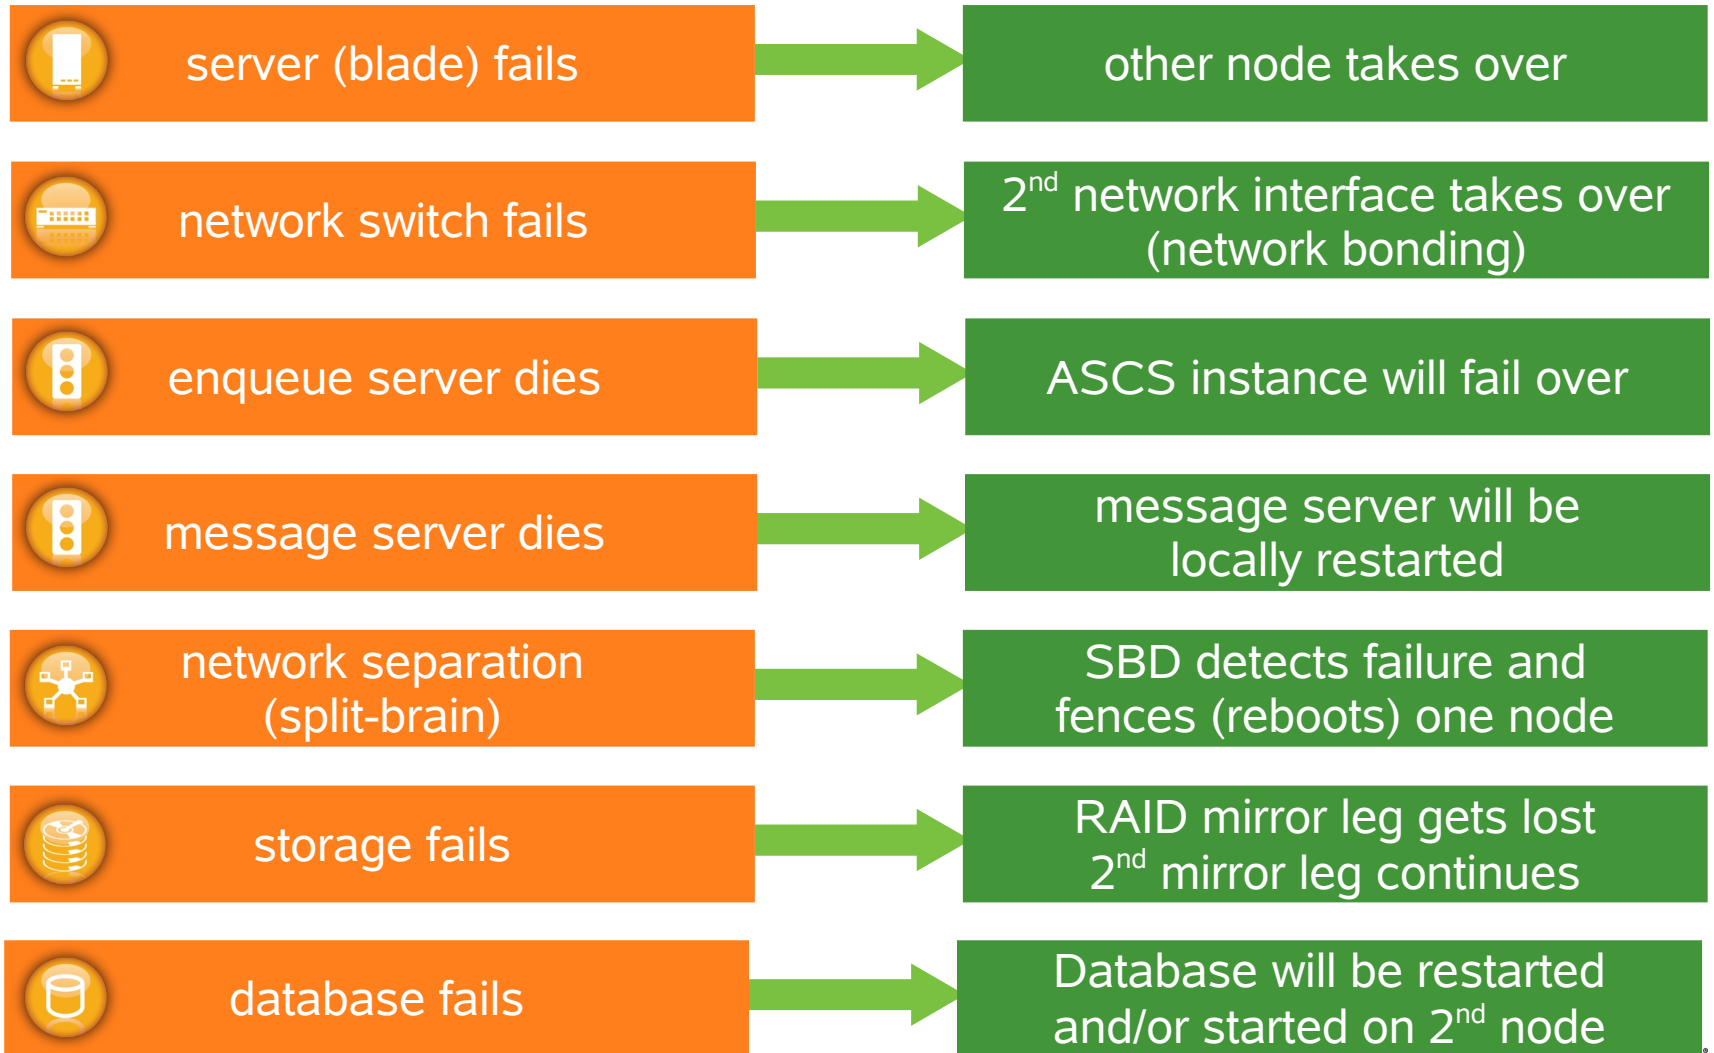


For productive clusters a test plan for all possible failure scenarios is necessary. In the workshop, some important basic tests are used:

Test Case	Expected Result
Set one node to standby	Database and CI instance group is running on remaining node. Virtual IPs are running on remaining node.
Shutdown one node gracefully	The node leaves the cluster. The enqueue replication table is moved. Database and CI instance group start on remaining host.
Turn off one node (power off)	The remaining node will try STONITH. If the STONITH was successful, the remaining node takes over all resources.
Turn on node again, start openais	The node rejoins the cluster and starts resources. Also the CI host moves to the rejoined node.
Plug out User LAN	The second node takes over all resources. Depending on the configuration, the affected node gets STONITHed.
Plug out both SAN links	The monitoring of the SAN resources fail. The affected node gets STONITHed.
Plug out all network (split-brain)	The cluster stops doing anything (resource keep running), since both nodes can't determine the state of their counterpart.
Kill SAP CI instance	The CI instance restarts 3 times. After the 3rd try, it fails over to the second node
Kill database	The database instance restarts 3 times. After the 3rd try, it fails over to the second node
Shut down one SAN storage	The MD-Mirrors get degraded but continue to work.

Failures and Solutions

Some examples of “What happens when”



Cluster Tools

ClusterTools2 RPM is a collection of tools, used in projects by SUSE consulting. Examples:

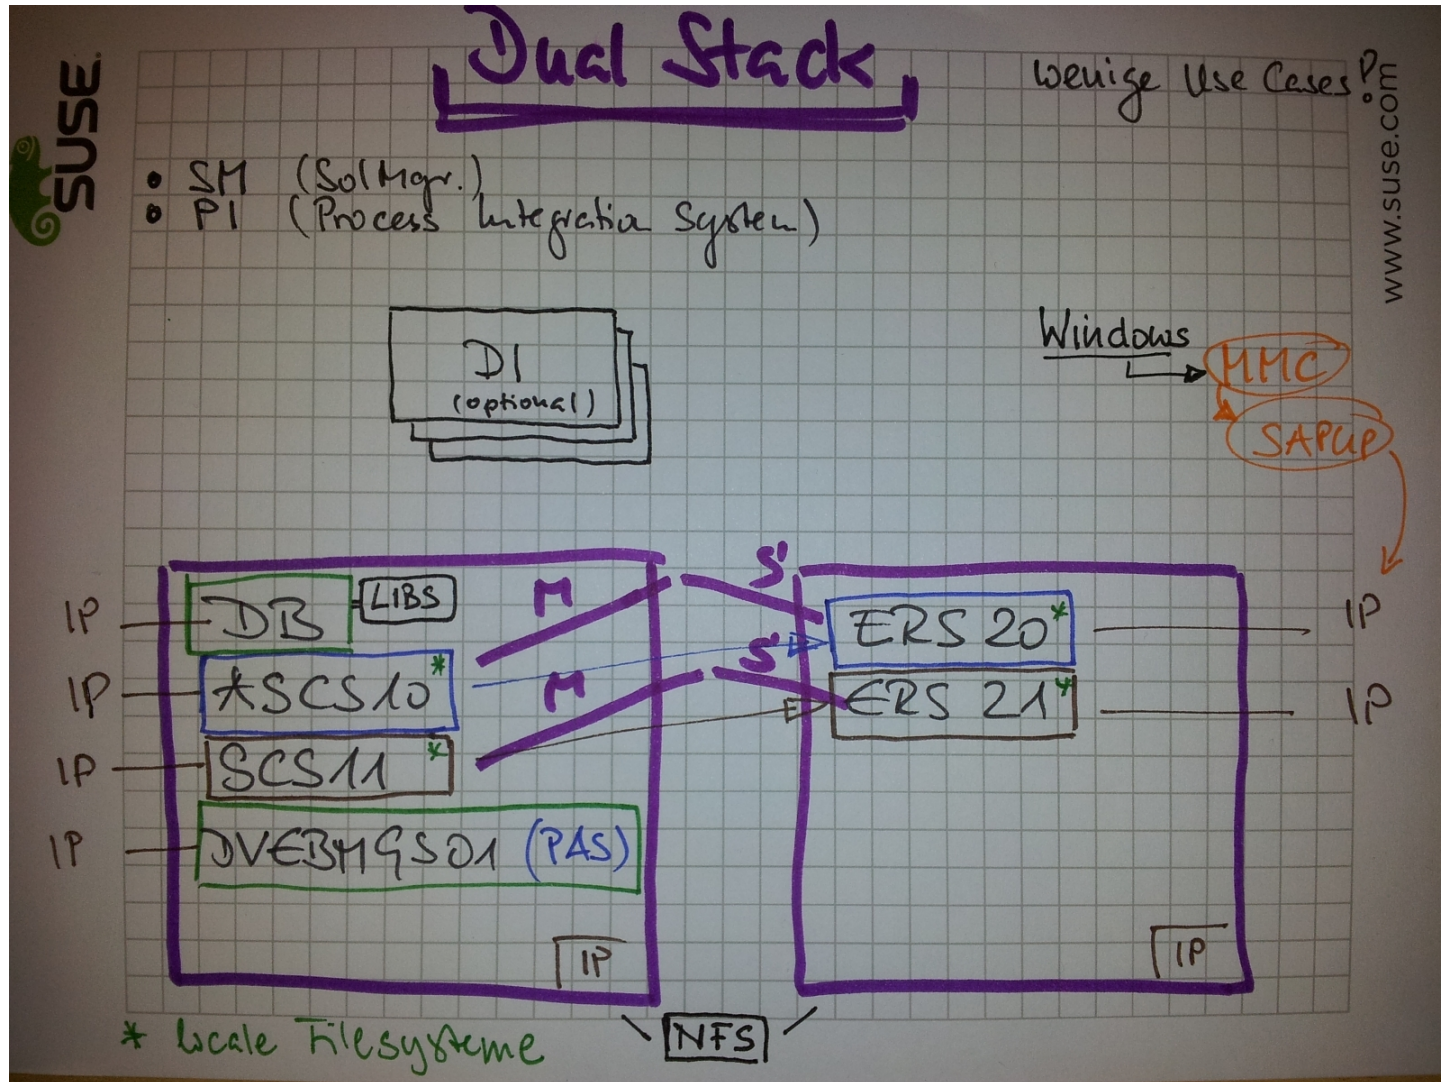
- ClusterService is a frontend for most used cluster operations. It is designed along the use cases that are common for operating a clustered SAP system.
- clusterstate is a script to analyse cluster status.
- Wow is a simple program, which helps to create cluster configurations. It is based on crm-snippets and scripts. The snippets may contain variable references which are resolved by the scripts and the wow engine.

SAP Dual Stack Enqueue Replication

SAP Dual Stack Enqueue Replication

Name	Status	Details
Cluster	have quorum	Openais & Pacemaker
cl3n01	online	
cl3n02	online (dc)	
Resources		
rsc_stonith_sbd	running on [cl3n02]	stonith::external/sbd
grp_sapdb_so1	group	
rsc_ip_SO1_sapso1db	running on [cl3n02]	ocf::heartbeat:IPAddr2
rsc_md_SO1_md0	running on [cl3n02]	ocf::heartbeat:Raid1
rsc_lvm_SO1_sapso1db	running on [cl3n02]	ocf::heartbeat:LVM
rsc_fs_SO1_sapdb	running on [cl3n02]	ocf::heartbeat:Filesystem
rsc_sapdb_SO1	running on [cl3n02]	ocf::heartbeat:SAPDatabase
grp_sapas_so1	group	
rsc_ip_SO1_sapso1as	running on [cl3n01]	ocf::heartbeat:IPAddr2
grp_sapci_so1	group	
rsc_ip_SO1_sapso1ci	running on [cl3n02]	ocf::heartbeat:IPAddr2
rsc_sapinst_SO1_DVEBMGS20_sapso1ci	running on [cl3n02]	ocf::heartbeat:SAPInstance
msl_sapas_SO1	master	
rsc_sapinst_SO1_ASCS00_sapso1as:0	running (Slave) on [cl3n02]	ocf::heartbeat:SAPInstance
rsc_sapinst_SO1_ASCS00_sapso1as:1	running (Master) on [cl3n01]	ocf::heartbeat:SAPInstance
msl_sapcs_SO1	master	
rsc_sapinst_SO1_SCS01_sapso1cs:0	running (Slave) on [cl3n01]	ocf::heartbeat:SAPInstance
rsc_sapinst_SO1_SCS01_sapso1cs:1	running (Master) on [cl3n02]	ocf::heartbeat:SAPInstance
grp_sapcs_so1	group	
rsc_ip_SO1_sapso1cs	running on [cl3n02]	ocf::heartbeat:IPAddr2
grp_saper_so1	group	
rsc_ip_SO1_sapso1er	running on [cl3n02]	ocf::heartbeat:IPAddr2
grp_sapes_so1	group	
rsc_ip_SO1_sapso1es	running on [cl3n01]	ocf::heartbeat:IPAddr2

SAP Dual Stack Enqueue Replication



Appendix

Links 1/2

SAP Notes

0171356	SAP on Linux
0816097	Availability of SAP on Linux for x86_64
1310037	SLES11
1275775	sapconf
1172419	Java VM on x86_64
1008828	SAP Kernel 4.6D
0877795	SAP sapstartsrv before 7.0
0995116	SAP backporting sapstartsrv
1014480	SAP Mgmt. Console
0618104	sapsysinfo
1552925	Linux: High Availability Cluster Solutions
1056161	SUSE Priority Support for SAP-Applications
1763512	Support details for SUSE Linux Enterprise High Availability

Links 2/2

Links

<http://www.sap.com/linux>

<http://wiki.sdn.sap.com/wiki/display/HOME/SAPonLinuxNotes>

<http://www.sap.com/partners/directories/searchpartner.epx>

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

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

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

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

<http://www.sdn.sap.com/irj/scn/index?rid=/library/uuid/d079c1f2-5c6b-2f10-dcbb-a29c20865e81>

<http://www.novell.com/de-de/partners/sap>

<http://www.novell.com/products/server/sap/matrix.html>

<http://www.suse.com/products/sles-for-sap/resource-library>

http://download.opensuse.org/repositories/home:/fmherschel/SLE_11_SP1/noarch/

<click>

Quiz Time

-



Corporate Headquarters
Maxfeldstrasse 5
90409 Nuremberg
Germany

+49 911 740 53 0 (Worldwide)
www.suse.com

Join us on:
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.

