



Geo Clustering Quick Start

SUSE Linux Enterprise High Availability Extension 15

Tanja Roth and Thomas Schraitle

Geo clustering protects workloads across globally distributed data centers. This document guides you through the basic setup of a Geo cluster, using the Geo bootstrap scripts provided by the `ha-cluster-bootstrap` package.

Publication Date: July 11, 2018

Contents

- 1 Conceptual Overview 3
- 2 Usage Scenario 4
- 3 Requirements 5
- 4 Overview of the Geo Bootstrap Scripts 6
- 5 Installation 7
- 6 Setting Up the First Site of a Geo Cluster 7
- 7 Adding Another Site to a Geo Cluster 9
- 8 Adding the Arbitrator 10
- 9 Monitoring the Cluster Sites 11
- 10 Next Steps 13
- 11 For More Information 15

- 12 Legal Notice 15
- A GNU Licenses 15
 - A.1 GNU Free Documentation License 15

1 Conceptual Overview

Geo clusters based on SUSE® Linux Enterprise High Availability Extension can be considered “overlay” clusters where each cluster site corresponds to a cluster node in a traditional cluster. The overlay cluster is managed by the booth cluster ticket manager (in the following called booth). Each of the parties involved in a Geo cluster runs a service, the `boothd`. It connects to the booth daemons running on the other sites and exchanges connectivity details. For making cluster resources highly available across sites, booth relies on cluster objects called tickets. A ticket grants the right to run certain resources on a specific cluster site. Booth guarantees that every ticket is granted to no more than one site at a time.

If the communication between two booth instances breaks down, it might be because of a network breakdown between the cluster sites *or* because of an outage of one cluster site. In this case, you need an additional instance (a third cluster site or an arbitrator) to reach consensus about decisions (such as failover of resources across sites). Arbitrators are single machines (outside of the clusters) that run a booth instance in a special mode. Each Geo cluster can have one or multiple arbitrators.

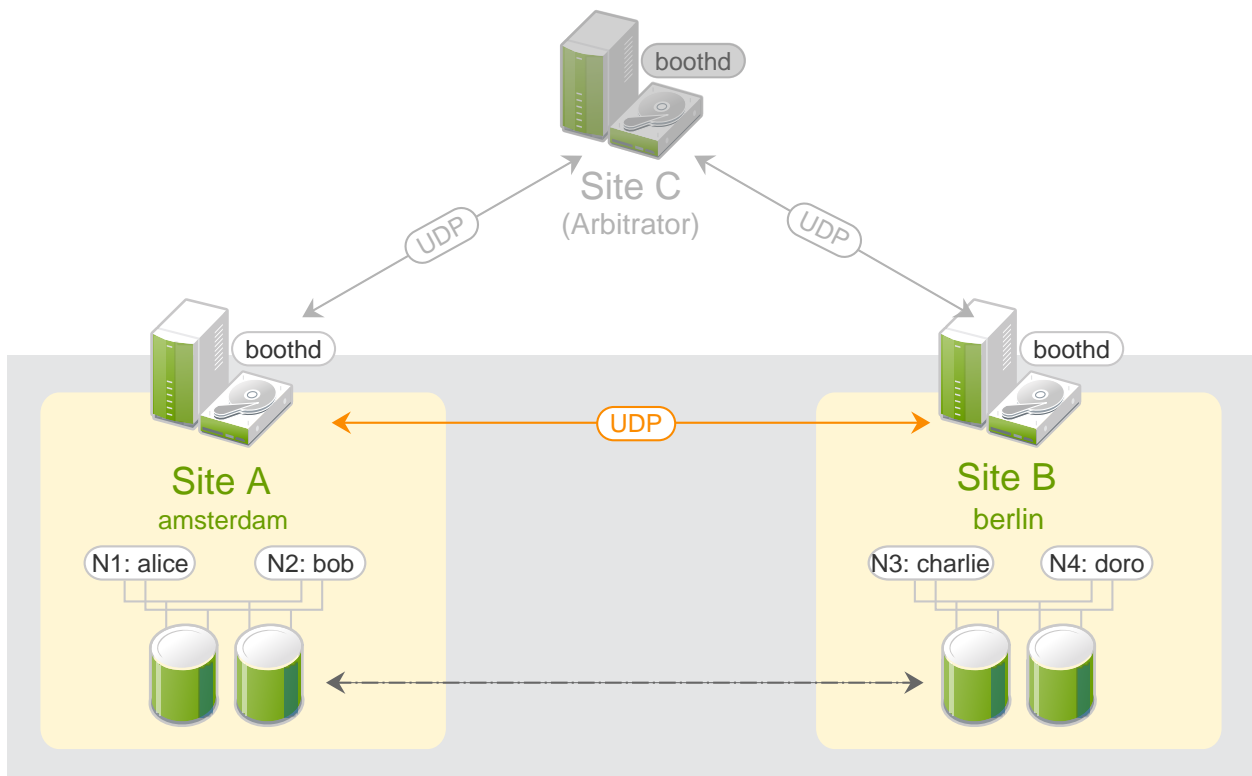


FIGURE 1: TWO-SITE CLUSTER—2X2 NODES + ARBITRATOR (OPTIONAL)

It is also possible to run a two-site Geo cluster *without* an arbitrator. In this case, a Geo cluster administrator needs to manually manage the tickets. If a ticket should be granted to more than one site at the same time, booth displays a warning.

For more details on the concept, components and ticket management used for Geo clusters, see the Book *“Geo Clustering Guide”*.

2 Usage Scenario

In the following, we will set up a basic Geo cluster with two cluster sites and one arbitrator:

- We assume the cluster sites are named amsterdam and berlin.
- We assume that each site consists of two nodes. The nodes alice and bob belong to the cluster amsterdam. The nodes charlie and doro belong to the cluster berlin.
- Site amsterdam will get the following virtual IP address: 192.168.201.100.
- Site berlin will get the following virtual IP address: 192.168.202.100.
- We assume that the arbitrator has the following IP address: 192.168.203.100.

Before you proceed, make sure the following requirements are fulfilled:

REQUIREMENTS

Two Existing Clusters

You have at least two existing clusters that you want to combine into a Geo cluster. (If you need to set up two clusters first, follow the instructions in the *Installation and Setup Quick Start*.)

Meaningful Cluster Names

Each cluster has a meaningful cluster name defined in /etc/corosync/corosync.conf that reflects its location.

Arbitrator

You have installed a third machine that is not part of any existing clusters and is to be used as arbitrator.

For detailed requirements on each item, see also *Section 3, “Requirements”*.

3 Requirements

All machines (cluster nodes and arbitrators) that will be part of the cluster need at least the following modules and extensions:

- Base System Module 15
- Server Applications Module 15
- SUSE Linux Enterprise High Availability Extension 15

When installing the machines, select `HA GEO Node` as `system role`. This leads to the installation of a minimal system where the packages from the pattern `Geo Clustering for High Availability (ha_geo)` are installed by default.

NETWORK REQUIREMENTS

- The virtual IPs to be used for each cluster site must be accessible across the Geo cluster.
- The sites must be reachable on one UDP and TCP port per booth instance. That means any firewalls or IPsec tunnels in between must be configured accordingly.
- Other setup decisions may require to open more ports (for example, for DRBD or database replication).

OTHER REQUIREMENTS AND RECOMMENDATIONS

- All cluster nodes on all sites should synchronize to an NTP server outside the cluster. For more information, see the *Administration Guide* for SUSE Linux Enterprise Server 15, available at <http://www.suse.com/documentation/sles>. Refer to the chapter *Time Synchronization with NTP*.

If nodes are not synchronized, log files and cluster reports are very hard to analyze.

- Use an *uneven* number of sites in your Geo cluster. In case the network connection breaks down, this makes sure that there still is a majority of sites (to avoid a split brain scenario). In case you have an even number of cluster sites, use an arbitrator for handling automatic failover of tickets. If you do not use an arbitrator, you need to handle ticket failover manually.
- The cluster on each site has a meaningful name, for example: `amsterdam` and `berlin`. The cluster names for each site are defined in the respective `/etc/corosync/corosync.conf` files:

```
totem {
```

```
[...]  
cluster_name: amsterdam  
}
```

Change the name with following crmsh command:

```
root # crm cluster rename NEW_NAME
```

Stop and start the `pacemaker` service for the changes to take effect:

```
root # systemctl stop pacemaker  
root # systemctl start pacemaker
```

4 Overview of the Geo Bootstrap Scripts

- With `ha-cluster-geo-init`, turn a cluster into the first site of a Geo cluster. The script takes some parameters like the names of the clusters, the arbitrator, and one or multiple tickets and creates `/etc/booth/booth.conf` from them. It copies the booth configuration to all nodes on the current cluster site. It also configures the cluster resources needed for booth on the current cluster site.

For details, see *Section 6, "Setting Up the First Site of a Geo Cluster"*.

- With `ha-cluster-geo-join`, add the current cluster to an existing Geo cluster. The script copies the booth configuration from an existing cluster site and writes it to `/etc/booth/booth.conf` on all nodes on the current cluster site. It also configures the cluster resources needed for booth on the current cluster site.

For details, see *Section 7, "Adding Another Site to a Geo Cluster"*.

- With `ha-cluster-geo-init-arbitrator`, turn the current machine into an arbitrator for the Geo cluster. The script copies the booth configuration from an existing cluster site and writes it to `/etc/booth/booth.conf`.

For details, see *Section 8, "Adding the Arbitrator"*.

All bootstrap scripts log to `/var/log/ha-cluster-bootstrap.log`. Check the log file for any details of the bootstrap process. Any options set during the bootstrap process can be modified later (by modifying booth settings, modifying resources etc.). For details, see the [Geo Clustering Guide](#). It is available at <https://www.suse.com/documentation>.

5 Installation

Support for using High Availability clusters across unlimited distances is available with Geo Clustering for SUSE Linux Enterprise High Availability Extension.

For setting up a Geo cluster you need the packages included in the following installation patterns:

- High Availability (named sles_ha on the command line)
- Geo Clustering for High Availability (named ha_geo on the command line)

Both patterns are only available if you have registered your system at SUSE Customer Center (or a local registration server) and have added the respective modules or installation media as an extension. For information on how to install extensions, see the *SUSE Linux Enterprise 15 Deployment Guide*, available at <http://www.suse.com/documentation/sles>. Refer to chapter *Installing Modules, Extensions, and Third Party Add-On Products*.

To install the packages from both patterns via command line, use Zypper:

```
root # zypper install -t pattern ha_sles ha_geo
```

Important: Installing Software Packages on all Parties

The software packages needed for High Availability and Geo clusters are *not* automatically copied to the cluster nodes.

- Install SUSE Linux Enterprise Server 15 and the ha_sles and ha_geo patterns on *all* machines that will be part of your Geo cluster.
- Instead of manually installing the packages on all machines that will be part of your cluster, use AutoYaST to clone existing nodes. Find more information at *Book "Administration Guide", Chapter 3 "Installing the High Availability Extension", Section 3.2 "Mass Installation and Deployment with AutoYaST"*.

6 Setting Up the First Site of a Geo Cluster

Use the ha-cluster-geo-init script to turn an existing cluster into the first site of a Geo cluster.

PROCEDURE 1: SETTING UP THE FIRST SITE (amsterdam) WITH `ha-cluster-geo-init`

1. Define a virtual IP per cluster site that can be used to access the site. We assume using `192.168.201.100` and `192.168.202.100` for this purpose. You do not need to configure the virtual IPs as cluster resources yet. This will be done by the bootstrap scripts.
2. Define the name of at least one ticket that will grant the right to run certain resources on a cluster site. Use a meaningful name that reflects the resources that will depend on the ticket (for example, `ticket-nfs`). The bootstrap scripts only need the ticket name—you can define the remaining details (ticket dependencies of the resources) later on, as described in *Section 10, "Next Steps"*.
3. Log in to a node of an existing cluster (for example, on node `alice` of the cluster `amsterdam`).
4. Run `ha-cluster-geo-init`. For example, use the following options:

```
root # ha-cluster-geo-init \  
--clusters ❶ "amsterdam=192.168.201.100 berlin=192.168.202.100" \  
--tickets ❷ ticket-nfs \  
--arbitrator ❸ 192.168.203.100
```

- ❶ The names of the cluster sites (as defined in `/etc/corosync/corosync.conf`) and the virtual IP addresses you want to use for each cluster site. In this case, we have two cluster sites (`amsterdam` and `berlin`) with a virtual IP address each.
- ❷ The name of one or multiple tickets.
- ❸ The host name or IP address of a machine outside of the clusters.

The bootstrap script creates the booth configuration file and synchronizes it across the cluster sites. It also creates the basic cluster resources needed for booth. *Step 4* of *Procedure 1* would result in the following booth configuration and cluster resources:

EXAMPLE 1: BOOTH CONFIGURATION CREATED BY `ha-cluster-geo-init`

```
# The booth configuration file is "/etc/booth/booth.conf". You need to  
# prepare the same booth configuration file on each arbitrator and  
# each node in the cluster sites where the booth daemon can be launched.  
  
# "transport" means which transport layer booth daemon will use.  
# Currently only "UDP" is supported.  
transport="UDP"  
port="9929"
```



```
arbitrator="192.168.203.100"
site="192.168.201.100"
site="192.168.202.100"
authfile="/etc/booth/authkey"
ticket="ticket-nfs"
expire="600"
```

EXAMPLE 2: CLUSTER RESOURCES CREATED BY `ha-cluster-geo-init`

```
primitive ❶ booth-ip IPAddr2 \  
  params rule #cluster-name eq amsterdam ip=192.168.201.100 \  
  params rule #cluster-name eq berlin ip=192.168.202.100 \  
primitive ❷ booth-site ocf:pacemaker:booth-site \  
  meta resource-stickiness=INFINITY \  
  params config=booth \  
  op monitor interval=10s  
group ❸ g-booth booth-ip booth-site \  
meta target-role=Stopped ❹
```

- ❶ A virtual IP address for each cluster site. It is required by the booth daemons who need a persistent IP address on each cluster site.
- ❷ A primitive resource for the booth daemon. It communicates with the booth daemons on the other cluster sites. The daemon can be started on any node of the site. To make the resource stay on the same node, if possible, `resource-stickiness` is set to `INFINITY`.
- ❸ A cluster resource group for both primitives. With this configuration, each booth daemon will be available at its individual IP address, independent of the node the daemon is running on.
- ❹ The cluster resource group is not started by default. After verifying the configuration of your cluster resources (and adding the resources you need to complete your setup), you need to start the resource group. See *Required Steps to Complete the Geo Cluster Setup* for details.

7 Adding Another Site to a Geo Cluster

After you have initialized the first site of your Geo cluster, add the second cluster with `ha-cluster-geo-join`, as described in *Procedure 2*. The script needs SSH access to an already configured cluster site and will add the current cluster to the Geo cluster.

PROCEDURE 2: ADDING THE SECOND SITE (berlin) WITH `ha-cluster-geo-join`

1. Log in to a node of the cluster site you want to add (for example, on node `charlie` of the cluster `berlin`).

2. Run the `ha-cluster-geo-join` command. For example:

```
root # ha-cluster-geo-join \  
--cluster-node ① 192.168.201.100\  
--clusters ② "amsterdam=192.168.201.100 berlin=192.168.202.100"
```

- ① Specifies where to copy the booth configuration from. Use the IP address or host name of a node in an already configured Geo cluster site. You can also use the virtual IP address of an already existing cluster site (like in this example). Alternatively, use the IP address or host name of an already configured arbitrator for your Geo cluster.
- ② The names of the cluster sites (as defined in `/etc/corosync/corosync.conf`) and the virtual IP addresses you want to use for each cluster site. In this case, we have two cluster sites (`amsterdam` and `berlin`) with a virtual IP address each.

The `ha-cluster-geo-join` script copies the booth configuration from ①, see [Example 1](#). In addition, it creates the cluster resources needed for booth (see [Example 2](#)).

8 Adding the Arbitrator

After you have set up all sites of your Geo cluster with `ha-cluster-geo-init` and `ha-cluster-geo-join`, set up the arbitrator with `ha-cluster-geo-init-arbitrator`.

PROCEDURE 3: SETTING UP THE ARBITRATOR WITH `ha-cluster-geo-init-arbitrator`

1. Log in to the machine you want to use as arbitrator.
2. Run the following command. For example:

```
root # ha-cluster-geo-init-arbitrator --cluster-node ① 192.168.201.100
```

- ① Specifies where to copy the booth configuration from. Use the IP address or host name of a node in an already configured Geo cluster site. Alternatively, use the virtual IP address of an already existing cluster site (like in this example).

The `ha-cluster-geo-init-arbitrator` script copies the booth configuration from ①, see [Example 1](#). It also enables and starts the booth service on the arbitrator. Thus, the arbitrator is ready to communicate with the booth instances on the cluster sites when the booth services are running there too.

9 Monitoring the Cluster Sites

To view both cluster sites with the resources and the ticket that you have created during the bootstrapping process, use Hawk2. The Hawk2 Web interface allows you to monitor and manage multiple (unrelated) clusters and Geo clusters.

PREREQUISITES

- All clusters to be monitored from Hawk2's *Dashboard* must be running SUSE Linux Enterprise High Availability Extension 15.
- If you did not replace the self-signed certificate for Hawk2 on every cluster node with your own certificate (or a certificate signed by an official Certificate Authority) yet, do the following: Log in to Hawk2 on *every* node in *every* cluster at least once. Verify the certificate (or add an exception in the browser to bypass the warning). Otherwise Hawk2 cannot connect to the cluster.

PROCEDURE 4: USING THE HAWK2 DASHBOARD

1. Start a Web browser and enter the virtual IP of your first cluster site, amsterdam:

```
https://192.168.201.100:7630/
```

Alternatively, use the IP address or host name of alice or bob. If you have set up both nodes with the bootstrap scripts, the hawk service should run on both nodes.

2. Log in to the Hawk2 Web interface.
3. From the left navigation bar, select *Dashboard*.

Hawk2 shows an overview of the resources and nodes on the current cluster site. In addition, it shows any *Tickets* that have been configured for the Geo cluster. If you need information about the icons used in this view, click *Legend*.

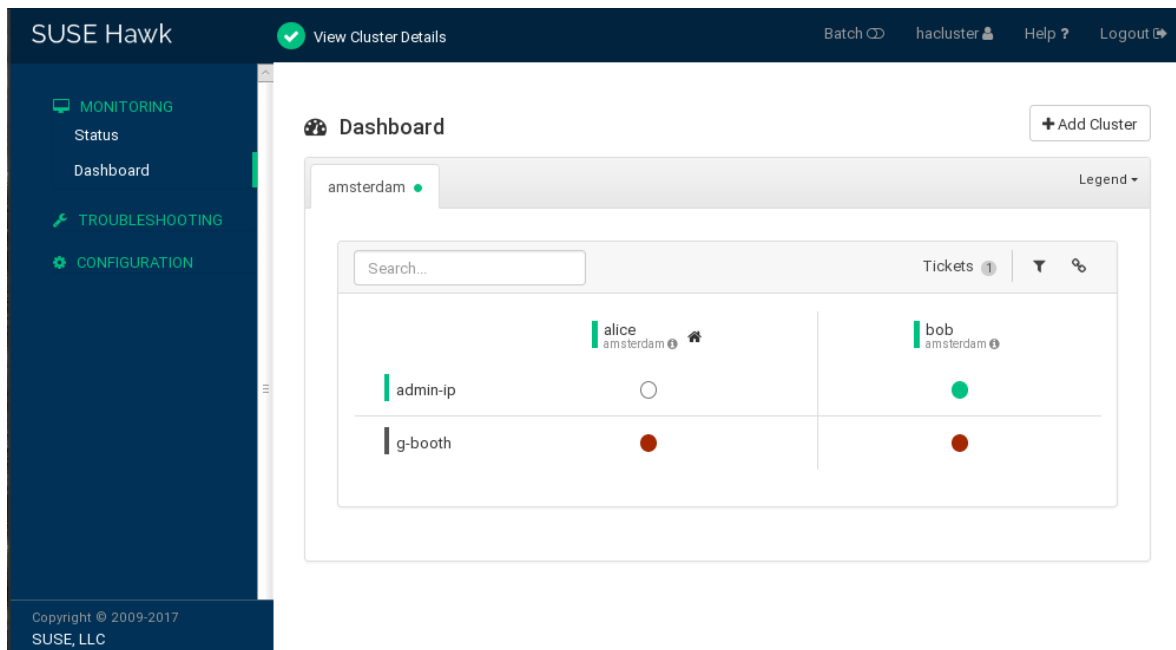


FIGURE 2: HAWK2 DASHBOARD WITH ONE CLUSTER SITE (amsterdam)

4. To add a dashboard for the second cluster site, click *Add Cluster*.
 - a. Enter the *Cluster name* with which to identify the cluster in the *Dashboard*. In this case, berlin.
 - b. Enter the fully qualified host name of one of the cluster nodes (in this case, charlie or doro).
 - c. Click *Add*. Hawk2 will display a second tab for the newly added cluster site with an overview of its nodes and resources.

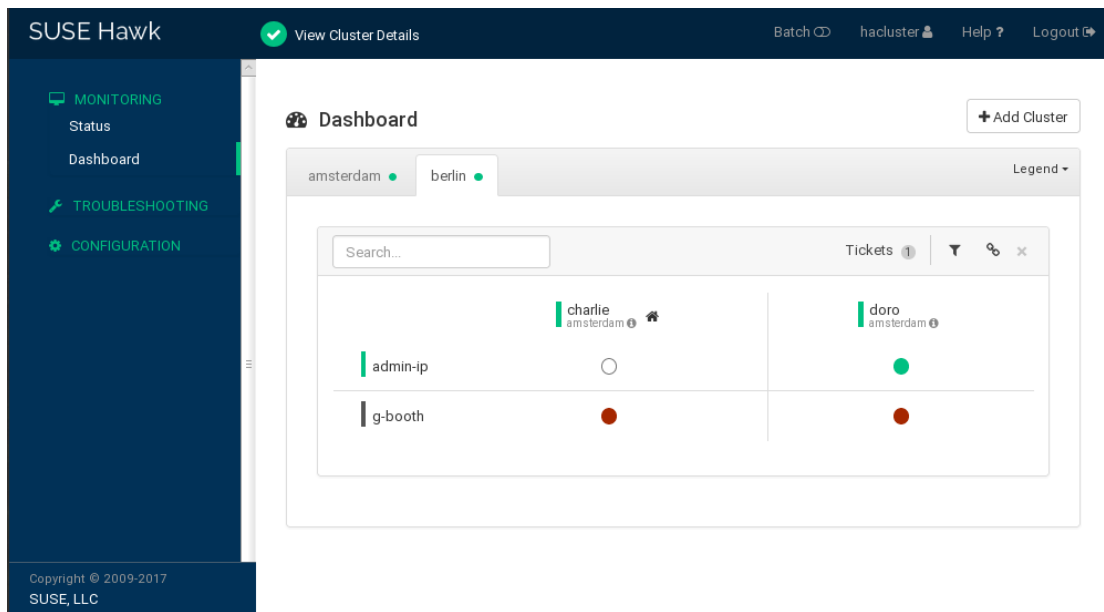


FIGURE 3: HAWK2 DASHBOARD WITH BOTH CLUSTER SITES

- To view more details for a cluster site or to manage it, switch to the site's tab and click the chain icon.

Hawk2 opens the *Status* view for this site in a new browser window or tab. From there, you can administer this part of the Geo cluster.

10 Next Steps

The Geo clustering bootstrap scripts provide a quick way to set up a basic Geo cluster that can be used for testing purposes. However, to move the resulting Geo cluster into a functioning Geo cluster that can be used in production environments, more steps are required—see *Required Steps to Complete the Geo Cluster Setup*.

REQUIRED STEPS TO COMPLETE THE GEO CLUSTER SETUP

Starting the Booth Services on Cluster Sites

After the bootstrap process, the arbitrator booth service cannot communicate with the booth services on the cluster sites yet, because they are not started by default.

The booth service for each cluster site is managed by the booth resource group `g-booth` (see *Example 2, "Cluster Resources Created by ha-cluster-geo-init"*). To start one instance of the booth service per site, start the respective booth resource group on each cluster site. This enables all booth instances to communicate with each other.

Configuring Ticket Dependencies and Ordering Constraints

To make resources depend on the ticket that you have created during the Geo cluster bootstrap process, configure constraints. For each constraint, set a `loss-policy` that defines what should happen to the respective resources if the ticket is revoked from a cluster site. For details, see *Book "Geo Clustering Guide", Chapter 6 "Configuring Cluster Resources and Constraints"*.

Initially Granting a Ticket to a Site

Before booth can manage a certain ticket within the Geo cluster, you initially need to *grant* it to a site manually. You can use either the booth client command line tool or Hawk2 to grant a ticket.

For details, see *Book "Geo Clustering Guide", Chapter 8 "Managing Geo Clusters"*.

The bootstrap scripts create the same booth resources on both cluster sites, and the same booth configuration files on all sites, including the arbitrator. If you extend the Geo cluster setup (to move to a production environment), you will probably fine-tune the booth configuration and change the configuration of the booth-related cluster resources. Afterward, you need to synchronize the changes to the other sites of your Geo cluster to take effect.



Note: Synchronizing Changes Across Cluster Sites

- To synchronize changes in the booth configuration to all cluster sites (including the arbitrator), use Csync2. Find more information at *Book "Geo Clustering Guide", Chapter 5 "Synchronizing Configuration Files Across All Sites and Arbitrators"*.
- The CIB (Cluster Information Database) is not automatically synchronized across cluster sites of a Geo cluster. That means any changes in resource configuration that are required on all cluster sites need to be transferred to the other sites manually. Do so by tagging the respective resources, exporting them from the current CIB, and importing them to the CIB on the other cluster sites. For details, see *Book "Geo Clustering Guide", Chapter 6 "Configuring Cluster Resources and Constraints", Section 6.4 "Transferring the Resource Configuration to Other Cluster Sites"*.

11 For More Information

- More documentation for this product is available at <http://www.suse.com/documentation/sle-ha-geo>. The documentation also includes a comprehensive Geo Clustering Guide. Refer to it for further configuration and administration tasks.
- A document with detailed information how to on data replication via DRBD across Geo clusters has been published in the SUSE Best Practices series: <http://www.suse.com/documentation/suse-best-practices/sbp-drbd/data/sbp-drbd.html>.

12 Legal Notice

Copyright© 2006– 2018 SUSE LLC and contributors. All rights reserved.

Permission is granted to copy, distribute and/or modify this document under the terms of the GNU Free Documentation License, Version 1.2 or (at your option) version 1.3; with the Invariant Section being this copyright notice and license. A copy of the license version 1.2 is included in the section entitled “GNU Free Documentation License”.

For SUSE trademarks, see <http://www.suse.com/company/legal/>. All other third-party trademarks are the property of their respective owners. Trademark symbols (®, ™ etc.) denote trademarks of SUSE and its affiliates. Asterisks (*) denote third-party trademarks.

All information found in this book has been compiled with utmost attention to detail. However, this does not guarantee complete accuracy. Neither SUSE LLC, its affiliates, the authors, nor the translators shall be held liable for possible errors or the consequences thereof.

A GNU Licenses

This appendix contains the GNU Free Documentation License version 1.2.

GNU Free Documentation License

Copyright (C) 2000, 2001, 2002 Free Software Foundation, Inc. 51 Franklin St, Fifth Floor, Boston, MA 02110-1301 USA. Everyone is permitted to copy and distribute verbatim copies of this license document, but changing it is not allowed.

0. PREAMBLE

The purpose of this License is to make a manual, textbook, or other functional and useful document “free” in the sense of freedom: to assure everyone the effective freedom to copy and redistribute it, with or without modifying it, either commercially or noncommercially. Secondly, this License preserves for the author and publisher a way to get credit for their work, while not being considered responsible for modifications made by others.

This License is a kind of “copyleft”, which means that derivative works of the document must themselves be free in the same sense. It complements the GNU General Public License, which is a copyleft license designed for free software.

We have designed this License in order to use it for manuals for free software, because free software needs free documentation: a free program should come with manuals providing the same freedoms that the software does. But this License is not limited to software manuals; it can be used for any textual work, regardless of subject matter or whether it is published as a printed book. We recommend this License principally for works whose purpose is instruction or reference.

1. APPLICABILITY AND DEFINITIONS

This License applies to any manual or other work, in any medium, that contains a notice placed by the copyright holder saying it can be distributed under the terms of this License. Such a notice grants a world-wide, royalty-free license, unlimited in duration, to use that work under the conditions stated herein. The "Document", below, refers to any such manual or work. Any member of the public is a licensee, and is addressed as "you". You accept the license if you copy, modify or distribute the work in a way requiring permission under copyright law.

A "Modified Version" of the Document means any work containing the Document or a portion of it, either copied verbatim, or with modifications and/or translated into another language.

A "Secondary Section" is a named appendix or a front-matter section of the Document that deals exclusively with the relationship of the publishers or authors of the Document to the Document's overall subject (or to related matters) and contains nothing that could fall directly within that overall subject. (Thus, if the Document is in part a textbook of mathematics, a Secondary Section may not explain any mathematics.) The relationship could be a matter of historical connection with the subject or with related matters, or of legal, commercial, philosophical, ethical or political position regarding them.

The "Invariant Sections" are certain Secondary Sections whose titles are designated, as being those of Invariant Sections, in the notice that says that the Document is released under this License. If a section does not fit the above definition of Secondary then it is not allowed to be designated as Invariant. The Document may contain zero Invariant Sections. If the Document does not identify any Invariant Sections then there are none.

The "Cover Texts" are certain short passages of text that are listed, as Front-Cover Texts or Back-Cover Texts, in the notice that says that the Document is released under this License. A Front-Cover Text may be at most 5 words, and a Back-Cover Text may be at most 25 words.

A "Transparent" copy of the Document means a machine-readable copy, represented in a format whose specification is available to the general public, that is suitable for revising the document straightforwardly with generic text editors or (for images composed of pixels) generic paint programs or (for drawings) some widely available drawing editor, and that is suitable for input to text formatters or for automatic translation to a variety of formats suitable for input to text formatters. A copy made in an otherwise Transparent file format whose markup, or absence of markup, has been arranged to thwart or discourage subsequent modification by readers is not Transparent. An image format is not Transparent if used for any substantial amount of text. A copy that is not "Transparent" is called "Opaque".

Examples of suitable formats for Transparent copies include plain ASCII without markup, Texinfo input format, LaTeX input format, SGML or XML using a publicly available DTD, and standard-conforming simple HTML, PostScript or PDF designed for human modification. Examples of transparent image formats include PNG, XCF and JPG. Opaque formats include proprietary formats that can be read and edited only by proprietary word processors, SGML or XML for which the DTD and/or processing tools are not generally available, and the machine-generated HTML, PostScript or PDF produced by some word processors for output purposes only.

The "Title Page" means, for a printed book, the title page itself, plus such following pages as are needed to hold, legibly, the material this License requires to appear in the title page. For works in formats which do not have any title page as such, "Title Page" means the text near the most prominent appearance of the work's title, preceding the beginning of the body of the text.

A section "Entitled XYZ" means a named subunit of the Document whose title either is precisely XYZ or contains XYZ in parentheses following text that translates XYZ in another language. (Here XYZ stands for a specific section name mentioned below, such as "Acknowledgements", "Dedications", "Endorsements", or "History".) To "Preserve the Title" of such a section when you modify the Document means that it remains a section "Entitled XYZ" according to this definition.

The Document may include Warranty Disclaimers next to the notice which states that this License applies to the Document. These Warranty Disclaimers are considered to be included by reference in this License, but only as regards disclaiming warranties: any other implication that these Warranty Disclaimers may have is void and has no effect on the meaning of this License.

2. VERBATIM COPYING

You may copy and distribute the Document in any medium, either commercially or noncommercially, provided that this License, the copyright notices, and the license notice saying this License applies to the Document are reproduced in all copies, and that you add no other conditions whatsoever to those of this License. You may not use technical measures to obstruct or control the reading or further copying of the copies you make or distribute. However, you may accept compensation in exchange for copies. If you distribute a large enough number of copies you must also follow the conditions in section 3.

You may also lend copies, under the same conditions stated above, and you may publicly display copies.

3. COPYING IN QUANTITY

If you publish printed copies (or copies in media that commonly have printed covers) of the Document, numbering more than 100, and the Document's license notice requires Cover Texts, you must enclose the copies in covers that carry, clearly and legibly, all these Cover Texts: Front-Cover Texts on the front cover, and Back-Cover Texts on the back cover. Both covers must also clearly and legibly identify you as the publisher of these copies. The front cover must present the full title with all words of the title equally prominent and visible. You may add other material on the covers in addition. Copying with changes limited to the covers, as long as they preserve the title of the Document and satisfy these conditions, can be treated as verbatim copying in other respects.

If the required texts for either cover are too voluminous to fit legibly, you should put the first ones listed (as many as fit reasonably) on the actual cover, and continue the rest onto adjacent pages.

If you publish or distribute Opaque copies of the Document numbering more than 100, you must either include a machine-readable Transparent copy along with each Opaque copy, or state in or with each Opaque copy a computer-network location from which the general network-using public has access to download using public-standard network protocols a complete Transparent copy of the Document, free of added material. If you use the latter option, you must take reasonably prudent steps, when you begin distribution of Opaque copies in quantity, to ensure that this Transparent copy will remain thus accessible at the stated location until at least one year after the last time you distribute an Opaque copy (directly or through your agents or retailers) of that edition to the public.

It is requested, but not required, that you contact the authors of the Document well before redistributing any large number of copies, to give them a chance to provide you with an updated version of the Document.

4. MODIFICATIONS

You may copy and distribute a Modified Version of the Document under the conditions of sections 2 and 3 above, provided that you release the Modified Version under precisely this License, with the Modified Version filling the role of the Document, thus licensing distribution and modification of the Modified Version to whoever possesses a copy of it. In addition, you must do these things in the Modified Version:

- A. Use in the Title Page (and on the covers, if any) a title distinct from that of the Document, and from those of previous versions (which should, if there were any, be listed in the History section of the Document). You may use the same title as a previous version if the original publisher of that version gives permission.
- B. List on the Title Page, as authors, one or more persons or entities responsible for authorship of the modifications in the Modified Version, together with at least five of the principal authors of the Document (all of its principal authors, if it has fewer than five), unless they release you from this requirement.

- C. State on the Title page the name of the publisher of the Modified Version, as the publisher.
- D. Preserve all the copyright notices of the Document.
- E. Add an appropriate copyright notice for your modifications adjacent to the other copyright notices.
- F. Include, immediately after the copyright notices, a license notice giving the public permission to use the Modified Version under the terms of this License, in the form shown in the Addendum below.
- G. Preserve in that license notice the full lists of Invariant Sections and required Cover Texts given in the Document's license notice.
- H. Include an unaltered copy of this License.
- I. Preserve the section Entitled "History", Preserve its Title, and add to it an item stating at least the title, year, new authors, and publisher of the Modified Version as given on the Title Page. If there is no section Entitled "History" in the Document, create one stating the title, year, authors, and publisher of the Document as given on its Title Page, then add an item describing the Modified Version as stated in the previous sentence.
- J. Preserve the network location, if any, given in the Document for public access to a Transparent copy of the Document, and likewise the network locations given in the Document for previous versions it was based on. These may be placed in the "History" section. You may omit a network location for a work that was published at least four years before the Document itself, or if the original publisher of the version it refers to gives permission.
- K. For any section Entitled "Acknowledgements" or "Dedications", Preserve the Title of the section, and preserve in the section all the substance and tone of each of the contributor acknowledgements and/or dedications given therein.
- L. Preserve all the Invariant Sections of the Document, unaltered in their text and in their titles. Section numbers or the equivalent are not considered part of the section titles.
- M. Delete any section Entitled "Endorsements". Such a section may not be included in the Modified Version.
- N. Do not retitle any existing section to be Entitled "Endorsements" or to conflict in title with any Invariant Section.
- O. Preserve any Warranty Disclaimers.

If the Modified Version includes new front-matter sections or appendices that qualify as Secondary Sections and contain no material copied from the Document, you may at your option designate some or all of these sections as invariant. To do this, add their titles to the list of Invariant Sections in the Modified Version's license notice. These titles must be distinct from any other section titles. You may add a section Entitled "Endorsements", provided it contains nothing but endorsements of your Modified Version by various parties—for example, statements of peer review or that the text has been approved by an organization as the authoritative definition of a standard.

You may add a passage of up to five words as a Front-Cover Text, and a passage of up to 25 words as a Back-Cover Text, to the end of the list of Cover Texts in the Modified Version. Only one passage of Front-Cover Text and one of Back-Cover Text may be added by (or through arrangements made by) any one entity. If the Document already includes a cover text for the same cover, previously added by you or by arrangement made by the same entity you are acting on behalf of, you may not add another; but you may replace the old one, on explicit permission from the previous publisher that added the old one.

The author(s) and publisher(s) of the Document do not by this License give permission to use their names for publicity for or to assert or imply endorsement of any Modified Version.

5. COMBINING DOCUMENTS

You may combine the Document with other documents released under this License, under the terms defined in section 4 above for modified versions, provided that you include in the combination all of the Invariant Sections of all of the original documents, unmodified, and list them all as Invariant Sections of your combined work in its license notice, and that you preserve all their Warranty Disclaimers.

The combined work need only contain one copy of this License, and multiple identical Invariant Sections may be replaced with a single copy. If there are multiple Invariant Sections with the same name but different contents, make the title of each such section unique by adding at the end of it, in parentheses, the name of the original author or publisher of that section if known, or else a unique number. Make the same adjustment to the section titles in the list of Invariant Sections in the license notice of the combined work.

In the combination, you must combine any sections Entitled "History" in the various original documents, forming one section Entitled "History"; likewise combine any sections Entitled "Acknowledgements", and any sections Entitled "Dedications". You must delete all sections Entitled "Endorsements".

6. COLLECTIONS OF DOCUMENTS

You may make a collection consisting of the Document and other documents released under this License, and replace the individual copies of this License in the various documents with a single copy that is included in the collection, provided that you follow the rules of this License for verbatim copying of each of the documents in all other respects.

You may extract a single document from such a collection, and distribute it individually under this License, provided you insert a copy of this License into the extracted document, and follow this License in all other respects regarding verbatim copying of that document.

7. AGGREGATION WITH INDEPENDENT WORKS

A compilation of the Document or its derivatives with other separate and independent documents or works, in or on a volume of a storage or distribution medium, is called an "aggregate" if the copyright resulting from the compilation is not used to limit the legal rights of the compilation's users beyond what the individual works permit. When the Document is included in an aggregate, this License does not apply to the other works in the aggregate which are not themselves derivative works of the Document.

If the Cover Text requirement of section 3 is applicable to these copies of the Document, then if the Document is less than one half of the entire aggregate, the Document's Cover Texts may be placed on covers that bracket the Document within the aggregate, or the electronic equivalent of covers if the Document is in electronic form. Otherwise they must appear on printed covers that bracket the whole aggregate.

8. TRANSLATION

Translation is considered a kind of modification, so you may distribute translations of the Document under the terms of section 4. Replacing Invariant Sections with translations requires special permission from their copyright holders, but you may include translations of some or all Invariant Sections in addition to the original versions of these Invariant Sections. You may include a translation of this License, and all the license notices in the Document, and any Warranty Disclaimers, provided that you also include the original English version of this License and the original versions of those notices and disclaimers. In case of a disagreement between the translation and the original version of this License or a notice or disclaimer, the original version will prevail.

If a section in the Document is Entitled "Acknowledgements", "Dedications", or "History", the requirement (section 4) to Preserve its Title (section 1) will typically require changing the actual title.

9. TERMINATION

You may not copy, modify, sublicense, or distribute the Document except as expressly provided for under this License. Any other attempt to copy, modify, sublicense or distribute the Document is void, and will automatically terminate your rights under this License. However, parties who have received copies, or rights, from you under this License will not have their licenses terminated so long as such parties remain in full compliance.

10. FUTURE REVISIONS OF THIS LICENSE

The Free Software Foundation may publish new, revised versions of the GNU Free Documentation License from time to time. Such new versions will be similar in spirit to the present version, but may differ in detail to address new problems or concerns. See <http://www.gnu.org/copyleft/>.

Each version of the License is given a distinguishing version number. If the Document specifies that a particular numbered version of this License "or any later version" applies to it, you have the option of following the terms and conditions either of that specified version or of any later version that has been published (not as a draft) by the Free Software Foundation. If the Document does not specify a version number of this License, you may choose any version ever published (not as a draft) by the Free Software Foundation.

ADDENDUM: How to use this License for your documents

```
Copyright (c) YEAR YOUR NAME.  
Permission is granted to copy, distribute and/or modify this document  
under the terms of the GNU Free Documentation License, Version 1.2  
or any later version published by the Free Software Foundation;  
with no Invariant Sections, no Front-Cover Texts, and no Back-Cover Texts.  
A copy of the license is included in the section entitled "GNU  
Free Documentation License".
```

If you have Invariant Sections, Front-Cover Texts and Back-Cover Texts, replace the "with...Texts." line with this:

```
with the Invariant Sections being LIST THEIR TITLES, with the  
Front-Cover Texts being LIST, and with the Back-Cover Texts being LIST.
```

If you have Invariant Sections without Cover Texts, or some other combination of the three, merge those two alternatives to suit the situation.

If your document contains nontrivial examples of program code, we recommend releasing these examples in parallel under your choice of free software license, such as the GNU General Public License, to permit their use in free software.