



Cisco TelePresence VCS Cluster Creation and Maintenance

Deployment Guide

First Published: December 2009

Last Updated: January 2017

Cisco VCS X8.9.1

Cisco TMS 13.2 or later

Contents

Preface	3
Change History	3
Introduction	5
Form a Cluster	7
Clustering Prerequisites	7
Create a New Cluster of VCS Peers	9
Add a Peer to a Cluster	13
Peer-Specific Items	18
Connect a Cluster	20
Neighboring Between VCS Clusters	20
Configure Endpoints to Work With a Cluster	21
Add a VCS to Cisco TMS	24
Upgrade a Cluster	26
Upgrade an X7.1 (or Later) Cluster to X8.9.1	26
Change a Cluster	29
Remove a Live Peer From a Cluster (Permanently)	29
Remove a Dead Peer From a Cluster (Permanently)	31
Disband a Cluster	33
Change the Primary Peer	35
Change the Address of a Peer	36
Replace a Peer	37
Troubleshooting	38
VCS Alarms and Warnings	38
Cisco TMS Warnings	39
Reference	41
Appendix 1: IP Ports and Protocols	41
Appendix 2: Cluster Name and DNS SRV Records	42
Appendix 3: Clusters in Isolated Networks	45
Appendix 4: NAPTR Records	46
Appendix 5: Impact of Clustering on Other VCS Applications	48
Appendix 6: Upgrade Earlier VCS Versions to X8.9.1	49

Preface

Cisco Legal Information	51
Cisco Trademark	51

Preface

Change History

Table 1 VCS Cluster Deployment Guide Change History

Date	Change	Reason
January 2017	Removed section on downgrades as they are not supported.	X8.9.1 release
December 2016	Added section on clusters in isolated networks in relation to TLS.	X8.9 release
June 2016	Cluster communications now use TLS. Registrations, FindMe, TMSPE support introduced on Expressway.	X8.8 release
November 2015	Updated for X8.7.	
July 2015	Updated for X8.6. New procedure for replacing a peer.	
April 2015	Menu path changes for X8.5 onwards. Republished with X8.5.2.	
December 2014	Updated for X8.5.	
June 2014	Republished for X8.2.	
April 2014	Updated for VCS X8.1.1: <ul style="list-style-type: none"> ■ New 'Replacing a VCS peer' section ■ Updates to 'IP ports and protocols' appendix 	
December 2013	Updated for VCS X8.1, including the removal of support for TMS Agent legacy mode provisioning.	
November 2012	Removed recommendation to use DNS name of the VCS cluster for endpoint SIP outbound configuration.	
August 2012	Updated for VCS X7.2.	
March 2012	Updated for VCS X7.1.	
January 2012	Various minor revisions and clarifications applied.	
October 2011	Updated for VCS X7.0.n.	
September 2011	Updated Appendix 3 - IP port and protocol numbers.	
August 2011	Updated for VCS X7.0.	

Table 1 VCS Cluster Deployment Guide Change History (continued)

Date	Change	Reason
May 2011	Updated for VCS X6.1.	
February 2011	Clarify this document refers to VCS X6.0 only, and to Cisco TMS 12.6 or later.	
December 2010	Updated for VCS X6.	
October 2010	New document styles applied.	
August 2010	Updated for VCS X5.2. Replace previous Appendix 12 (Upgrading debian packages - only required for X5.1.1) with a new Appendix 12 - NAPTR records overview.	
May 2010	Updated for VCS X5.1.1 and Cisco TMS 12.6.	
December 2009	Updated for VCS X5.	
Up to December 2009	Releases for earlier versions of VCS.	

Introduction

Cisco TelePresence Video Communication Server (VCS) clusters are designed to extend the resilience and capacity of a VCS installation. VCS peers in a cluster share bandwidth usage as well as routing, zone, FindMe™ and other configuration among themselves. Endpoints can register to any of the peers in the cluster; if they lose connection to their initial peer, they can re-register to another peer in the cluster.

Capacity licensing is carried out on a per-cluster basis. Any capacity licenses that have been installed on a cluster peer are available for use by any peer within the cluster. If a cluster peer becomes unavailable, the license capacity installed on that peer remains available to the rest of the cluster for two weeks after it lost contact with the peer. This will maintain the overall license capacity of the cluster. Note that each peer is always limited by its physical capacity; the license capacity borrowing is only intended to give you time to repair your cluster.

"Capacity" includes the following license types:

- Traversal and non-traversal call licenses
- TURN relay licenses

Every VCS peer in the cluster must have the same routing capabilities – if any VCS can route a call to a destination it is assumed that all VCS peers in that cluster can route a call to that destination. If the routing is different on different VCS peers, then separate VCSs / VCS clusters must be used.

This guide describes how to create and maintain VCS clusters. It provides information on:

- [Upgrade an X7.1 \(or Later\) Cluster to X8.9.1, page 26](#)
- [Create a New Cluster of VCS Peers, page 9](#)
- [Add a Peer to a Cluster, page 13](#)
- [Peer-Specific Items, page 18](#)
- [Neighboring Between VCS Clusters, page 20](#)
- [Configure Endpoints to Work With a Cluster, page 21](#)
- [Add a VCS to Cisco TMS, page 24](#)
- [Remove a Live Peer From a Cluster \(Permanently\), page 29](#)
- [Remove a Dead Peer From a Cluster \(Permanently\), page 31](#)
- [Disband a Cluster, page 33](#)
- [Change the Primary Peer, page 35](#)
- [Change the Address of a Peer, page 36](#)
- [Replace a Peer, page 37](#)

Cisco TMS is required if FindMe configuration is to be sourced from Cisco TMS and also if endpoints are to be provisioned, but is not essential for clustering.

FindMe can be configured locally on the VCS cluster if required, but be aware that if FindMe data is stored on the VCS cluster and not sourced by Cisco TMS, there is no way to later move that FindMe information to Cisco TMS.

Enabling provisioning and creating a cluster are two separate processes. If you intend to enable provisioning on your cluster, either:

- follow the instructions in this guide to create the cluster of VCSs (without provisioning enabled), and then follow the instructions in *Cisco TMS Provisioning Extension Deployment Guide* to enable provisioning across the cluster, or
- follow the instructions in *Cisco TMS Provisioning Extension Deployment Guide* to enable provisioning on what will be the primary VCS, and then follow the instructions in this guide to create the cluster of VCSs

Note that VCS X8 and later supports only Cisco TMS Provisioning Extension mode.

Introduction

For more information about TMS Agent legacy mode provisioning, or for creating, modifying, and troubleshooting clusters that will remain running on earlier software versions, see earlier versions of this deployment guide. Also, see [Appendix 6: Upgrade Earlier VCS Versions to X8.9.1, page 49](#) for information about upgrading these clusters to X8.9.1.

Form a Cluster

Clustering Prerequisites

Before setting up a cluster of X8.9.1 VCS peers or adding an X8.9.1 VCS to a cluster, ensure that:

Platform and Software Versions Match

- All clusters peers are running the same version of code. The only occasion where different peers are allowed to run different versions of code is for the short period of time while a cluster is being upgraded from one version of code to another, during which time the cluster will operate in a partitioned fashion.
- Each peer is using a hardware platform (appliance or virtual machine) with equivalent capabilities; for example, you can cluster peers that are running on standard appliances with peers running on 2 core Medium VMs, but you cannot cluster a peer running on a standard appliance with peers running on 8 core Large VMs.

Network Conditions Are Met

- Each peer has a different LAN configuration (a different IPv4 address and a different IPv6 address, where enabled).
- Each peer in a cluster is within a 15ms hop (30ms round trip delay) of each and every other VCS in or to be added to the cluster.
- Each peer in a cluster is directly routable to each and every other VCS in or to be added to the cluster. (There must be no NAT between cluster peers – if there is a firewall ensure that the required ports are opened.)

Basic Configuration Is Done

- Each peer has a different system name to all other peers.
- Each peer has a certificate that identifies it to other peers, creating unauthenticated TLS connections between peers (default of **TLS verification mode** set to *Permissive*).

If you wish to have authenticated TLS connections, the certificate must also be valid and be issued by an authority that is trusted by all peers (**TLS Verification mode** set to *Enforce*).

We recommend populating the CN of all peer certificates with the same cluster FQDN, and populating each peer certificate's SAN with that peer's FQDN.

- All peers have the same set of option keys installed, with the following exceptions:
 - Traversal and non-traversal call licenses
 - TURN relay licenses

All other license keys must be identical on each peer.

Note: Installing some types of option keys requires you to restart the VCS.

- H.323 mode is enabled on each peer (**Configuration > Protocols > H.323**, and for **H.323 mode** select *On*).
The cluster uses H.323 signaling between peers to determine the best route for calls, even if all endpoints are SIP endpoints.

DNS Configuration is Done

DNS server configuration does not replicate so you must enter the DNS server address(es) on each peer.

- The DNS servers used by the VCS peers must support both forward and reverse DNS lookups of Cisco TMS and all VCS peer addresses; the DNS servers must also provide address lookup for any other DNS functionality required, such as:
 - NTP servers or the external manager if they configured using DNS names
 - Microsoft FE Server FQDN lookup
 - LDAP server forward and reverse lookup (reverse lookups are frequently provided through PTR records)

Form a Cluster

Note: In an isolated network such as a DMZ, peers may not be able to resolve each others' FQDNs to private IP addresses. This could prevent the cluster from forming.

See [Appendix 3: Clusters in Isolated Networks, page 45](#).

- A DNS SRV record is available for the cluster which contains A or AAAA records for each peer of the cluster. This configuration is advised for video interoperability and business to business (B2B) video calling, but is **not required for Mobile and Remote Access**.
- (For MRA) Create a `collab-edge` SRV record for each peer in the VCS Expressway cluster
- (For B2B only) The VCS Expressway cluster has a DNS SRV record that defines all cluster peers

TMS Has Been Configured (if necessary)

- Cisco TMS, if used, is running version 13.2 or later (12.6 or later is permitted if you are not using Cisco TMS for provisioning or FindMe).
- If Cisco TMS is to be used for replicating FindMe and/or Provisioning data, ensure that Provisioning Extension mode functionality has been enabled on Cisco TMS (see [Cisco TMS Provisioning Extension Deployment Guide](#) for details).

Create a New Cluster of VCS Peers

Note: This procedure requires a period of downtime for the VCS service. Ensure that these instructions are followed in a scheduled maintenance window.

You **must** use Cisco TMS if:

- You want to use Device Provisioning with the VCS cluster
- You want to use FindMe with the VCS cluster

If you are not using Device Provisioning or FindMe with your cluster, then Cisco TMS is optional but recommended.

Process Summary

This process initiates a cluster of a single VCS. Do not use this process if the cluster already exists.

To complete the cluster containing multiple VCSs, when this section is complete, follow the instructions in [Add a Peer to a Cluster, page 13](#) to add the subordinate peers to the cluster.

Prerequisites

- Cisco TMS, if used, is running version 13.2 or later (12.6 or later is permitted if you are not using Cisco TMS for provisioning or FindMe).
- If you're using Cisco TMS for Device Provisioning or FindMe, then this VCS and Cisco TMS must be proven to be operating in Provisioning Extension mode.
- If a firewall exists between cluster peers, it must be configured to permit the traffic documented in [Appendix 1: IP Ports and Protocols, page 41](#).
- All VCSs to be included in the cluster must be running the same version (X8.9.1) of VCS software.

Configure the Primary Peer

This process sets up the first (primary) peer of this new cluster – additional peers are added afterwards using the [Add a Peer to a Cluster, page 13](#) process.

Before proceeding, the VCS that will be the primary must be determined. The primary VCS will be the source of the configuration information for all VCS peers in the cluster. Subordinate VCS peers will have their configuration deleted and replaced by that from the primary.

On other VCSs:

Check that no other VCS (anywhere in the organization) has this VCS's IP address in its clustering peers list.

On this primary VCS:

1. Check that the VCS is running X8.9.1 software.
2. Backup the VCS (**Maintenance > Backup and restore**).

Form a Cluster

3. Using the web interface, review the configuration to ensure that the VCS has:

- A valid Ethernet speed (**System > Network interfaces > Ethernet**).
- Valid IP address and IP gateway (**System > Network interfaces > IP**).
- The same set of option keys installed as those that will be installed on all other peers of the cluster (**Maintenance > Option keys**).

Note that the number of call license keys may be different on different peers; all other license keys must be identical on each peer.

- At least one valid DNS server configured, and that if unqualified DNS names are used elsewhere (e.g. for the NTP server), that the correct **Domain name** is also configured (**Domain name** is added as a suffix to an unqualified DNS name to make it into an FQDN) (**System > DNS**).
 - A valid and working NTP server configured (**System > Time**; in the Status section, the State should be "Synchronized").
 - No peers configured (**System > Clustering** – all Peer N address fields on this page should be blank. If not, clear the fields and click **Save**).
 - **H.323 Mode** set to *On* (**Configuration > Protocols > H.323**)
4. Ensure that this VCS does not list any of the VCSs that are to be peers in this new cluster in any of its neighbor zones or traversal zones (**Configuration > Zones > Zones** then check each neighbor and traversal zone).
5. Set the **H.323 Time to live** to an appropriate value for the size of your deployment. A smaller number, like 60 (seconds), means that if a VCS becomes inaccessible, the endpoint will quickly register with another peer (**Configuration > Protocols > H.323**).

Note: By reducing the registration time to live too much, you risk flooding the VCS with registration requests, which will severely impact performance. This impact is proportional to the number of endpoints, so you should balance the need for occasional quick failover against the need for continuous good performance.

6. Go to **System > DNS** and ensure that **System host name** is the DNS hostname for this VCS (typically the same as the **System name** in **System > Administration**, but excluding spaces, and unique for each VCS in the cluster). If it is not configured correctly, set it up appropriately and click **Save**.

Note: <System host name>.<DNS domain name> = FQDN of this VCS

7. Go to **Configuration > Call routing** and set **Call signaling optimization** to *On*.

8. Click **Save**.

9. Enable maintenance mode:

- a. Go to **Maintenance > Maintenance mode**.
- b. Set **Maintenance mode** to *On*.
- c. Click **Save** and click **OK** on the confirmation dialog.

10. Wait for all calls to clear and registrations to timeout on this peer.

- If necessary, manually remove any calls on this peer that do not clear automatically (using the web browser go to **Status > Calls**, then select the check box next to the calls you want to terminate and click **Disconnect**).
- If necessary, manually remove any registrations from this peer that do not clear automatically (using the web browser go to **Status > Registrations > By device**, then select the check box next to the devices you want to remove and click **Unregister**).

You can leave the registration for the Conference Factory – this will not be the source of calls, and even if deleted will not roll over to another peer, as other peers have their own Conference Factory registration (if enabled).

Form a Cluster

11. (Not applicable to MRA) Go to **System > Clustering** and ensure that **Cluster name** is the routable Fully Qualified Domain Name used in SRV records that address this VCS cluster, for example `cluster1.example.com`. (See [Appendix 2: Cluster Name and DNS SRV Records, page 42](#)).

Change the **Cluster name** if necessary.

12. Click **Save**.
13. On the **Clustering** page configure the fields as follows:

Configuration primary	1
Cluster IP version	Choose <i>IPv4</i> or <i>IPv6</i> to match the underlying network addressing scheme.
TLS verification mode	Choose <i>Permissive</i> or <i>Enforce</i> . <i>Permissive</i> means that the peers do not validate each others' certificates when establishing intracluster TLS connections. <i>Enforce</i> is more secure, but requires that each peer has a valid certificate and that the signing CA is trusted by all other peers.
Peer 1 address	Enter the address of this VCS (the primary peer). If TLS verification mode is set to <i>Enforce</i> , then you must enter an FQDN that matches the subject CN or a SAN on this peer's certificate.

14. Click **Save**.
To the right of the **Peer 1 address** field the words "**This system**" should appear (though this may require the page to be refreshed before they appear).
15. Restart the VCS (go to **Maintenance > Restart options**, then click **Restart** and confirm **OK**).

After the restart, on the primary VCS web interface:

- Check that configuration data exists as expected:
 - If FindMe is in use, check that the expected FindMe entries still exist (**Status > Applications > TMS Provisioning Extension Services > FindMe > Accounts**).
 - Check configuration for items from the **System**, **Configuration** and **Application** menus.
- Backup the VCS (**Maintenance > Backup and restore**).

On other devices:

If you have any other VCSs neighbored (or connected via a traversal zone) to this primary VCS peer, ensure that their zone configuration for this cluster is updated to only include the address of this primary VCS.

On this primary VCS peer:

- Log in to the web browser of this VCS.
- Go to **Status > Alarms**.
If there is an alarm that the VCS must be restarted, go to **Maintenance > Restart options** and then click **Restart**.

Form a Cluster

3. If the VCS did not need to be restarted, ensure that maintenance mode is disabled.
 - a. Go to **Maintenance > Maintenance mode**.
 - b. Set **Maintenance mode** to *Off*.
 - c. Click **Save**.

Next Steps

Go to **Status > Alarms** and ensure that all alarms are acted upon and cleared.

The setup of the new cluster (of one VCS) is complete; add other VCSs to the cluster using [Add a Peer to a Cluster](#), page 13.

If you are integrating the cluster with Microsoft infrastructure, see [Appendix 5: Impact of Clustering on Other VCS Applications](#), page 48.

Add a Peer to a Cluster

Follow this process if you have an existing X8.9.1 cluster (of one or more peers) to which you want to add another VCS peer. If you do not have an existing cluster, following the instructions in the section [Create a New Cluster of VCS Peers, page 9](#).

This process will add an X8.9.1 VCS to the cluster and replicate the cluster primary's configuration onto that VCS. Note that:

- You can have up to 6 VCSs, including the primary VCS, in a cluster.
- Only one VCS must be added to the cluster at a time.

You should only make configuration changes on the primary VCS.

Caution: Do not adjust any cluster-wide configuration until the cluster is stable with all peers running. Cluster database replication will be negatively impacted if any peers are upgrading, restarting, or out of service when you change the cluster's configuration.

Any changes made on other peers are not reflected across the cluster, and will be overwritten the next time the primary's configuration is replicated across the peers. The only exceptions to this are some [peer-specific configuration items](#).

You may need to wait up to one minute before changes are updated across all peers in the cluster.

Prerequisites

- If Device Provisioning or FindMe are enabled, the existing cluster of 1 or more peers must already be running using Provisioning Extension mode.
- All VCSs to be included in the cluster must be running the same version (X8.9.1) of VCS software.

Add the New Peer

On the primary VCS:

Ensure that the primary VCS does not list this new VCS peer in any of its neighbor zones or traversal zones (**Configuration > Zones > Zones**).

Note: The primary VCS will be the source of the configuration for this new VCS peer and all other VCS peers in the cluster. When a VCS is added to the cluster, its configuration will be deleted and replaced by that from the primary.

On the VCS to be added to the cluster:

1. Check that no other VCS (anywhere in the organization) has this VCS's IP address in its clustering peers list.
2. Backup the VCS (**Maintenance > Backup and restore**).

Form a Cluster

3. Using the web interface, review the configuration to ensure that the VCS has:
 - A valid Ethernet speed (**System > Network interfaces > Ethernet**).
 - Valid IP address and IP gateway (**System > Network interfaces > IP**).
 - The same set of option keys installed as those that will be installed on all other peers of the cluster (**Maintenance > Option keys**).

Note that the number of call license keys may be different on different peers; all other license keys must be identical on each peer.

 - At least one valid DNS server configured, and that if unqualified DNS names are used elsewhere (e.g. for the NTP server), that the correct **Domain name** is also configured (**Domain name** is added as a suffix to an unqualified DNS name to make it into an FQDN) (**System > DNS**).
 - A valid and working NTP server configured (**System > Time**; in the Status section, the State should be “Synchronized”).
 - No peers configured (**System > Clustering** – all Peer N address fields on this page should be blank. If not, clear the fields and click **Save**).
 - **H.323 Mode** set to *On* (**Configuration > Protocols > H.323**)

If the cluster is managed by Cisco TMS, this VCS being added to the cluster must also be managed by Cisco TMS.

If Cisco TMS is used, on the VCS to be added to the cluster:

1. Ensure that the VCS can see Cisco TMS.
To do this, go to **System > External manager** and in the Status section, ensure that the **State** is **Active**.
If it is not active, follow the process in [Add a VCS to Cisco TMS, page 24](#).
2. Ensure that Cisco TMS can communicate with this VCS.
To do this, on Cisco TMS go to **Systems > Navigator** (and any required sub folders) then click on the name of the VCS and ensure that it says:
“✓ System has no open or acknowledged tickets”
If it does not say this, follow the process in [Add a VCS to Cisco TMS, page 24](#).

On the VCS being added to the cluster:

Go to **Status > Alarms**. If there is an alarm that the VCS must be restarted, go to **Maintenance > Restart options** and then click **Restart**.

If multiple peers need restarting: restart each peer in turn, waiting for the peer to be accessible through the web interface before restarting the next peer.

If Cisco TMS is used:

1. For the VCS to be added to the cluster, ensure that the Host Name of the VCS is set up in Cisco TMS:
 - a. Go to **Systems > Navigator** (and any required sub folders).
 - b. Select this VCS.
 - c. Select the **Connection** tab.
 - d. Set **Host Name** to be the FQDN of this subordinate peer, for example vcs3.uk.company.com.
 - e. Click **Save/Try**.
You can ignore any error messages such as “DNS config failure resolving <DNS name>: Did not find system IP address () in DNS: <Server IP>”
 - f. Ensure that Cisco TMS updates its DNS.
 1. Select the **Settings** tab.
 2. Click **Force Refresh**.

Form a Cluster

2. Delete FindMe accounts from the VCS.

If Cisco TMS is configured to replicate FindMe information with this VCS being added to the cluster, then you should delete any user accounts that existed on the VCS.

To check: On Cisco TMS, go to **Systems > Navigator** (and any required sub folders), select this VCS, select the **Provisioning** tab. There you can see if the FindMe services box is checked.

- If FindMe replication is not enabled (box not checked), then no FindMe data is stored on Cisco TMS. You do not need to delete the FindMe entries from the VCS.
- If FindMe data is provisioned from Cisco TMS (box is checked), delete FindMe user accounts from the VCS as follows:
 - a. Go to **Users > FindMe accounts**.
 - b. Select all of the accounts shown and click **Delete**.

On the VCS being added to the cluster:**1. Enable maintenance mode:**

- a. Go to **Maintenance > Maintenance mode**.
- b. Set **Maintenance mode** to *On*.
- c. Click **Save** and click **OK** on the confirmation dialog.

2. Wait for all calls to clear and registrations to timeout on this peer.

- If necessary, manually remove any calls on this peer that do not clear automatically (using the web browser go to **Status > Calls**, then select the check box next to the calls you want to terminate and click **Disconnect**).
- If necessary, manually remove any registrations from this peer that do not clear automatically (using the web browser go to **Status > Registrations > By device**, then select the check box next to the devices you want to remove and click **Unregister**).

You can leave the registration for the Conference Factory – this will not be the source of calls, and even if deleted will not roll over to another peer, as other peers have their own Conference Factory registration (if enabled).

On the primary VCS:**1. Go to System > Clustering.**

One or more of the **Peer N address** fields should be empty.

2. In the first empty field, enter the address of the new VCS peer.**Notes:**

- Add one VCS to the cluster at a time (repeat the process for each peer you add).
- If the new VCS peer has the Advanced Networking option, the **Peer N address** MUST specify the LAN 1 interface. This interface must not have NAT enabled on it.

3. Click Save.

Peer 1 should indicate 'This system'. The new peer may indicate 'Unknown' and then with a refresh should indicate 'Failed' because it has not fully joined the cluster yet.

Notes:

- A cluster communication failure alarm will be raised on the primary and on other peers already in the cluster advising that this new VCS peer is not communicating – this will clear later.
- Cluster configuration replication is suspended until the new VCS has properly joined the cluster. Configuration changes will not replicate yet.

Form a Cluster

On every other subordinate VCS already in the cluster (not the VCS being added):

1. Go to **System > Clustering** and configure the fields as follows:

Cluster name	Identical to the Cluster name configured on the primary VCS
Configuration primary	Same number as chosen on the primary VCS
Cluster IP version	Same version as chosen on the the primary VCS
TLS verification mode	Same setting as chosen on the primary VCS
Peer 1 address ...Peer 6 address	The addresses should be the same, and in the same order, as those entered on the primary VCS

2. Click **Save**.

On the additional subordinate VCS being added to the cluster:

1. Log in as admin on an SSH or other CLI interface. At the VCS command prompt, type:

```
xcommand DefaultValueSet Level: 2
```

```
xcommand DefaultLinksAdd
```

Note: This command removes any LDAP authentication configuration – ensure that you have the web admin password before executing this command.

2. Go to **Users > Administrator accounts**.
3. Delete all entries except the default admin account.
4. Go to **System > DNS** and ensure that **System host name** is the DNS hostname for this VCS (typically the same as the **System name** in **System > Administration**, but excluding spaces, and unique for each VCS in the cluster). If it is not configured correctly, set it up appropriately and click **Save**.

Note: <System host name>.<DNS domain name> = FQDN of this VCS

5. Go to **System > Clustering** and configure the fields as follows:

Cluster name	Identical to the Cluster name configured on the primary VCS
Configuration primary	Same number as chosen on the primary VCS
Cluster IP version	Same version as chosen on the the primary VCS
TLS verification mode	Same setting as chosen on the primary VCS
Peer 1 address ...Peer 6 address	The addresses should be the same, and in the same order, as those entered on the primary VCS

6. Click **Save**.

Note that a cluster communication failure alarm will be raised on this VCS peer advising that this new VCS is not communicating – this will clear after the restart.

7. Restart the VCS (go to **Maintenance > Restart options**, then click **Restart** and confirm **OK**).

If multiple peers need restarting: restart each peer in turn, waiting for the peer to be accessible through the web interface before restarting the next peer.

8. After the restart, wait approximately 2 minutes – this is the frequency with which configuration is copied from the primary.

Form a Cluster

9. Go to **Status > Alarms**. If there is an alarm that the VCS must be restarted, go to **Maintenance > Restart options** and then click **Restart**.

If multiple peers need restarting: restart each peer in turn, waiting for the peer to be accessible through the web interface before restarting the next peer.

10. Check that configuration data exists as expected:

- If FindMe is in use, check that the expected FindMe entries still exist (**Status > Applications > TMS Provisioning Extension Services > FindMe > Accounts**).
- Check configuration for items from the **System**, **Configuration** and **Application** menus.

On other devices:

If you have any other VCSs neighbored (or connected via a traversal zone) to this cluster of VCS peers, ensure that their zone configuration for this cluster is updated to include the address of this new peer.

On each VCS peer (including the primary and this new VCS peer):

Go to **Status > Alarms**. If there is an alarm that the VCS must be restarted, go to **Maintenance > Restart options** and then click **Restart**.

If multiple peers need restarting: restart each peer in turn, waiting for the peer to be accessible through the web interface before restarting the next peer.

Check Replication Status

It can take 5 or more minutes before VCS reports the successful status. If problems are seen, refresh the screen after waiting 5 minutes.

On VCS:

1. Go to **System > Clustering** and check that the cluster database status reports as Active.
2. If Cisco TMS is being used, ensure that Cisco TMS has all the correct settings for this upgraded VCS by forcing a refresh of Cisco TMS.

On Cisco TMS:

1. For every VCS in the cluster (including the primary VCS):
 - a. Select **Systems > Navigator** (and any required sub folders) and click on the name of the VCS.
 - b. Select the **Settings** tab.
 - c. Click **Force Refresh**.
2. Repeat for all VCS peers in the cluster.

Next Steps

If Conference Factory (Multiway™) is to be used, see [Appendix 5: Impact of Clustering on Other VCS Applications, page 48](#).

If the cluster has non-default trusted CA certificate and / or non default server certificate ensure that the added peer is configured with the required trusted CA certificate and an appropriate server certificate.

Adding an X8.9.1 VCS to a VCS X8.9.1 cluster is now complete.

Peer-Specific Items

Most items of configuration are applied via the primary peer to all peers in a cluster. However, the following items (marked with a † on the web interface) must be specified separately on each cluster peer.

Note: You should not modify configuration data that applies to all peers on any peer other than the primary peer. At best it will result in the changes being overwritten from the primary; at worst it will cause cluster replication to fail.

Cluster configuration (System > Clustering)

The list of **Peer N addresses** (including the peer's own IP address) that make up the cluster has to be specified on each peer and they must be identical on each peer.

The **Cluster name**, **Configuration primary**, **Cluster IP version**, and **TLS verification mode** must be specified on each peer and must be identical for all peers.

Ethernet speed (System > Network interfaces > Ethernet)

The **Ethernet speed** is specific to each peer. Each peer may have slightly different requirements for the connection to their Ethernet switch.

IP configuration (System > Network interfaces > IP)

LAN configuration is specific to each peer.

- Each peer must have a different **IPv4 address** and a different **IPv6 address**.
- **IP gateway** configuration is peer-specific. Each peer can use a different gateway.

Note that the IP protocol is applied to all peers, because each peer must support the same protocols.

IP static routes (System > Network interfaces > Static routes)

Any static routes you add are peer-specific and you may create different routes on different peers if required. If you want all peers in the cluster to be able to use the same static route, you must create the route on each peer.

System name (System > Administration)

The **System name** must be different for each peer in the cluster.

DNS servers and DNS host name (System > DNS)

DNS servers are specific to each peer. Each peer can use a different set of DNS servers.

The **System host name** and **Domain name** are specific to each peer.

NTP servers and time zone (System > Time)

The **NTP servers** are specific to each peer. Each peer may use one or more different NTP servers.

The **Time zone** is specific to each peer. Each peer may have a different local time.

SNMP (System > SNMP)

SNMP settings are specific to each peer. They can be different for each peer.

Logging (Maintenance > Logging)

The Event Log and Configuration Log on each peer only report activity for that particular VCS. The **Log level** and the list of **Remote syslog servers** are specific to each peer. We recommend that you set up a remote syslog server to which the logs of all peers can be sent. This allows you to have a global view of activity across all peers in the cluster.

Security certificates (Maintenance > Security certificates)

The trusted CA certificate, server certificate and certificate revocation lists (CRLs) used by the VCS must be uploaded individually per peer.

Form a Cluster

Administration access (System > Administration)

The following system administration access settings are specific to each peer:

- Serial port / console
- SSH service
- Web interface (over HTTPS)
- Redirect HTTP requests to HTTPS
- Automated protection service

Option keys (Maintenance > Option keys)

Option keys that control features are specific to the peer where they are applied. Option keys that control licenses are pooled for use by the whole cluster.

Each peer must have an identical set of feature option keys installed, which means you must purchase a key for each peer in the cluster.

License option keys can be applied to one or more peers in the cluster, and the sum of the installed licenses is available across the cluster. This license pooling behavior includes the following option keys:

- Traversal calls
- Non-traversal calls
- TURN relays

Note: In some cases a peer will raise an alarm that it has no key to enable licenses the peer needs, even though there are licenses available in the cluster. You can acknowledge and ignore this category of alarm, unless the only peer that has the required licenses is out of service.

Active Directory Service (Configuration > Authentication > Devices > Active Directory Service)

When configuring the connection to an Active Directory Service for device authentication, the **NetBIOS machine name (override)**, and domain administrator **Username** and **Password** are specific to each peer.

Conference Factory template (Applications > Conference Factory)

The template used by the Conference Factory application to route calls to the MCU is peer-specific, as it must be unique for each peer in the cluster.

VCS front panel display mode (configurable through CLI only)

The `xConfiguration Administration LCDPanel Mode CLI` setting is specific to each peer.

Connect a Cluster

Neighboring Between VCS Clusters

You can neighbor your local VCS (or VCS cluster) to a remote VCS cluster; this remote cluster could be a neighbor, traversal client, or traversal server to your local VCS. In this case, when a call is received on your local VCS and is passed via the relevant zone to the remote cluster, it will be routed to whichever peer in that neighboring cluster has the lowest resource usage. That peer will then forward the call as appropriate to one of its:

- locally registered endpoints (if the endpoint is registered to that peer)
- peers (if the endpoint is registered to another peer in that cluster)
- external zones (if the endpoint has been located elsewhere)

Lowest resource usage is determined by comparing the number of available traversal calls (maximum - current use) on the peers, and choosing the peer with the highest number. Peers that are in maintenance mode are not considered.

When configuring a connection to a remote cluster, you create a single zone and configure it with details of all the peers in the cluster. Adding this information to the zone ensures that the call is passed to that cluster regardless of the status of the individual peers.

You also need to enter the IP address of all peers in the remote cluster when the connection is via a **neighbor** or **traversal client** zone. You do not do this for **traversal server** zones, as these connections are not configured by specifying the remote system's IP address.

Note: Systems that are configured as peers must **not** also be configured as neighbors to each other, and vice versa.

Neighboring your clusters

To neighbor your local VCS (or VCS cluster) to a remote VCS cluster, you create a single zone to represent the cluster and configure it with the details of all the peers in that cluster:

1. On your local VCS (or, if the local VCS is a cluster, on the primary peer), create a zone of the appropriate type. This zone will represent the connection to the cluster.
2. In the **Location** section, enter the IP address or FQDN of each peer in the remote cluster in the **Peer 1 to Peer 6** address fields.

Note that:

- Ideally you should use FQDNs in these fields. Each FQDN must be different and must resolve to a single IP address for each peer. With IP addresses, you may not be able to use TLS verification, because many CAs will not supply certificates to authenticate an IP address.
- The order in which the peers in the remote VCS cluster are listed here does not matter.
- Whenever you add an extra VCS to a cluster (to increase capacity or improve redundancy, for example) you will need to modify any VCSs which neighbor to that cluster to let them know about the new cluster peer.

Configure Endpoints to Work With a Cluster

When configuring endpoints it is desirable for them to know about all the VCS peers in a cluster, so that at initial registration or later, if they lose connection to their VCS peer, they have the ability to register with and use another peer in the VCS cluster.

SIP and H.323 endpoints behave differently – the following sections show the methods that can be used, and list them in preferred order.

For additional details about DNS SRV and round-robin DNS see the URI dialing section in *VCS Administrator Guide*.

Also see [Appendix 2: Cluster Name and DNS SRV Records, page 42](#).

SIP Endpoints

The options below are listed in preference order for providing resilience of connectivity of endpoints to a cluster of VCSs where 1 or more VCS cluster peers become inaccessible. The choice of option will depend on what functionality the endpoint you are using supports.

Option 1 - SIP Outbound (preferred)

SIP outbound allows an endpoint to be configured to register to 2 or more VCS peers simultaneously. The benefit of this is that if the connection between one peer and the endpoint gets broken, then a connection from the endpoint to the other peer remains. With the endpoint registering to both peers simultaneously, there is no break in service while the endpoint realizes that its registration has failed, before it registers to a different peer. Thus, at no time is the endpoint unreachable.

Configuration of SIP outbound is endpoint specific, but typically will be:

- Proxy 1
 - Server discovery = Manual
 - Server Address =
DNS name of cluster peer or
IP address of cluster peer
- Proxy 2
 - Server discovery = Manual
 - Server Address =
DNS name of a different cluster peer or
IP address of a different cluster peer
- Outbound = On

Option 2 - DNS SRV (2nd choice)

To use this option, there must be a DNS SRV record available for the DNS name of the VCS cluster that defines an equal weighting and priority for each cluster peer.

On each SIP endpoint configure the SIP Settings as:

- Server discovery = Manual
- Server Address = DNS name of the VCS cluster

If the endpoint supports DNS SRV, on startup the endpoint issues a DNS SRV request and receives a DNS SRV record back defining an equal weighting and priority for each cluster peer.

The endpoint then tries to register with a relevant cluster peer (having taken into account the priority / weightings). If that peer is not available, the endpoint will try and register to another listed peer at the same priority, or if all peers at that priority have been tried, a peer at the next lower priority.

This will be repeated until the endpoint can register with a VCS.

Connect a Cluster

The endpoint will continue to use the first VCS that it registered to for re-registrations and for calls. If it ever loses connection to its VCS, it will use the DNS SRV entry to find a new VCS to register to, starting at the highest priority.

DNS SRV cache timeout should be set to a fairly long time (e.g. 24 hours) to minimize DNS traffic.

Option 3 - DNS Round-Robin (3rd choice)

To use this option, there must be a DNS A-record available for the DNS name of the VCS cluster that supplies a round-robin list of IP addresses.

On each SIP endpoint configure the SIP Settings as:

- Server discovery = Manual
- Server Address = DNS name of the VCS cluster

If the endpoint does not support DNS SRV, on startup the endpoint will perform a DNS A-record lookup. The DNS server will have been configured to support round-robin DNS, with each of the cluster peer members defined in the round-robin list.

The endpoint will take the address given by the DNS lookup and will then try and register with the relevant cluster peer. If that is not available, then the endpoint will perform another DNS lookup and will try to connect to the new VCS peer that it is given. (The DNS server will have supplied the next cluster peer's IP address.)

This will be repeated until the endpoint can register with a VCS.

The endpoint will continue to use the first VCS that it registered to for re-registrations and for calls. If it ever loses connection to its VCS it will perform another DNS lookup to find a new VCS to register to (the DNS server providing a VCS in the round-robin sequence).

DNS cache timeout should be set to a fairly short time (e.g. 1 minute or less) so that if a VCS is not accessible the endpoint is quickly pointed at a different VCS.

Option 4 - Static IP (least preferred)

Use this option if the VCS cluster does not have a DNS name.

On each SIP endpoint configure the SIP Settings as:

- Server discovery = Manual
- Server Address = IP address of a VCS peer

On startup the endpoint will try and register with the VCS at the specified IP address. If that is not available, then the endpoint will continue trying at regular intervals.

This will be repeated until the endpoint can register with the VCS.

The endpoint will continue to use the first VCS that it registered to for re-registrations and for calls. If it ever loses connection then it will keep on trying to register to that VCS until it is accessible again.

H.323 Endpoints

The options below are listed in preference order for providing resilience of connectivity of endpoints to a cluster of VCSs where 1 or more VCS cluster peers become inaccessible. The choice of option will depend on what functionality the endpoint you are using supports.

Option 1 - DNS SRV (preferred)

To use this option, there must be a DNS SRV record available for the DNS name of the VCS cluster that defines an equal weighting and priority for each cluster peer.

On each H.323 endpoint, configure the Gatekeeper Settings as:

- Discovery = Manual
- IP Address = DNS name of the VCS cluster

Connect a Cluster

If the endpoint supports DNS SRV, on startup the endpoint issues a DNS SRV request and receives a DNS SRV record back defining an equal weighting and priority for each cluster peer.

The endpoint then tries to register with a relevant cluster peer (having taken into account the priority / weightings). If that peer is not available, the endpoint will try and register to another listed peer at the same priority, or if all peers at that priority have been tried, a peer at the next lower (higher numbered) priority.

This will be repeated until the endpoint can register with a VCS. On registering with the VCS, the VCS will respond with the H.323 “Alternate Gatekeepers” list containing the list of VCS cluster peer members.

The endpoint will continue to use the first VCS that it registered to for re-registrations and for calls. If it ever loses connection to its VCS then it will select an “Alternate Gatekeeper” from the list it was supplied with.

DNS SRV cache timeout should be set to a fairly long time (e.g. 24 hours) to minimize DNS traffic.

Option 2 - DNS Round-Robin (2nd choice)

To use this option, there must be a DNS A-record available for the DNS name of the VCS cluster that supplies a round-robin list of IP addresses.

On each H.323 endpoint configure the Gatekeeper Settings as:

- Discovery = Manual
- IP Address = DNS name of the VCS cluster

If the endpoint does not support DNS SRV, on startup the endpoint will perform a DNS A-record lookup. The DNS server will have been configured to support round-robin DNS, with each of the cluster peer members defined in the round-robin list.

The endpoint will take the address given by the DNS lookup and will then try and register with the relevant cluster peer. If that peer is not available, then the endpoint will perform another DNS lookup and will try to connect to the new VCS peer that it is given. (The DNS server will have supplied the next cluster peer’s IP address.)

This will be repeated until the endpoint can register with a VCS. On registering with the VCS, the Cisco will respond with the H.323 ‘Alternate Gatekeepers’ list containing the list of VCS cluster peer members.

The endpoint will continue to use the first VCS that it registered to for re-registrations and for calls. If it ever loses connection then it will select an “Alternate Gatekeeper” from the list it was supplied with.

DNS cache timeout should be set to a fairly short time (e.g. 1 minute or less) so that on failure to reach a VCS at startup, the endpoint is quickly pointed at a different VCS.

Option 3 - Static IP (least preferred)

Use this option if the VCS cluster does not have a DNS name.

On each H.323 endpoint configure the Gatekeeper Settings as:

- Discovery = Manual
- IP Address = IP address of a VCS peer

On startup the endpoint will try and register with the VCS at the specified IP address. If that is not available, then the endpoint will continue trying at regular intervals.

This will be repeated until the endpoint can register with the VCS. On registering with the VCS, the VCS will respond with the H.323 “Alternate Gatekeepers” list containing the list of VCS cluster peer members.

The endpoint will continue to use the first VCS that it registered to for re-registrations and for calls. If it ever loses connection then it will select an “Alternate Gatekeeper” from the list it was supplied with.

Add a VCS to Cisco TMS

On the VCS:

- Go to **System > SNMP** and ensure that:
 - SNMP mode** is set to *v3 plus Cisco TMS support* or *v2c*.
 - Community name** is set to *public*.


(If SNMP was previously disabled, an alarm may appear indicating the need for a restart. If so, restart the system via **Maintenance > Restart options**.)

If multiple peers need restarting: restart each peer in turn, waiting for the peer to be accessible through the web interface before restarting the next peer.
- Go to **System > External manager** and ensure that:
 - Address** is set to the IP address or FQDN of Cisco TMS.
 - Path** is set to *tms/public/external/management/SystemManagementService.asmx*.
 - If the **Protocol** is *HTTPS* and **Certificate verification mode** is *On* then you must load the relevant certificates before the connection can become 'Active'. See *Implementing Secure Management* (document reference D50520) for details.
(If the **Protocol** is *HTTP* or **Certificate verification mode** is *Off*, no certificates need to be loaded.)
- Click **Save**.

The **Status** section of the **External manager** page should show a **State** of 'Active' or 'Initialising'.¹

On Cisco TMS:

- Select **Systems > Navigator**.
- Select (or create) an appropriate folder in which to put the VCS (in the example below the folder has been called "Cluster"):



- Click **Add Systems**.
 - In section 1. **Specify Systems by IP addresses or DNS names**, enter the IP address or DNS name of the VCS.
 - Click **Next**.
 - Look for ✓ System added.
If an error occurs, such as "Wrong Password", click on the **Edit System** link and correct the problem (enter the root password of the VCS).
- Note:** When you add an Expressway to TMS, the TMS UI shows it as a VCS. This is a known issue.
- Click **Finish Adding Systems**, **Add System despite warnings** or **Add More Systems** as appropriate.

¹Cisco TMS may force the protocol to be HTTPS. The configuration for this is found in **Administrative Tools > Configuration > Network settings**. The protocol will be forced to HTTPS if, in the **TMS Services** section **Enforce Management Settings on Systems** is set to *On* and in the **Secure-Only Device Communication** section **Secure-Only Device Communication** is set to *On*.

Connect a Cluster

8. If 'Could not connect to system. Details: No SNMP response' is reported, go to the **Connection** tab and type 'public' into the SNMP Get Community Name and select **Save/Try**.
9. If the VCS password is not default, set this up in the **Connection** tab or in **Settings > Edit Settings**.

If the VCS was already configured in Cisco TMS, ensure that it has the correct IP address (in Cisco TMS, go to **Systems > Navigator** (and any required sub folders), select the VCS, and from the **Connection** tab check the **IP Address** field).

On VCS:

Go to **System > External manager** and check that the **State** now shows **Active**.

Upgrade a Cluster

Upgrade an X7.1 (or Later) Cluster to X8.9.1

This procedure describes how to upgrade an existing X7.1 or later cluster to an X8.9.1 cluster.

When you select an upgrade window:

- Choose a period of low activity
- Allow time to upgrade all peers in the same window. The cluster will not reform correctly until the software versions match on all peers

Prerequisites

If you are using Cisco TMS for device provisioning or FindMe, this procedure assumes that:

- Cisco TMS is already running version 13.2 or later software and is operational with the X7.1 or later cluster.
- The VCS cluster and Cisco TMS are configured and operating using Cisco TMS Provisioning Extension mode. If Provisioning Extension mode is not already operational, upgrade to Provisioning Extension mode, using the steps in *Cisco TelePresence Management Suite Provisioning Extension Deployment Guide* before performing this upgrade.

For each VCS peer (including the primary), check the **Alarms** page (**Status > Alarms**) and ensure that all alarms are acted upon and cleared.

Upgrade VCS Cluster Peers to X8.9.1

Note: To avoid the risk of configuration data being lost, in addition to maintaining service continuity, it is important that you follow the upgrade instructions in the order specified here. In summary, we recommend that you:

1. Upgrade and restart the primary peer first.
2. Upgrade and restart the subordinate peers one at a time.

Full instructions are provided below.

Upgrade the Primary Peer

1. Log in to the primary peer as admin on the web interface.
2. Backup the VCS (**Maintenance > Backup and restore**).
You should backup your system before upgrading. In case you need to restore the configuration for any reason.
3. Enable maintenance mode:
 - a. Go to **Maintenance > Maintenance mode**.
 - b. Set **Maintenance mode** to *On*.
 - c. Click **Save** and click **OK** on the confirmation dialog.

This will stop the primary peer processing any new calls (existing calls will be dropped only if a restart occurs; other cluster peers will continue processing calls)

Upgrade a Cluster

4. Wait for all calls to clear and registrations to timeout on this peer.
 - If necessary, manually remove any calls on this peer that do not clear automatically (using the web browser go to **Status > Calls**, then select the check box next to the calls you want to terminate and click **Disconnect**).
 - If necessary, manually remove any registrations from this peer that do not clear automatically (using the web browser go to **Status > Registrations > By device**, then select the check box next to the devices you want to remove and click **Unregister**).

You can leave the registration for the Conference Factory – this will not be the source of calls, and even if deleted will not roll over to another peer, as other peers have their own Conference Factory registration (if enabled).

5. Upgrade and restart the primary VCS (**Maintenance > Upgrade**).

For any further details see the "Upgrading Software" section of *VCS Administrator Guide*.

Note that the web browser interface may timeout during the restart process, after the progress bar has reached the end. This may occur if VCS carries out a disk file system check – which it does approximately once every 30 restarts, or if it has not checked its disks in the last 6 months.

Upgrading of the software on the primary VCS is now complete.

Cluster-related alarms

You can ignore any cluster-related alarms that occur during the upgrade process, for example:

- "Cluster communication failure" alarm on the primary or any subordinate peers – this is expected.
- "Cluster replication error: cannot find primary or this subordinate's peer configuration file, manual synchronization of configuration is required"

Similarly you can ignore any Cisco TMS warnings and emails about cluster problems created during this upgrade. These alarms and warnings are expected and will clear when all cluster peers are upgraded and cluster data synchronization has occurred (they should clear within 10 minutes of the complete upgrade).

On the Primary VCS Peer:

1. Check that configuration data exists as expected:
 - If FindMe is in use, check that the expected FindMe entries still exist (**Status > Applications > TMS Provisioning Extension Services > FindMe > Accounts**).
 - Check configuration for items from the **System, Configuration** and **Application** menus.
2. Backup the VCS (**Maintenance > Backup and restore**).

Note: While VCS peers are running different versions of code, do not make any changes except what you need to do to upgrade. The cluster does not replicate changes to any peers that are on different versions to the primary VCS.

Upgrade Subordinate Peers

You can ignore any cluster-related alarms that occur during the upgrade process – these are expected.

1. Log in to the subordinate peer as admin on the web interface.
2. Backup the VCS (**Maintenance > Backup and restore**).

You should backup your system before upgrading. In case you need to restore the configuration for any reason.

Upgrade a Cluster

3. Enable maintenance mode:
 - a. Go to **Maintenance > Maintenance mode**.
 - b. Set **Maintenance mode** to *On*.
 - c. Click **Save** and click **OK** on the confirmation dialog.
4. Wait for all calls to clear and registrations to timeout on this peer.
 - If necessary, manually remove any calls on this peer that do not clear automatically (using the web browser go to **Status > Calls**, then select the check box next to the calls you want to terminate and click **Disconnect**).
 - If necessary, manually remove any registrations from this peer that do not clear automatically (using the web browser go to **Status > Registrations > By device**, then select the check box next to the devices you want to remove and click **Unregister**).

You can leave the registration for the Conference Factory – this will not be the source of calls, and even if deleted will not roll over to another peer, as other peers have their own Conference Factory registration (if enabled).
5. Upgrade and restart the VCS (**Maintenance > Upgrade**).

If multiple peers need restarting: restart each peer in turn, waiting for the peer to be accessible through the web interface before restarting the next peer.

For any further details see the "Upgrading Software" section of *VCS Administrator Guide*.

Note that the web browser interface may timeout during the restart process, after the progress bar has reached the end. This may occur if VCS carries out a disk file system check – which it does approximately once every 30 restarts, or if it has not checked its disks in the last 6 months.

Upgrading the software on this subordinate VCS peer is now complete.

Cluster-related alarms

You can ignore any cluster-related alarms that occur during the upgrade process, for example:

- "Cluster communication failure" alarm on the primary or any subordinate peers – this is expected.
- "Cluster replication error: cannot find primary or this subordinate's peer configuration file, manual synchronization of configuration is required"

Similarly you can ignore any Cisco TMS warnings and emails about cluster problems created during this upgrade. These alarms and warnings are expected and will clear when all cluster peers are upgraded and cluster data synchronization has occurred (they should clear within 10 minutes of the complete upgrade).

On this subordinate VCS peer:

1. Go to **Status > Alarms** and ensure that all alarms are acted upon and cleared.
2. Check that configuration data exists as expected:
 - If FindMe is in use, check that the expected FindMe entries still exist (**Status > Applications > TMS Provisioning Extension Services > FindMe > Accounts**).
 - Check configuration for items from the **System**, **Configuration** and **Application** menus.
3. Repeat these steps for each subordinate peer.

Check Cluster Status

Do this after you have completed upgrading all peers .

On each VCS (including the primary), go to **System > Clustering** and check that the cluster database status reports as Active.

Upgrading an X7.1 or later cluster to an X8.9.1 cluster is now complete.

Change a Cluster

Remove a Live Peer From a Cluster (Permanently)

This process removes one VCS peer from an existing cluster. FindMe and configuration replication to this VCS will be stopped and the VCS will no longer be included in the list of peers in the cluster. Provisioning will also be disabled on the removed VCS.

- If the whole cluster is to be disbanded then use the [Disband a Cluster, page 33](#) procedure instead.
- If the cluster peer to be removed is not accessible, use the procedure defined in [Remove a Dead Peer From a Cluster \(Permanently\), page 31](#).

Before starting:

1. Ensure that the VCS to be removed from the cluster is not indicated as the primary peer. If it is the primary, see [Change the Primary Peer, page 35](#) for instructions.


On the VCS that is being removed:

1. Log into the web interface.
2. Go to **System > Clustering**:
 - a. Change the **Cluster name** to a unique ID for this VCS (ideally to the routable Fully Qualified Domain Name used in SRV records that address this individual VCS, for example "vcse1.example.com". See [Appendix 2: Cluster Name and DNS SRV Records, page 42.](#))
 - b. Delete all entries in the **Peer N address** fields.

3. Click **Save**.

Note that:

- All previous FindMe users will still be available on this removed VCS. Remove them manually (**Users > FindMe accounts**) if required.
- An alarm similar to that shown below may appear on the web interface and CLI of the VCS being removed. This is not a problem, the alarm will be cleared when the VCS is restarted:



4. If Microsoft Interoperability was used in the cluster:
 - a. Go to **Applications > B2BUA > Microsoft Interoperability > Configuration**.
 - b. If **Microsoft Interoperability** is *Enabled*, change it to *Disabled* and click **Save**.
5. Restart the VCS (go to **Maintenance > Restart options**, then click **Restart** and confirm **OK**).

On the primary VCS:

1. Go to **System > Clustering**.
2. Delete the IP address of the VCS that has been removed.
3. If the VCS being removed is not the last field in the list, move any other addresses up the list so that there are no empty fields between entries.
4. If the primary VCS peer's address has been moved up the list in the previous step, alter the **Configuration primary** value to match its new location.
5. Click **Save**.

Change a Cluster

On all the remaining subordinate VCS peers:

1. Go to **System > Clustering**.
2. Edit the **Peer N address** and **Configuration primary** fields so that they are identical to those configured on the primary VCS.
3. Click **Save**.
4. Repeat for all remaining subordinate VCS peers until they all have identical clustering configuration.

On other devices:

1. If you have any other systems neighbored (or connected via a traversal zone) to this cluster of VCS peers, ensure that their zone configuration for this cluster is updated to exclude the address of the removed peer.
2. If you have any endpoints registering to the VCS that has now been removed, change the configuration of the endpoint (or the configuration of the DNS server entry that points to the cluster peers) so that they register to one of the remaining clustered VCS peers instead.

Caution: The removed VCS will retain its configuration at the time it is removed from the cluster, and will continue to function as a non-clustered VCS. After it has been removed from the cluster, we recommend that it is taken out of service (e.g. perform a factory reset as documented in [Remove a Dead Peer From a Cluster \(Permanently\)](#), page 31 before reconnecting the out-of-service VCS back to the network) or an alternative configuration is applied to the VCS, so that other devices no longer try to use it as a cluster peer.

Removing a live VCS from a cluster is now complete.

Remove a Dead Peer From a Cluster (Permanently)

Use the following procedure if:

- the VCS is dead and needs to be RMA'd, or
- the VCS cannot be accessed for some other reason

If the whole cluster is to be disbanded then use the procedure defined in [Disband a Cluster, page 33](#).

If the cluster peer to be removed is accessible, use the procedure defined in [Remove a Live Peer From a Cluster \(Permanently\), page 29](#) which clears up the removed VCS as well as its previous peers.

Note: This procedure does not delete clustering configuration from the removed VCS. When removed, you must not reconnect the out-of-service VCS without first deleting all of its peers and stopping FindMe and configuration replication (see [Clear Configuration From This Peer, page 31](#) below).

Before starting:

Ensure that the VCS to be removed from the cluster is not indicated as the primary VCS on Cisco TMS.

If it is the primary VCS, see [Change the Primary Peer, page 35](#) for instructions on how to make a different peer the primary.

On the primary VCS:

1. Go to **System > Clustering**.
2. Delete the IP address of the VCS that has been removed.
3. If the VCS being removed is not the last field in the list, move any other addresses up the list so that there are no empty fields between entries.
4. If the primary VCS peer's address has been moved up the list in the previous step, alter the **Configuration primary** value to match its new location.
5. Click **Save**.

On all the remaining subordinate VCS peers:

1. Go to **System > Clustering**.
2. Edit the **Peer N address** and **Configuration primary** fields so that they are identical to those configured on the primary VCS.
3. Click **Save**.
4. Repeat for all remaining subordinate VCS peers until they all have identical clustering configuration.

On other devices:

1. If you have any other systems neighbored (or connected via a traversal zone) to this cluster of VCS peers, ensure that their zone configuration for this cluster is updated to exclude the address of the removed peer.
2. If you have any endpoints registering to the VCS that has now been removed, change the configuration of the endpoint (or the configuration of the DNS server entry that points to the cluster peers) so that they register to one of the remaining clustered VCS peers instead.

Removing an out-of-service VCS from a cluster is now complete.

Clear Configuration From This Peer

If you ever recover the peer that you removed, you must clear its configuration before you reconnect it to the network. The easiest way to do this is to perform a factory reset.

Change a Cluster

The following procedure must be performed from the serial console (or via a direct connection to the appliance with a keyboard and monitor). This is because the network settings will be rewritten, so any SSH session used to initiate the reset would be dropped and the output of the procedure would not be seen.

The process takes approximately 20 minutes.

1. Log in to the system as **root**.
2. Type `factory-reset`
3. Answer the questions as required:

The recommended responses will reset the system completely to a factory default state.

Prompt	Recommended response
Keep option keys [YES/NO]?	NO
Keep FIPS140 configuration [YES/NO]?	NO
Keep IP configuration [YES/NO]?	NO
Keep ssh keys [YES/NO]?	NO
Keep ssl certificates and keys [YES/NO]?	NO
Keep root and admin passwords [YES/NO]?	NO
Save log files [YES/NO]?	NO

4. Confirm that you want to proceed.
5. After the serial boots, you will be taken to the Install Wizard. Some of the questions in the wizard may be skipped depending on your responses in step 3.
6. Once you complete the Install Wizard the system will apply the configuration and reboot. You can now log in using the admin password you set.

This ensures that the configuration of the recovered VCS is returned to default and it will not interact with its ex-peers.

Disband a Cluster

This process removes all VCS peers from an existing cluster. FindMe and configuration replication will be stopped, as will provisioning, and the cluster will be deleted from Cisco TMS.

Each VCS will retain its configuration at the time the cluster was disbanded, and will function as a stand-alone VCS.

Caution: If any of the VCSs are left in operation after being removed from the cluster, calls between endpoints registered to different VCSs that were once part of the same cluster will not succeed. This is because the disbanded cluster has no link between the two VCSs over which to route calls. If you want to restore calling between two disbanded peers, create neighbor relationships between them.

Before you start:

If any VCS is not accessible, firstly remove it using the procedure [Remove a Dead Peer From a Cluster \(Permanently\)](#), page 31.

If Cisco TMS is used (in Provisioning Extension mode):


1. Select **Systems > Navigator** (and any required sub folders), then click on any VCS in the cluster.
2. Select the **Provisioning** tab.
3. Disable all 4 services (clear checkboxes).
4. Click **Save**.

On each subordinate VCS peer:

1. Log into the web interface.
2. Go to **System > Clustering**:
 - a. Change the **Cluster name** to a unique ID for this VCS (ideally to the routable Fully Qualified Domain Name used in SRV records that address this individual VCS, for example "vcse1.example.com". See [Appendix 2: Cluster Name and DNS SRV Records](#), page 42.)
 - b. Delete all entries in the **Peer N address** fields.
3. Click **Save**.

Note that:

- All previous FindMe users will still be available on this removed VCS. Remove them manually (**Users > FindMe accounts**) if required.
- An alarm similar to that shown below may appear on the web interface and CLI of the VCS being removed. This is not a problem, the alarm will be cleared when the VCS is restarted:



4. If Microsoft Interoperability was used in the cluster:
 - a. Go to **Applications > B2BUA > Microsoft Interoperability > Configuration**.
 - b. If **Microsoft Interoperability** is *Enabled*, change it to *Disabled* and click **Save**.
5. Restart the VCS (go to **Maintenance > Restart options**, then click **Restart** and confirm **OK**).
If multiple peers need restarting: restart each peer in turn, waiting for the peer to be accessible through the web interface before restarting the next peer.
6. Repeat for each subordinate VCS.


Change a Cluster

On the primary VCS:

1. Log into the web interface.
2. Go to **System > Clustering**:
 - a. Optionally, change the **Cluster name** to a unique ID for this VCS (ideally to the routable Fully Qualified Domain Name used in SRV records that address this individual VCS, for example "vcse1.example.com". (See [Appendix 2: Cluster Name and DNS SRV Records, page 42.](#))
 - b. Delete all entries in the **Peer x IP address** fields.
3. Click **Save**.

Note that:

- All previous FindMe users will still be available on this removed VCS
- An alarm similar to that shown below may appear on the web interface and CLI of the VCS being removed. This is not a problem, the alarm will be cleared when the VCS is restarted:



4. Restart the VCS (go to **Maintenance > Restart options**, then click **Restart** and confirm **OK**).

On other devices:

If you have any other systems neighbored (or connected via a traversal zone) to this cluster of VCS peers, ensure that they have those zones removed, or modified appropriately.

If you have any endpoints registering to this VCS cluster, change the configuration of the endpoints (or the configuration of the DNS server entry that points to the cluster peers) so that they now register with an appropriate VCS.

Disband a cluster of VCSs is now complete.

Change the Primary Peer

You can do this procedure even if the current primary peer is not accessible.

Note: Follow this process in the order listed, to avoid putting the cluster in a state where multiple peers are contending to be the primary.

On the "new" primary VCS:

1. Go to **System > Clustering**.
2. From the **Configuration primary** drop-down menu select the ID number of the peer entry that says 'This system'.
3. Click **Save**.

While changing the primary peer, ignore any alarms on VCS that report 'Cluster primary mismatch' or 'Cluster replication error' – they will be rectified as part of this procedure.

On all other VCS peers, starting with the "old" primary peer (if it is still accessible):

1. Go to **System > Clustering**.
2. From the **Configuration primary** drop-down menu select the ID number of the "new" primary VCS.
3. Click **Save**.

On each VCS in the cluster (including the primary):

1. Confirm that the change to the **Configuration primary** has been accepted by going to **System > Clustering** and refreshing the page.
2. If any VCSs have not accepted the change, repeat the steps above.
3. Check that the cluster database status reports as Active.

Note that any alarms raised on the VCS peers that relate to 'Cluster primary mismatch' and 'Cluster replication error' should clear automatically after approximately 2 minutes.

Cisco TMS

No changes are required; Cisco TMS will see the primary change on the VCS cluster and report this appropriately.

If the old primary is not available

If you are changing the primary peer because the "old" primary is not accessible, see [Remove a Dead Peer From a Cluster \(Permanently\)](#), page 31 procedure.

If the "old" primary was not accessible when you changed the cluster's primary peer, but later that VCS becomes available, you must clear its configuration and then bring it back into the cluster using the [Add a Peer to a Cluster](#), page 13 procedure.

Changing the primary peer of a VCS cluster is now complete.

Change the Address of a Peer

To change the address of a VCS peer you must remove the VCS from the cluster, change its address, and then add the VCS back into the cluster.

The process is as follows:

1. Ensure that the VCS whose address you want to change is not the primary VCS.
If it is the primary VCS, follow the steps in [Change the Primary Peer, page 35](#) to make a different peer the primary.
2. Carry out the process documented in [Remove a Live Peer From a Cluster \(Permanently\), page 29](#).
3. Change the address of the VCS.
4. Carry out the process documented in [Add a Peer to a Cluster, page 13](#).

Replace a Peer

This section summarizes the procedure for replacing a cluster peer VCS with a different unit.

1. Ensure that the VCS to be replaced is not the primary VCS.
If it is the primary VCS, follow the steps in [Change the Primary Peer, page 35](#) to make a different peer the primary.
2. Remove the existing peer from the cluster:
 - a. If the cluster peer to be replaced is not accessible, use the procedure defined in [Remove a Dead Peer From a Cluster \(Permanently\), page 31](#).
 - b. If the cluster peer to be replaced is accessible, use the procedure defined in [Remove a Live Peer From a Cluster \(Permanently\), page 29](#).
3. Add the replacement peer to the cluster using the procedure defined in [Add a Peer to a Cluster, page 13](#).

Replace a Peer And Keep Its Configuration

This procedure assumes that you are replacing an accessible VCS peer with a different VCS.

1. Ensure that the VCS to be replaced is not the primary VCS.
If it is the primary VCS, follow the steps in [Change the Primary Peer, page 35](#) to make a different peer the primary
2. **Remove the peer safely**, as described in [Remove a Live Peer From a Cluster \(Permanently\), page 29](#).
You must clear the clustering configuration from the peer, and update the clustering configuration on all the other peers (primary first), before you take the backup
3. Backup the configuration of the removed peer
4. Generate and apply option keys for the new VCS. You must apply the same set of keys that are applied to the other peers
5. Restore the backup onto the new VCS
6. Check the DNS configuration of the new VCS is the same as the other peers, and synchronize it with the same NTP servers
7. Add the replacement peer to the cluster using the procedure defined in [Add a Peer to a Cluster, page 13](#)
You already have much of the configuration described in the linked procedure. The most important steps described there are summarized here:
 - a. Add the new peer's address to the clustering configuration on the primary
 - b. Add the new peer's address to the clustering configuration on other existing peers
 - c. Enter the clustering configuration on the new peer (cluster name, shared secret, ordered peer list)
8. Restart the new peer
9. Wait for approximately five minutes, then check the cluster status and resolve any alarms

Note: If you replaced a peer in a cluster of VCS Controls, and that cluster was configured for SSO of MRA clients, then SSO will fail some of the time until you update the IDP with the cluster's new SAML metadata.

This is because one of the peers has a new serial number, and the serial numbers of the peers are used to generate the cluster's metadata.

You must export the cluster's SAML metadata and copy it to the IDP.

See *Mobile and Remote Access Through Cisco Video Communication Server* on the [VCS configuration guides page](#).

Troubleshooting

VCS Alarms and Warnings

" Cluster name not configured: if FindMe or clustering are in use a cluster name must be defined"

Ensure that the same cluster name is configured on each VCS in the cluster.

The **Cluster name** should be the routable Fully Qualified Domain Name used in SRV records that address this VCS cluster, for example "cluster1.example.com". (See [Appendix 2: Cluster Name and DNS SRV Records](#), page 42).

" Cluster replication error: <details> manual synchronization of configuration is required"

This may be:

- " Cluster replication error: manual synchronization of configuration is required"
- " Cluster replication error: cannot find primary or this subordinate's peer configuration file, manual synchronization of configuration is required"
- " Cluster replication error: configuration primary ID is inconsistent, manual synchronization of configuration is required"
- " Cluster replication error: this peer's configuration conflicts with the primary's configuration, manual synchronization of configuration is required"

If a subordinate VCS reports an alarm: " Cluster replication error - <details> synchronization of configuration"

On that subordinate VCS:

1. Log in as admin on an SSH or other CLI interface.
2. At the command prompt type:

```
xcommand ForceConfigUpdate
```

This will delete the subordinate VCS configuration and then force it to update its configuration from the primary VCS.

Caution: Only use this command if the configuration on the primary VCS is known to be in a good state. We recommend that you take a backup before running this command.

" Cluster replication error: the NTP server is unreachable"

Configure an accessible NTP server on the VCS **System > Time** page.

" Cluster replication error: the local VCS does not appear in the list of peers"

Check and correct the list of peers for this VCS on the primary VCS, and copy to all other VCS peers (**System > Clustering**).

" Cluster replication error: automatic replication of configuration has been temporarily disabled because an upgrade is in progress"

Wait until the upgrade has completed.

" Invalid clustering configuration: H.323 mode must be turned On - clustering uses H.323 communications between peers"

Ensure that H.323 mode is on (see **Configuration > Protocols > H.323**).

"VCS database failure: Please contact your Cisco support representative"

The support representative will help you work through the following steps:

1. Take a system snapshot and provide it to your support representative.
2. Remove the VCS from the cluster using: [Remove a Live Peer From a Cluster \(Permanently\)](#), page 29.
3. Restore that VCS's database by restoring a backup taken on that VCS previously.
4. Add the VCS back to the cluster using [Add a Peer to a Cluster](#), page 13.

A second method is possible if the database does not recover:

1. Take a system snapshot and provide it to TAC.
2. Remove the VCS from the cluster using: [Remove a Live Peer From a Cluster \(Permanently\)](#), page 29.
3. Log in as root and run `clusterdb_destroy_and_purge_data.sh`
4. Restore that VCS's database by restoring a backup taken on that VCS previously.
5. Add the VCS back to the cluster using [Add a Peer to a Cluster](#), page 13.

Note: `clusterdb_destroy_and_purge_data.sh` is as dangerous as it sounds – only use this command in conjunction with instructions from your support representative.

Cisco TMS Warnings

Cisco TMS Cluster Diagnostics

If Cisco TMS cluster diagnostics reports a difference in configuration on VCS peers, it is comparing the output of `https://<ip address>/alternatesconfiguration.xml` for each VCS.

To manually check the differences, on a Unix / Linux system, run:

```
wget --user=admin --password=<password> --no-check-certificate https://<IP or FQDN of VCS>/alternatesconfiguration.xml
```

for each of the VCS peers, then use `diff` to check for differences.

Conference Factory Template Does Not Replicate

This is by design; the Conference Factory %% value is NOT shared between cluster peers and the Conference Factory application configuration is NOT replicated across a cluster.

See [Appendix 5: Impact of Clustering on Other VCS Applications](#), page 48.

VCS's External Manager Protocol Keeps Getting Set to HTTPS

Cisco TMS can be configured to force specific management settings on connected systems. This includes ensuring that a VCS uses HTTPS for feedback. If enabled, Cisco TMS will (on a time period defined by Cisco TMS) re-configure the VCS's **System > External manager Protocol** to *HTTPS*.

If HTTPS must be used for VCS to supply feedback to Cisco TMS, see [Add a VCS to Cisco TMS](#), page 24 for information about how to set up certificates.

Cisco TMS will force HTTPS on VCS if:

- **TMS Services > Enforce Management Settings on Systems = On (Administrative Tools > Configuration > Network Settings)**
and

Troubleshooting

- **Secure-Only Device Communication > Secure-Only Device Communication = On (Administrative Tools > Configuration > Network Settings)**

Set **Enforce Management Settings on Systems** to *Off* if Cisco TMS does not need to force the management settings.

Set **Secure-Only Device Communication** to *Off* if it is unnecessary for VCS to provide feedback to Cisco TMS using HTTPS (if HTTP is sufficient).

Reference

Appendix 1: IP Ports and Protocols

External Firewalls Between Peers

It is unusual to have any sort of external firewall between cluster peers, but if there is, the IP protocols and ports that must be open between each and every VCS peer in the cluster are listed below.

For cluster communications between VCS peers:

- TCP port 4371 is used for cluster recovery (over TLS)
- TCP port 4372 is used for cluster database synchronization (over TLS)

For calls between VCS peers:

- Standard SIP and H.323 signaling ports are used for calls
- UDP port 1719 is used for bandwidth updates between VCS peers

Firewall Rules on the Peers

If you are using the VCS's built-in **Firewall rules** feature, make sure that your rules allow the following connections:

Table 2 Clustering Connections

Purpose	Protocol	Source	Port	Destination	Port
Cluster communication	TCP/TLS	Other peers	Ephemeral	This peer	4372
Cluster recovery	TCP/TLS	Other peers	Ephemeral	This peer	4371

Cisco TelePresence Management Suite Provisioning Extension

For cluster communications between VCS peers and a Cisco TMS when running in Provisioning Extension mode:

- VCS ephemeral port to port 443 on Cisco TMS (secure) or
- VCS ephemeral port to port 80 on Cisco TMS

Note: Ports 443 and 80 are the default values; they can be configured in the Cisco TMS IIS, and VCS if different ports are required.

Appendix 2: Cluster Name and DNS SRV Records

Using DNS SRV to convert a domain to an IP address has a number of benefits:

- The structure of the lookup includes service type and protocol as well as the domain, so that a common domain can be used to reference multiple different services which are hosted on different machines (e.g. HTTP, SIP, H.323).
- The DNS SRV response includes priority and weighting values which allow the specification of primary, secondary, tertiary etc groups of servers, and within each priority group, the weighting defines the proportion of accesses that should use each server.
- As the DNS SRV response contains details about priorities and weights of multiple servers, the receiving device can use a single lookup to search for an in-service server (where some servers are inaccessible) without the need to repeatedly query the DNS server. (This is in contrast to using round robin DNS which does require repeated lookups into the DNS server if initial servers are found to be inaccessible.)

The generic format of a DNS SRV query is:

- `_service._protocol.<fully.qualified.domain>`

The DNS SRV response is a set of records in the format:

- `_service._protocol.<fully.qualified.domain>. TTL Class SRV Priority Weight Port Target`
where Target is an A-record defining the destination.

Further details on DNS SRV can be found in *VCS Administrator Guide* and *RFC 2782*.

DNS SRV Configuration for Mobile and Remote Access

This section summarizes the public (external) and local (internal) DNS requirements. For more information, see the *Cisco Jabber Planning Guide* (for your version) on the [Jabber Install and Upgrade Guides page](#).

Public DNS

The public (external) DNS must be configured with `_collab-edge._tls.<domain>` SRV records so that endpoints can discover the VCS Expressways to use for Mobile and Remote Access. SIP service records are also required (for general deployment, not specifically for Mobile and Remote Access). For example, for a cluster of 2 VCS Expressway systems:

Domain	Service	Protocol	Priority	Weight	Port	Target host
example.com	collab-edge	tls	10	10	8443	vcse1.example.com
example.com	collab-edge	tls	10	10	8443	vcse2.example.com
example.com	sips	tcp	10	10	5061	vcse1.example.com
example.com	sips	tcp	10	10	5061	vcse2.example.com

Local DNS

The local (internal) DNS requires `_cisco-uds._tcp.<domain>` SRV records. For example:

Domain	Service	Protocol	Priority	Weight	Port	Target host
example.com	cisco-uds	tcp	10	10	8443	cucmserver1.example.com
example.com	cisco-uds	tcp	10	10	8443	cucmserver2.example.com

Reference

Notes:

- **Important! From version X8.8 onward, you must create forward and reverse DNS entries for all VCS Expressway systems, so that systems making TLS connections to them can resolve their FQDNs and validate their certificates.**
- Ensure that the `cisco-uds` SRV records are NOT resolvable outside of the internal network, otherwise the Jabber client will not start Mobile and Remote Access negotiation via the VCS Expressway.
- You must create internal DNS records, for both forward and reverse lookups, for all Unified Communications nodes used with Mobile and Remote Access. This allows VCS Control to find the nodes when IP addresses or hostnames are used instead of FQDNs.

DNS SRV Configuration for Video Conferencing

The format of DNS SRV queries for sip (RFC 3263) and H.323 used by VCS are:

- `_sips._tcp.<fully.qualified.domain>`
- `_sip._tcp.<fully.qualified.domain>`
- `_sip._udp.<fully.qualified.domain>` - not recommended for video calls, only use for audio-only calls
- `_h323ls._udp.<fully.qualified.domain>` - for UDP location (RAS) signaling, such as LRQ
- `_h323cs._tcp.<fully.qualified.domain>` - for H.323 call signaling

The format of DNS SRV queries for sip (RFC 3263) and H.323 typically used by an endpoint are:

- `_sips._tcp.<fully.qualified.domain>`
- `_sip._tcp.<fully.qualified.domain>`
- `_sip._udp.<fully.qualified.domain>` - not recommended for video calls, only use for audio-only calls
- `_h323ls._udp.<fully.qualified.domain>` - for UDP location (RAS) signaling, such as LRQ
- `_h323cs._tcp.<fully.qualified.domain>` - for H.323 call signaling
- `_h323rs._udp.<fully.qualified.domain>` - for H.323 registrations

Note that UDP is not a recommended transport medium for video signaling; SIP messaging for video system is too large to be reliably carried on datagram-based (rather than stream-based) transports.

The VCS **Cluster name** (configured on the **System > Clustering** page) should be the `<fully.qualified.domain>` specified in the DNS SRV records that point to the VCS cluster.

Example

DNS SRV records for 2 peers of a VCS-E cluster for `example.com`

where:

- FQDN of VCS Expressway peer 1: `vcse1.example.com`
- FQDN of VCS Expressway peer 2: `vcse2.example.com`
- FQDN of VCS Expressway cluster: `example.com`

```
_sips._tcp.example.com. 86400 IN SRV 1 1 5061 vcse1.example.com.
_sips._tcp.example.com. 86400 IN SRV 1 1 5061 vcse2.example.com.

_sip._tcp.example.com. 86400 IN SRV 1 1 5060 vcse1.example.com.
_sip._tcp.example.com. 86400 IN SRV 1 1 5060 vcse2.example.com.

_h323ls._udp.example.com. 86400 IN SRV 1 1 1719 vcse1.example.com.
_h323ls._udp.example.com. 86400 IN SRV 1 1 1719 vcse2.example.com.

_h323cs._tcp.example.com. 86400 IN SRV 1 1 1720 vcse1.example.com.
_h323cs._tcp.example.com. 86400 IN SRV 1 1 1720 vcse2.example.com.
```

Reference

```
_h323rs._udp.example.com. 86400 IN SRV 1 1 1719 vcse1.example.com.
_h323rs._udp.example.com. 86400 IN SRV 1 1 1719 vcse2.example.com.
```

Note that:

- Priorities are all the same. Only use different priorities if you have different clusters allowing failover from one primary cluster to another (secondary) cluster. In that case the primary cluster's peers should have one value and the other (secondary) cluster's peers a (larger) value.
- Weights should be the same – so that there is equal use of each peer.

Checking DNS SRV Settings

Checking DNS SRV settings via web interface

1. Go to **Maintenance > Tools > Network utilities > DNS lookup**.
2. Enter the SRV path in the **Host** field.
3. Click **Lookup**.

nslookup

```
nslookup -query=SRV _sip._tcp.example.com
```

Dig

```
dig _sip._tcp.example.com SRV

; <<>> DiG 9.4.1 <<>> _sip._tcp.example.com SRV
;; global options: printcmd
;; Got answer:
;; ->>HEADER<<- opcode: QUERY, status: NOERROR, id: 44952
;; flags: qr rd ra; QUERY: 1, ANSWER: 2, AUTHORITY: 2, ADDITIONAL: 4

;; QUESTION SECTION:
;_sip._tcp.example.com.      IN      SRV

;; ANSWER SECTION:
_sip._tcp.example.com. 1183    IN      SRV     1 0 5060 vcse1.example.com.
_sip._tcp.example.com. 1183    IN      SRV     1 0 5060 vcse2.example.com.

;; AUTHORITY SECTION:
example.com.           87450   IN      NS       ns1.mydyndns.org.
example.com.           87450   IN      NS       ns2.mydyndns.org.

;; ADDITIONAL SECTION:
vcse1.example.com.    1536    IN      A        194.73.59.53
vcse2.example.com.    1376    IN      A        194.73.59.54
ns1.mydyndns.org.     75      IN      A        204.13.248.76
ns2.mydyndns.org.    10037   IN      A        204.13.249.76

;; Query time: 0 msec

~ #
```

Appendix 3: Clusters in Isolated Networks

As of X8.8, VCS peers use TLS to communicate with each other. You have the option of permissive TLS - the certificates are not verified - or enforced TLS where the certificates are verified.

In the latter case, each peer will need to DNS look up the common name (CN), and perhaps also subject alternate names (SANs), that they read from their peers' certificates. They compare the returned IP addresses against the IP addresses that gave them the certificates and if they match, the connection is authenticated.

In isolated networks, the peers will not typically be able to reach the internal DNS servers, because that would require unsolicited inbound requests. In a dual-NIC setup, you probably also don't want to put the peers' private IP addresses into the public DNS.

The issue is compounded by not being able to use IP addresses as common names or subject alternate names on server certificates: certificate authorities do not advocate this and probably will not issue such certificates.

The workaround for this issue depends on how your peers are configured. The following options presume that you are forming a cluster of VCS Expressway peers in the same subnet:

VCS Expressway peers have dual NICs, with no static NAT

You can enforce TLS between cluster peers:

1. Enter public DNS servers on the DNS configuration of each peer.
2. Choose which of the LAN interfaces will take the public facing address.
3. Configure the public DNS to resolve each peer's FQDN to its public IP address.
4. Populate the CN of all peer certificates with the same cluster FQDN, and populating each peer certificate's SAN with that peer's FQDN.
5. Enter the cluster FQDN and peer FQDNs on the clustering configuration page and set the **TLS verification mode** to *Enforce*.

The peers will now use the public DNS to verify each others' identities, as presented on their certificates.

VCS Expressway peers have dual NICs, with static NAT enabled

In addition to its private IP address in the isolated network, you can give one of the NICs a public IP address that translates to its private address. In this case, you cannot use FQDNs to form the cluster.

This is because the public DNS record for each peer's FQDN would match its translated (public) IP address, but the peers would see each other's private addresses when swapping certificates. The mismatch between IP addresses would prevent the TLS connection being established, and the cluster would not form.

To form the cluster:

1. Enter public DNS servers on the DNS configuration of each peer.
2. Choose which of the LAN interfaces on each peer will have static NAT enabled.
3. Enter the private IP addresses of the other LAN interfaces on the clustering configuration pages, and set the TLS mode to *Permissive*.

The peers will now use the private IP addresses to form the cluster, but will not check the contents of the certificates against the DNS records.

Note: We intend to deliver a more secure solution to this issue in future, so we recommend creating valid certificates if you want to enforce TLS and use static NAT.

Appendix 4: NAPTR Records

NAPTR records are typically used to specify various methods to connect to a destination URI, for example by email, by SIP, by H.323. They can also be used to specify the priority to use for those connection types, for example to use SIP TLS in preference over using SIP TCP or SIP UDP.

NAPTR records are also used in ENUM, when converting a telephone number into a dialable URI. (For further details on ENUM see [ENUM Dialing on VCS Deployment Guide](#)).

NAPTR Record Format

Example: SIP access to example.com, and for enum lookups for 557120, 557121, and 557122.

\$ORIGIN example.com.

```
IN NAPTR 10 100 "s" "SIPS+D2T" "" _sips._tcp.example.com.
IN NAPTR 12 100 "s" "SIP+D2T" "" _sip._tcp.example.com.
IN NAPTR 14 100 "s" "SIP+D2U" "" _sip._udp.example.com.
```

\$ORIGIN www.example.com.

```
IN NAPTR 10 100 "s" "http+I2R" "" _http._tcp.example.com.
IN NAPTR 10 100 "s" "ftp+I2R" "" _ftp._tcp.example.com.
```

\$ORIGIN 0.2.1.7.5.5.enum.lookup.com.

```
IN NAPTR 10 100 "u" "E2U+sip" "!^.*$!john.smith@tandberg.com!" .
IN NAPTR 12 100 "u" "E2U+h323" "!^.*$!john.smith@tandberg.com!" .
IN NAPTR 10 100 "u" "mailto+E2U" "!^.*$!mailto:john.smith@tandberg.com!" .
```

\$ORIGIN 1.2.1.7.5.5.enum.lookup.com.

```
IN NAPTR 10 100 "u" "E2U+sip" "!^.*$!mary.jones@tandberg.com!" .
```

\$ORIGIN 2.2.1.7.5.5.enum.lookup.com.

```
IN NAPTR 10 100 "u" "E2U+h323" "!^.*$!peter.archibald@myco.com!" .
```

IN = Internet routing

NAPTR = record type

10 = order value (use lowest order value first)

100 = preference value if multiple entries have the same order value

"u" = the result is a routable URI

"s" = the result is a DNS SRV record

"a" = the result is an 'A' or 'AAAA' record

"E2U+sip" to make SIP call

"E2U+h323" to make h.323 call

Regular expression:

! = delimiter

"" = no expression used

... usual Regex expressions can be used

Replace field; . = not used

Looking up an ENUM NAPTR record

```
dig 4.3.7.8.enum4.example.com. NAPTR
```

```
; <<> ;; global options: printcmd
```

```
;; Got answer:
```

```
;; ->>HEADER<<- opcode: QUERY, status: NOERROR, id: 38428
```

```
;; flags: qr aa rd ra; QUERY: 1, ANSWER: 2, AUTHORITY: 1, ADDITIONAL: 2
```

```
;; QUESTION SECTION:
```

```
;4.3.7.8.enum4.example.com. IN NAPTR
```

```
;; ANSWER SECTION:
```

Reference

```

4.3.7.8.enum4.example.com. 60 IN NAPTR 10 100 "u" "E2U+sip" "!^.*$!bob@example.com!" .
4.3.7.8.enum4.example.com. 60 IN NAPTR 10 100 "u" "E2U+h323" "!^.*$!bob@example.com!" .

;; AUTHORITY SECTION:
enum4.example.com. 60      IN      NS      int-server1.example.com.

;; ADDITIONAL SECTION:
int-server1.example.com. 3600 IN A      10.44.9.144
int-server1.example.com. 3600 IN AAAA   3ffe:80ee:3706::9:144

;; Query time: 0 msec

```

Looking up a domain NAPTR record

Example: NAPTR record allowing endpoints to detect that they are in the public (external) network. The flag “s” is extended to “se” to indicate that it is “external”.

```

~ # dig -t NAPTR example.com

; <<>> DiG 9.4.1 <<>> -t NAPTR example.com
;; global options: printcmd
;; Got answer:
;; ->>HEADER<<- opcode: QUERY, status: NOERROR, id: 1895
;; flags: qr rd ra; QUERY: 1, ANSWER: 3, AUTHORITY: 6, ADDITIONAL: 4

;; QUESTION SECTION:
;example.com.      IN      NAPTR

;; ANSWER SECTION:
example.com.      2      IN      NAPTR   50 50 "se" "SIPS+D2T" "" _sips._tcp.example.com.
example.com.      2      IN      NAPTR   90 50 "se" "SIP+D2T" "" _sip._tcp.example.com.
example.com.      2      IN      NAPTR  100 50 "se" "SIP+D2U" "" _sip._udp.example.com.

;; AUTHORITY SECTION:
example.com.      320069 IN      NS      nserver2.example.com.
example.com.      320069 IN      NS      nserver.euro.example.com.
example.com.      320069 IN      NS      nserver.example.com.
example.com.      320069 IN      NS      nserver3.example.com.
example.com.      320069 IN      NS      nserver4.example.com.
example.com.      320069 IN      NS      nserver.asia.example.com.

;; ADDITIONAL SECTION:
nserver.example.com. 56190 IN      A      17.111.10.50
nserver2.example.com. 57247 IN      A      17.111.10.59
nserver3.example.com. 57581 IN      A      17.22.14.50
nserver4.example.com. 57452 IN      A      17.22.14.59

;; Query time: 11 msec

```

Appendix 5: Impact of Clustering on Other VCS Applications

Conference Factory (Multiway™)

When using Conference Factory (Multiway) in a cluster, note that:

- The Conference Factory application configuration is NOT replicated across a cluster.
- The Conference Factory template MUST be DIFFERENT on each of the VCS peers.

When configuring a cluster to support Multiway:

1. Set up the **same** Conference Factory alias (the alias called by the endpoint to initiate a Multiway conference) on each peer.
2. Set up a **different** Conference Factory template on each peer (so that each peer generates unique Multiway conference IDs).

For example, if the MCU service prefix for ad hoc conferences is 775 then the primary VCS may have a template of 775001%%@domain, peer 2 a template of 775002%%@domain, and peer 3 a template of 775003%%@domain. In this way, whichever VCS serves the conference ID, it cannot serve a conference ID that any other VCS could have served.

The same applies across a network. If there is more than one VCS or VCS cluster that provides Conference Factory functionality in a network, each and every VCS must provide values in a unique range, so that no two VCSs can serve the same conference ID.

See [Cisco TelePresence Multiway Deployment Guide](#) for more information.

Microsoft Interoperability

If Microsoft infrastructure is deployed with a VCS cluster, see [VCS and Microsoft Infrastructure Deployment Guide](#).

Appendix 6: Upgrade Earlier VCS Versions to X8.9.1

This section describes how to upgrade earlier VCS versions to X8.9.1.

Upgrading an X7.1 or X7.2 Cluster in TMS Agent Legacy Mode to X8.9.1

As from X8, TMS Agent legacy mode is no longer supported. You cannot upgrade directly from a cluster running in TMS Agent legacy mode to an X8 or later cluster.

To upgrade to X8.9.1 you must:

1. Upgrade your X7.1 or X7.2 cluster to Provisioning Extension mode as described in *Cisco TelePresence Management Suite Provisioning Extension Deployment Guide*.
2. Upgrade to X8.9.1 as described in [Upgrade an X7.1 \(or Later\) Cluster to X8.9.1, page 26](#).

Upgrading an X6.0, X6.1, X7.0, or X7.0.n Cluster to X8.9.1

If you are using Cisco TMS for Device Provisioning or FindMe:

- As from X8, TMS Agent legacy mode is no longer supported.
- You cannot upgrade directly from a cluster running in TMS Agent legacy mode to an X8 or later cluster.

To upgrade to X8.9.1 you must:

1. Upgrade the cluster to X7.2 as described in *VCS Cluster Creation and Maintenance Deployment Guide (X7.2)*.
2. If you are using Cisco TMS for Device Provisioning or FindMe, upgrade to Provisioning Extension mode as described in *Cisco TelePresence Management Suite Provisioning Extension Deployment Guide*.
3. Upgrade to X8.9.1 as described in [Upgrade an X7.1 \(or Later\) Cluster to X8.9.1, page 26](#).

Upgrading an X5.1, X5.1.1 or X5.2 Cluster to X8.9.1

An upgrade from an X5.1 / X5.1.1 / X5.2 cluster directly to an X8.9.1 cluster is not possible. You must upgrade in stages through previous VCS versions:

To upgrade to X8.9.1 you must:

1. Upgrade the cluster to X6.1 as described in *VCS Cluster Creation and Maintenance Deployment Guide (X6.1)*.
2. Upgrade the cluster to X7.2 as described in *VCS Cluster Creation and Maintenance Deployment Guide (X7.2)*.
3. If you are using Cisco TMS for Device Provisioning or FindMe, upgrade to Provisioning Extension mode as described in *Cisco TelePresence Management Suite Provisioning Extension Deployment Guide*.
4. Upgrade to X8.9.1 as described in [Upgrade an X7.1 \(or Later\) Cluster to X8.9.1, page 26](#).

Upgrading X1 / X2 Alternates or an X3 / X4 / X5.0 Cluster to X8.9.1

An upgrade from X1 / X2 alternates or an X3 / X4 / X5.0 cluster directly to an X8.9.1 cluster is not possible. You must upgrade in stages through previous VCS versions:

To upgrade to X8.9.1 you must:

1. Upgrade the cluster to X5.2 as described in *VCS Cluster Creation and Maintenance Deployment Guide (X5)*.
2. Upgrade the cluster to X6.1 as described in *VCS Cluster Creation and Maintenance Deployment Guide (X6.1)*.
3. Upgrade the cluster to X7.2 as described in *VCS Cluster Creation and Maintenance Deployment Guide (X7.2)*.
4. If you are using Cisco TMS for Device Provisioning or FindMe, upgrade to Provisioning Extension mode as

Reference

described in *Cisco TelePresence Management Suite Provisioning Extension Deployment Guide*.

5. Upgrade to X8.9.1 as described in [Upgrade an X7.1 \(or Later\) Cluster to X8.9.1, page 26](#).

Cisco Legal Information

THE SPECIFICATIONS AND INFORMATION REGARDING THE PRODUCTS IN THIS MANUAL ARE SUBJECT TO CHANGE WITHOUT NOTICE. ALL STATEMENTS, INFORMATION, AND RECOMMENDATIONS IN THIS MANUAL ARE BELIEVED TO BE ACCURATE BUT ARE PRESENTED WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED. USERS MUST TAKE FULL RESPONSIBILITY FOR THEIR APPLICATION OF ANY PRODUCTS.

THE SOFTWARE LICENSE AND LIMITED WARRANTY FOR THE ACCOMPANYING PRODUCT ARE SET FORTH IN THE INFORMATION PACKET THAT SHIPPED WITH THE PRODUCT AND ARE INCORPORATED HEREIN BY THIS REFERENCE. IF YOU ARE UNABLE TO LOCATE THE SOFTWARE LICENSE OR LIMITED WARRANTY, CONTACT YOUR CISCO REPRESENTATIVE FOR A COPY.

The Cisco implementation of TCP header compression is an adaptation of a program developed by the University of California, Berkeley (UCB) as part of UCB's public domain version of the UNIX operating system. All rights reserved. Copyright © 1981, Regents of the University of California.

NOTWITHSTANDING ANY OTHER WARRANTY HEREIN, ALL DOCUMENT FILES AND SOFTWARE OF THESE SUPPLIERS ARE PROVIDED "AS IS" WITH ALL FAULTS. CISCO AND THE ABOVE-NAMED SUPPLIERS DISCLAIM ALL WARRANTIES, EXPRESSED OR IMPLIED, INCLUDING, WITHOUT LIMITATION, THOSE OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT OR ARISING FROM A COURSE OF DEALING, USAGE, OR TRADE PRACTICE.

IN NO EVENT SHALL CISCO OR ITS SUPPLIERS BE LIABLE FOR ANY INDIRECT, SPECIAL, CONSEQUENTIAL, OR INCIDENTAL DAMAGES, INCLUDING, WITHOUT LIMITATION, LOST PROFITS OR LOSS OR DAMAGE TO DATA ARISING OUT OF THE USE OR INABILITY TO USE THIS MANUAL, EVEN IF CISCO OR ITS SUPPLIERS HAVE BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

Any Internet Protocol (IP) addresses and phone numbers used in this document are not intended to be actual addresses and phone numbers. Any examples, command display output, network topology diagrams, and other figures included in the document are shown for illustrative purposes only. Any use of actual IP addresses or phone numbers in illustrative content is unintentional and coincidental.

All printed copies and duplicate soft copies are considered un-Controlled copies and the original on-line version should be referred to for latest version.

Cisco has more than 200 offices worldwide. Addresses, phone numbers, and fax numbers are listed on the Cisco website at www.cisco.com/go/offices.

© 2017 Cisco Systems, Inc. All rights reserved.

Cisco Trademark

Cisco and the Cisco logo are trademarks or registered trademarks of Cisco and/or its affiliates in the U.S. and other countries. To view a list of Cisco trademarks, go to this URL: www.cisco.com/go/trademarks. Third-party trademarks mentioned are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (1110R)