



Clustering and Peers

This section describes how to set up a cluster of Expressway peers. Clustering is used to increase the capacity of your Expressway deployment and to provide resiliency.

- [About Clusters, on page 1](#)
- [Cluster License Usage and Capacity Guidelines, on page 3](#)
- [Managing Clusters and Peers, on page 5](#)
- [Troubleshooting Cluster Replication Problems, on page 14](#)
- [Troubleshooting System Key Related Issues, on page 15](#)

About Clusters


An Expressway can be part of a cluster of up to six Expressways. Each Expressway in the cluster is a peer of every other Expressway in that cluster. When creating a cluster you define a cluster name and nominate one peer as the primary from which configuration is replicated to the other peers. Clusters are used for the following reasons:

- **Capacity.** Increase the capacity of your Expressway deployment compared with a single Expressway.
- **Resilience.** Provide redundancy while an Expressway is in [maintenance mode](#), or in the rare case that it becomes inaccessible due to a network / power outage, or some other reason.



Note There is no capacity gain after four peers. So in a six-peer cluster for example, the 5th and 6th Expressways do not add extra call capacity to the cluster. Resilience is improved with the extra peers, but not capacity.

Peers share information with each other about their use of bandwidth, registrations, and user accounts. This allows the cluster to act as one large Expressway Local Zone, as shown in this example:



454315

Cluster License Usage and Capacity Guidelines

This section describes how licenses are used across a cluster and provides capacity guidelines. For ease of reference, the capacity guidelines for standalone systems are also included here.

The maximum supported capacities / sizing for Cisco Expressway Series (not Cisco VCS) are listed in the tables below. These figures are guidelines only and are NOT guaranteed, because many factors affect performance in real-life deployments. Expressway supports so many different use cases that it is not possible to provide capacity limits for individual, specific deployments.

Expressway sizing / capacity information is categorized on the basis of the number of supported concurrent registrations and/or calls.

Important Caveats

- The figures provided here assume all necessary software licenses are applied.
- The figures are tested for specific, dedicated Expressway scenarios. Based on an Expressway or cluster being used for a single service or scenario, such as just for MRA or just for B2B calling. It's not possible to provide tested capacity guidelines for multi-service deployments.
- Up to six Expressway systems can be clustered, but this only increases capacity **by a maximum factor of four** (except Small VMs, which have no gain).
- For Small VMs, clustering is only for redundancy and not for scale and **there is no capacity gain from clustering**.
- The figures provided for video calls and audio-only calls are alternatives - the stated capacity is available either for video or for audio, not for both.

Dependencies

The figures for calls refer to concurrent calls.

Concurrent calls and Rich Media Session (RMS) licenses do not have a one-to-one relationship. Various factors determine RMS license usage, which means that some calls may be “free” and others may use multiple licenses.

To support 6000 TURN relays on a large system (Large VM or CE1200) you need to enable “TURN Port Multiplexing on Large Expressway” (**Configuration > Traversal > TURN**).

Small VMs are supported on the Cisco Business Edition 6000 platform, or on general purpose hardware / ESXi which matches the Cisco Business Edition 6000 specification. The figures for Small VMs are for M5-based BE6000 appliances.

Figures for Standalone Systems

This table shows the base capacity for a standalone Expressway.

Table 1: Standalone Capacity Guidelines - Single Expressway

Platform	Registrations (room/desktop)	Calls (video or audio-only)	RMS Licenses	MRA Registrations (proxied)	TURN Relays
CE1200	5000	500 video or 1000 audio	500	5000	6000
Large VM	5000	500 video or 1000 audio	500	2500	6000
Medium VM	2500	100 video or 200 audio	100	2500	1800
Small VM	2000	40 non-MRA video, or 20 MRA video or 40 audio	75	200	1800

Figures for Clustered Systems

This table illustrates the increased capacity for a clustered system with six Expressways (the maximum cluster size for scale gain). Each cluster can have up to six Expressway nodes and a maximum of $N+2$ physical redundancy. All nodes are active in the cluster.

Table 2: Clustered Capacity Guidelines - Example for Cluster with 4 Expressway Peers

Platform	Registrations (room/desktop)	Calls (video or audio-only)	RMS Licenses	MRA Registrations (proxied)	TURN Relays
CE1200	20,000	2000 video or 4000 audio	2000	20,000	24,000
Large VM	20,000	2000 video or 4000 audio	2000	10,000	24,000
Medium VM	10,000	400 video or 800 audio	400	10,000	7200
Small VM	2000	40 non-MRA video, or 20 MRA video or 40 audio	75	200	1800

Example Deployment

Say you want to deploy a resilient cluster that can handle up to 750 concurrent desktop registrations and 250 Rich Media Sessions. In this case you could configure 4 peers as follows:

	Peer 1	Peer 2	Peer 3	Peer 4	Total cluster capacity
Desktop registration licenses	250	250	250	0	750
Rich Media Sessions	100	100	50	0	250

In this example it doesn't matter which peer an endpoint registers to, as the licenses are shared across all of the peers. If any one peer is temporarily taken out of service the full set of call licenses remain available to the entire cluster.

Intracluster Calls

License usage when endpoints are registered to different peers in the same cluster, depends on call media traversal across the cluster:

- If call media does not traverse the cluster peers, a call between the endpoints does not use any RMS licenses (it's a "Registered" call).
 - If any of the endpoint is not registered to Cisco infrastructure then calls will use RMS license.
- If call media does traverse the cluster peers, a call between the endpoints uses an RMS license on the Expressway where the B2BUA is engaged.
 - If both the endpoints are registered to Cisco infrastructure then call will not use RMS license.

More information about how licenses are used in clustered systems is provided in the licensing section of this guide.

Managing Clusters and Peers

Setting Up a Cluster

Before you Start

1. Make sure that all prerequisites listed in the *Cisco Expressway Cluster Creation and Maintenance Deployment Guide* for your version are complete (on the [Cisco Expressway Series Configuration Guides](#) page).
2. We recommend that you backup your Expressway data before setting up a cluster. Instructions are in the *Cisco Expressway Cluster Creation and Maintenance Deployment Guide*.

Process

To create the cluster you must first configure a primary peer and then add the other peers into the cluster one at a time.

Maintaining a Cluster

The **Clustering** page (**System** > **Clustering**) lists the IP addresses of all the peers in the cluster, to which this Expressway belongs, and identifies the configuration primary peer.

Basics of Cluster Configuration

- The **Cluster name** is used to identify one cluster of Expressways from another. Set it to the fully qualified domain name (FQDN) used in SRV records that address this Expressway cluster, for example **cluster1.example.com**.

The FQDN can comprise multiple levels. Each level's name can only contain letters, digits and hyphens, with each level separated by a period (dot). A level name cannot start or end with a hyphen, and the final level name must start with a letter. A cluster name is required if FindMe is enabled.

- All peers must agree on which is the **Configuration primary**. Use the same number on each peer, and keep the **Peer N address** list in the same order on all peers.
- All peers must use the same IP version. Set the **Cluster IP version** to the same value on all peers.
- All peers must use the same **TLS verification mode**. Choose *Enforce* for better security, but be aware that the peers must be able to verify each others' certificates against their trusted CAs.
- The **Cluster Address Mapping** option allows you to map Cisco Expressway-E peers' FQDNs to their private IP addresses. Cluster address mapping allows you to enforce TLS clustering of peers in an isolated network, because it does not require the use of the public DNS and the peers' public IP addresses.

For details, see the *Expressway Cluster Creation and Maintenance Deployment Guide* on the [Expressway Configuration Guides](#) page.

Other Configuration for the Cluster

You should only make configuration changes on the primary Expressway.



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.



Caution

Dbxsh is a python script that connects to a cluster database on the local loopback address over port 4370. The Dbxsh does not need to authenticate the database before executing the commands. The port is open for connection and is strictly for internal use only. This is accessible from root only.

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 Items in Clustered Systems](#).

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

Adding and Removing Peers From a Cluster

After a cluster has been set up you can add new peers to the cluster or remove peers from it. For details see the *Expressway Cluster Creation and Maintenance Deployment Guide*.

**Caution**

If you clear all the peer address fields from the clustering page and save the configuration, then the Expressway will factory reset itself the next time you do a restart. This means you will lose all existing configuration except basic networking for the LAN1 interface, including all configuration that you do between when you clear the fields and the next restart.

Changing the Primary Peer

Typically you only need to change the **Configuration primary** in the following cases:

- If the original primary peer fails. (If the primary fails, the remaining peers continue to function normally except that, as they are unable to copy their configuration from the primary, they may become out of sync with each other.)
- To take the primary Expressway unit out of service.

For details about how to change the primary peer, see the *Expressway Cluster Creation and Maintenance Deployment Guide*.


Monitoring Cluster Status

The status sections at the bottom of the **Clustering** page show you the current status of the cluster, and the time of the previous and next synchronization.

Troubleshooting Cluster Problems

See [Troubleshooting Cluster Replication Problems](#).

Peer-Specific Items in Clustered Systems

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.

Configuration data that applies to all peers should only be modified on the primary peer. Otherwise, at best the changes will be overwritten from the primary or at worst the cluster replication will fail.

Service setup wizard

Configuration settings made through the service setup wizard (including Select Type, Select Series, service selection, licensing for those services, and basic network settings) must be configured on each peer in a cluster.

Cluster configuration (System > Clustering)

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

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

If you need to enable cluster address mapping, we recommend forming the cluster on IP addresses first. Then you only need to add the mappings on one peer.

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 unique IP address, whether that is an **IPv4 address**, an **IPv6 address**, or both.
- **IP gateway** configuration is peer-specific. Each peer can use a different gateway.



Note 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 Expressway. 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)

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

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)

This section only applies to systems that use PAK-based licensing (option keys do not apply if your system uses Smart Licensing). Option keys can control licensing or specific features. They are gradually being phased out for Expressway and their use is diminishing.

Option keys that control **licenses** are pooled for use by the whole cluster.

Option keys that control **features** (such as advanced account security or Microsoft Interoperability) are specific to the peer where they are applied. Each peer must have an identical set of feature option keys installed, which means that if you use option keys for features 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:

- Rich media sessions
- Telepresence room systems
- Desktop systems



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 a conferencing server must be unique for each peer in the cluster.

Sharing Registrations Across Peers

When a cluster peer receives a search request (such as an INVITE), it checks its own and its peers' registration lists before responding. This allows all endpoints in the cluster to be treated as if they were registered with a single Expressway.

Peers are periodically queried to ensure they are still functioning.

H.323 registrations

All the peers in a cluster share responsibility for their H.323 endpoint community. When an H.323 endpoint registers with one peer, it receives a registration response which contains a list of alternate gatekeepers, populated with a randomly ordered list of the IP addresses of all the other peers in that cluster.

If the endpoint loses contact with the initial peer, it will seek to register with one of the other peers. The random ordering of the list of alternate peers ensures that endpoints that can only store a single alternate peer will failover evenly across the cluster.

When using a cluster, you may want to reduce the registration **Time to live** on all peers in the cluster from the default 30 minutes. This setting determines how often endpoints are *required* to re-register with their Expressway, and reducing it means that if a cluster peer is unavailable, the endpoint will failover more quickly to an available peer.



Note By reducing the registration time to live too much, you risk flooding the Expressway 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.

To change this setting, go to **Configuration > Protocols > H.323 > Gatekeeper > Time to live**.

SIP registrations

The Expressway supports multiple client-initiated connections (also referred to as “SIP Outbound”) as outlined in [RFC 5626](#).

This allows SIP endpoints that support *RFC 5626* to be simultaneously registered to multiple Expressway cluster peers. This provides extra resiliency: if the endpoint loses its connection to one cluster peer it will still be able to receive calls via one of its other registration connections.

You can also use DNS round-robin techniques to implement a registration failover strategy. Some SIP UAs, such as Jabber Video, can be configured with a SIP server address that is an FQDN. If the FQDN resolves to a round-robin DNS record populated with the IP addresses of all the peers in the cluster, then this could allow the endpoint to re-register with another peer if its connection to the original peer is lost.

Sharing Bandwidth Across Peers

When clustering has been configured, all peers share the bandwidth available to the cluster.

- Peers must be configured identically for all aspects of bandwidth control including subzones, links and pipes.

- Peers share their bandwidth usage information with all other peers in the cluster, so when one peer is consuming part or all of the bandwidth available within or from a particular subzone, or on a particular pipe, this bandwidth will not be available for other peers.

For general information on how the Expressway manages bandwidth, see the [bandwidth control](#) section.

Cluster Upgrades, Backup, and Restore

Upgrading a cluster

See the *Cisco Expressway Cluster Creation and Maintenance Deployment Guide*, for your version, on the [Cisco Expressway Series Configuration Guides](#) page.



Note If you are upgrading to X8.8 or later from an earlier version, clustering communications changed in X8.8 to use TLS connections between peers instead of IPSec. TLS verification is not enforced (by default) after you upgrade, and you'll see an alarm reminding you to enforce TLS verification.

Backing up a cluster

Use the [backup and restore](#) process to save cluster configuration information. The backup process saves all configuration information for the cluster, regardless of the Expressway used to make the backup.



Caution Do not take VMware snapshots of Cisco Expressway systems. The process interferes with database timing and negatively impacts performance.

Restoring a cluster

To restore previously backed up cluster configuration data, follow this process.



Important You can't restore data to an Expressway that is part of a cluster. As described here, first remove the Expressway peer from the cluster. Then do the restore. (After the restore you need to build a new cluster.)

1. Remove the Expressway peer from the cluster so that it becomes a standalone Expressway.
2. Restore the configuration data to the standalone Expressway. See [Restoring a Previous Backup](#) for details.
3. Build a new cluster using the Expressway that now has the restored data.
4. Take each of the other peers out of their previous cluster and add them to the new cluster. See [Setting Up a Cluster](#) for details.



Note No additional steps are required if you are using FQDN's and have a valid cluster address mapping configured. Mappings will be configured on a restore action.

Clustering and Cisco TMS

Cisco TMS version 13.2 or later is mandatory if your cluster is configured to use FindMe or Device Provisioning.

Size limitations for clusters and provisioning

An Expressway cluster of any size supports up to:

- 10,000 FindMe accounts
- 10,000 users for provisioning
- 200,000 phonebook entries



Note Even if the [Cluster License Usage and Capacity Guidelines](#) of your system is greater, you are limited to 10,000 FindMe accounts/users and 10,000 provisioned devices per cluster.

If you need to provision more than 10,000 devices, your network will require additional Expressway clusters with an appropriately designed and configured dial plan.

See the *Cisco Expressway Cluster Creation and Maintenance Deployment Guide*, for your version, on the [Cisco Expressway Series Configuration Guides](#) page.

About the Cluster Subzone

When two or more Expressways are clustered together, a new subzone is created within the cluster's Local Zone. This is the Cluster Subzone (see the diagram in the [About Clusters](#) section). Any calls between two peers in the cluster will briefly pass via this subzone during call setup.

The Cluster Subzone is (like the Traversal Subzone) a virtual subzone used for call routing only, and endpoints cannot register to this subzone. After a call has been established between two peers, the Cluster Subzone will no longer appear in the call route and the call will appear as having come from (or being routed to) the Default Subzone.

The two situations in which a call will pass via the Cluster Subzone are:

- Calls between two endpoints registered to different peers in the cluster.

For example, Endpoint A is registered in the Default Subzone to Peer 1. Endpoint B is also registered in the Default Subzone, but to Peer 2. When A calls B, the call route is shown on Peer 1 as **Default Subzone -> Cluster Subzone**, and on Peer 2 as **Cluster Subzone -> Default Subzone**.

- Calls received from outside the cluster by one peer, for an endpoint registered to another peer.

For example, we have a single Expressway for the Branch Office, which is neighbored to a cluster of 4 Expressways at the Head Office. A user in the Branch Office calls Endpoint A in the Head Office. Endpoint A is registered in the Default Subzone to Peer 1. The call is received by Peer 2, as it has the lowest resource usage at that moment. Peer 2 then searches for Endpoint A within the cluster's Local Zone, and finds that it is registered to Peer 1. Peer 2 then forwards the call to Peer 1, which forwards it to Endpoint A. In this case, on Peer 2 the call route will be shown as **Branch Office -> Default Subzone -> Cluster Subzone**, and on Peer 1 as **Cluster Subzone -> Default Subzone**.



Note If **Call signaling optimization** is set to *On* and the call is H.323, the call will not appear on Peer 2, and on Peer 1 the route will be **Branch Office > Default Subzone**.

Neighboring Between Expressway Clusters

You can neighbor your local Expressway (or Expressway cluster) to a remote Expressway cluster. The remote cluster might be a neighbor, traversal client, or traversal server to the local system. When a call is received on the local Expressway and is passed via the relevant zone to the remote cluster, it gets routed to whichever peer in that neighbor cluster has the lowest resource usage (peers in maintenance mode are not considered). That peer then forwards the call to one of the following:

- A locally registered endpoint, if the endpoint is registered to that peer.
- A peer, if the endpoint is registered to another peer in the cluster.
- An external zone, if the endpoint is located elsewhere.

Lowest resource usage is determined by comparing the number of available media sessions (maximum - current use) on the peers, and choosing the peer with the highest number.

Expressways that are configured as peers **must not also be configured as neighbors** to each other, or the other way round.

Process to Neighbor Clusters

You create a single zone on the local system to represent the connection to the remote cluster, and configure it with the details of all the peers in the remote cluster. Adding this information to the zone ensures that the call is passed to that cluster regardless of the status of the individual peers.

1. On the local Expressway (or on the primary peer for a cluster), create a zone of the appropriate type.
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. You do not do this for **traversal server** zones, as these connections are not configured by specifying the remote system's address.

Ideally, 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 Expressway cluster are listed here does not matter.



Note Whenever you add an extra Expressway to a cluster, you need to modify any Expressways which neighbor to that cluster to let them know about the new peer.

Troubleshooting Cluster Replication Problems

Cluster replication can fail for a variety of reasons. This section describes the most common problems and how to resolve them. For more detailed information:

See the *Cisco Expressway Cluster Creation and Maintenance Deployment Guide*, for your version, on the [Cisco Expressway Series Configuration Guides](#) page.

Some peers have a different primary peer defined

1. For each peer in the cluster, go to the **System > Clustering** page.
2. Ensure each peer identifies the same **Configuration primary**.

Unable to reach the cluster configuration primary peer

The Expressway operating as the primary peer could be unreachable for many reasons, including:

- Network access problems
- Expressway unit is powered down
- Incorrectly configured addresses
- TLS verification mode is set to Enforce but some peers have invalid or revoked certificates
- Different software versions on peers
- DNS settings not correct in cluster

“Manual synchronization of configuration is required” alarms are raised on subordinate peer Expressways

1. Log in to the peer as **admin** through the CLI (available by default over SSH and through the serial port on hardware versions).
2. Type **xCommand ForceConfigUpdate**.

This will delete the subordinate Expressway peer's configuration and force it to update its configuration from the primary Expressway.



Caution

Never issue this command on the primary Expressway because you will lose all configuration for the cluster.

“Cluster config error” alarms are raised on Expressway peer

You can specify a new configuration primary on the clustering page as per the description of the alarm raised.



Note

You can revert to the old configuration primary once all the replication alarms are lowered.

Incorrect IP to FQDN mappings

1. Go to the **System > Clustering** page on any peer.
2. Check that all FQDN and IP addresses have been entered correctly.

Firewall preventing the cluster communicating

- If you intended to cluster using public IP addresses, make sure your firewall isn't preventing cluster communication by blocking the clustering communications ports. If it is, consider whether you can change your firewall rules.
- If you intended to cluster with private addresses, ensure you have configured your cluster as per our recommendations, i.e. form a cluster using FQDN with IP address mappings, and TLS authentication.

Troubleshooting System Key Related Issues

This section describes the most common problems related to system key and how to resolve them.

“Failed to update key file” alarms are raised on Expressways (Single node scenario)

1. Log in as admin through the CLI (available by default over SSH and through the serial port on hardware versions).
2. Type **xCommand ForceSystemKeyUpdate**.

“Failed to update key file” alarms are raised on Expressways (Cluster scenario)

1. Log in to node as admin through the CLI (available by default over SSH and through the serial port on hardware versions) where this alarm is not raised.
2. Type **xCommand ForceSystemKeyUpdate**.



Note Make sure to address “Failed to update key file” alarm before adding the node to a cluster.
