



Enterprise Chat and Email Design Guide, Release 12.6(1)

**For Unified Contact Center Enterprise and Packaged Contact Center
Enterprise**

First Published: May, 2021

Last Updated: November, 2024

Americas Headquarters

Cisco Systems, Inc.
170 West Tasman Drive
San Jose, CA 95134-1706
USA
<http://www.cisco.com>
Tel: 408 526-4000
800 553-NETS (6387)
Fax: 408 527-0883

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.

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 <http://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)

Enterprise Chat and Email Design Guide: For Unified Contact Center Enterprise. November 8, 2024

Copyright © 2016–2021, Cisco Systems, Inc. All rights reserved.

Contents

- Preface5**
 - About This Guide 6
 - Change History 6
 - Related Documents 7
 - Communications, Services, and Additional Information 7
 - Cisco Bug Search Tool 7
 - Documentation Feedback 8
 - Field Alerts and Field Notices 8
 - Document Conventions 8

- Chapter 1: System Flow and Configuration Design.....9**
 - Email Routing Through Unified CCE 10
 - Chat Routing Through Unified CCE 11
 - Callback Routing Through Unified CCE 12
 - Designing Integrated Routing - An Example 12
 - Unified CCE Configuration 12
 - Configuration in ECE 13

- Chapter 2: Sizing Guidelines15**
 - Sizing Inputs 16
 - Planning for Database Growth 16
 - About Deploying on Virtual Servers 17
 - Sizing for Combined Email, Web, and Voice Scenarios 17
 - Important Information About Sizing 17
 - Support for up to 400 Standardized Concurrent Agents 17
 - Support for 401 to 1500 Standardized Concurrent Agents 18
 - Support for 401 to 2500 Standardized Concurrent Agents 18
 - Additional Configurations for Sizing 19
 - Expanding the Disk Size on Database Server Machines 19

Chapter 3: Fault Tolerance and Redundancy	20
Load Balancing Considerations	21
High Availability Options.....	22
Managing Failover	23
Chapter 4: Network Latency.....	25
Network Latency.....	26
Bandwidth Requirements	26
Geographic Server Distribution	26
Chapter 5: Firewall and Hardening	27
Firewall Considerations	28
Server Hardening Considerations	28
Default Windows and IIS Service Requirements for ECE	28
Guidelines for Microsoft SQL Server.....	29
Appendix: Failover Behavior.....	30
Customer and Agent Behavior During Failover.....	31

Preface

- ▶ [About This Guide](#)
- ▶ [Change History](#)
- ▶ [Related Documents](#)
- ▶ [Communications, Services, and Additional Information](#)
- ▶ [Documentation Feedback](#)
- ▶ [Field Alerts and Field Notices](#)
- ▶ [Document Conventions](#)

Welcome to the Enterprise Chat and Email (ECE) feature, which provides multichannel interaction software used by businesses all over the world as a core component to the Unified Contact Center Enterprise product line. ECE offers a unified suite of the industry’s best applications for chat and email interaction management to enable a blended agent for handling of web chat, email and voice interactions.

About This Guide

Enterprise Chat and Email Design Guide is intended for engineers, system architects, and other technical audience responsible for planning the deployment and maintenance of Enterprise Chat and Email for Cisco Unified Contact Center Enterprise (Unified CCE) and Cisco Packaged Contact Center Enterprise (PCCE).

The document is designed to provide an overview of the system flow for different types of interactions, sizing guidelines, load-balancing options, network latency considerations, firewall considerations, and interface boundaries.

For sizing guidelines for servers used in the Unified CCE deployment, refer to the Unified CCE SRND guide available here: <https://www.cisco.com/go/srnd>.

Change History

This table lists changes made to this guide. Most recent changes appear at the top.

Change	See	Date
Added information on data retention using purge jobs.	“About Deploying on Virtual Servers” on page 17	November, 2024
Updated information on how to manage data growth	“Important Information About Sizing” on page 17	
Added customer and agent side observations for the following action performed during a failover: ▶ Maintenance mode is executed on the UCCE side	“Customer and Agent Behavior During Failover” on page 31	October, 2024
Added more agent side observations for the following actions performed during a failover: ▶ Agent completes the chat ▶ Agent does not pick the activity and the activity is auto-pushbacked	“Customer and Agent Behavior During Failover” on page 31	May, 2024
Added more agent side and customer side observations for different actions performed during a failover.	“Customer and Agent Behavior During Failover” on page 31	March, 2024

Change	See	Date
Added a new section: Support for 401 to 2500 Standardized Concurrent Agents	“Support for 401 to 2500 Standardized Concurrent Agents” on page 18	July, 2023
Added a new chapter	“Appendix: Failover Behavior” on page 30	November, 2022
Added more information around installing web and application servers separately.	“Managing Failover” on page 23	

Related Documents

The latest versions of all Cisco documentation can be found online at <http://www.cisco.com>

Subject	Link
Complete documentation for Enterprise Chat and Email, for both Cisco Unified Contact Center Enterprise (UCCE) and Cisco Packaged Contact Center Enterprise (PCCE)	https://www.cisco.com/c/en/us/support/customer-collaboration/cisco-enterprise-chat-email/tsd-products-support-series-home.html

Communications, Services, and Additional Information

- ▶ To receive timely, relevant information from Cisco, sign up at [Cisco Profile Manager](#).
- ▶ To get the business impact you’re looking for with the technologies that matter, visit [Cisco Services](#).
- ▶ To submit a service request, visit [Cisco Support](#).
- ▶ To discover and browse secure, validated enterprise-class apps, products, solutions and services, visit [Cisco Marketplace](#).
- ▶ To obtain general networking, training, and certification titles, visit [Cisco Press](#).
- ▶ To find warranty information for a specific product or product family, access [Cisco Warranty Finder](#).

Cisco Bug Search Tool

[Cisco Bug Search Tool](#) (BST) is a web-based tool that acts as a gateway to the Cisco bug tracking system that maintains a comprehensive list of defects and vulnerabilities in Cisco products and software. BST provides you with detailed defect information about your products and software.

Documentation Feedback

To provide comments about this document, send an email message to the following address:
contactcenterproducts_docfeedback@cisco.com

We appreciate your comments.

Field Alerts and Field Notices

Cisco products may be modified or key processes may be determined to be important. These are announced through use of the Cisco Field Alerts and Cisco Field Notices. You can register to receive Field Alerts and Field Notices through the Product Alert Tool on Cisco.com. This tool enables you to create a profile to receive announcements by selecting all products of interest.

Log into www.cisco.com and then access the tool at <http://www.cisco.com/cisco/support/notifications.html>

Document Conventions

This guide uses the following typographical conventions.

Convention	Indicates
<i>Italic</i>	Emphasis. Or the title of a published document.
Bold	Labels of items on the user interface, such as buttons, boxes, and lists. Or text that must be typed by the user.
Monospace	The name of a file or folder, a database table column or value, or a command.
<i>Variable</i>	User-specific text; varies from one user or installation to another.

Document conventions

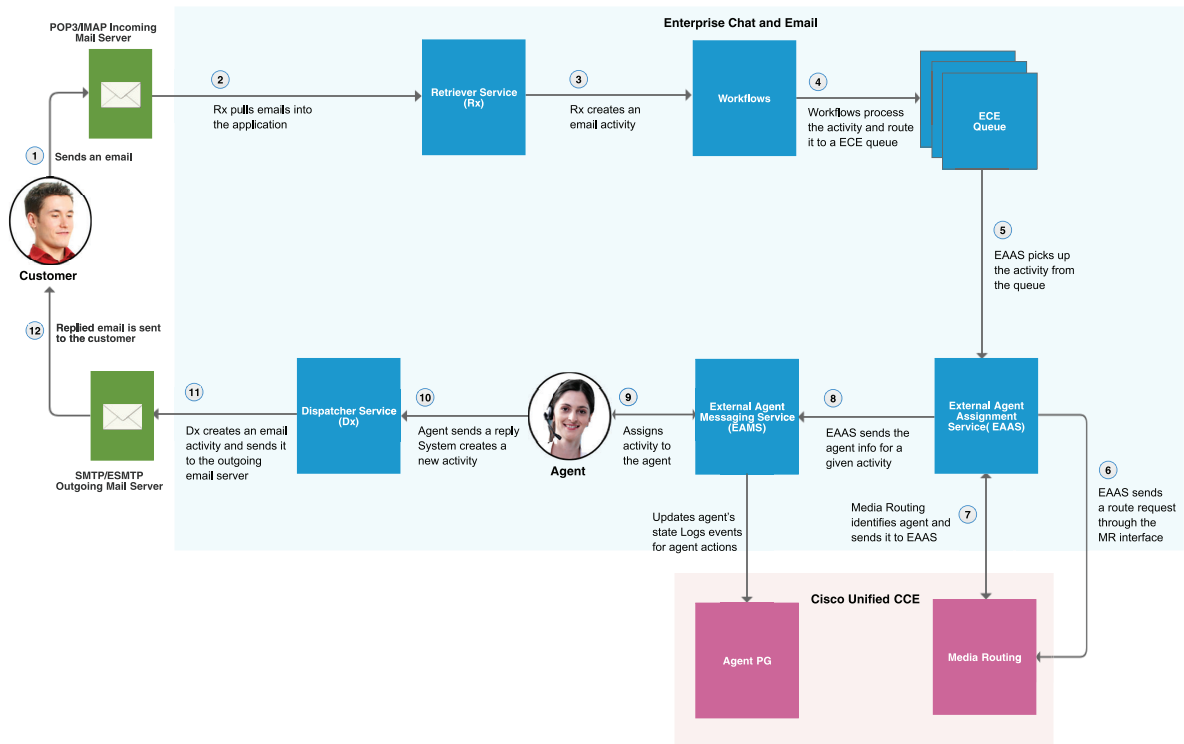
1 System Flow and Configuration Design

- ▶ [Email Routing Through Unified CCE](#)
- ▶ [Chat Routing Through Unified CCE](#)
- ▶ [Callback Routing Through Unified CCE](#)
- ▶ [Designing Integrated Routing - An Example](#)

This chapter illustrates the major components involved in the life-cycle of email, callback, and chat activities as they move through the system. It includes the various steps in the handling of an interaction from the point at which a customer query is received to it being routed to the appropriate user or queue.

Email Routing Through Unified CCE


There are various steps involved in efficiently responding to emails from customers. Emails are first retrieved into the system and routed to appropriate users or queues. Once a response is created, it is processed through the system and sent to the customer. The general system flow for integrated email routing is described in this section.



General system flow for email routing

Chat Routing Through Unified CCE


There are various steps involved in efficiently responding to chat requests received from customers. The general system flow for handling a chat request in a system integrated with Unified CCE (Enterprise and Packaged) is included here.



General system flow for chat routing

Callback Routing Through Unified CCE

Callback allows customers to submit a phone number using a web form. An agent who can handle the call is identified, and a call is placed, connecting the agent to the customer. The call back can happen right away, or after a period of time chosen by the customer.



System flow for web and delayed callback

Designing Integrated Routing - An Example

This section includes a sample integrated routing scenario and some best practices around creating MRDs and skill groups for a deployment with blended agents.

The following example shows an email deployment, but is also relevant for chat.

Unified CCE Configuration

The following objects are configured in Unified CCE.

1. One media routing domain for emails: **Email_MRD**.
2. Two peripherals: **NA_Peripheral**, and **EMEA_Peripheral**.

- The peripheral to skill group to agent relationship is as follows. All eight skill groups are associated with the Email_MRD.

Peripheral	Skill Group	Agents
NA_Peripheral	Domestic_Sales_NA	Clark, Bruce
	Global_Sales_NA	Clark
	Domestic_Support_NA	Peter
	Global_Support_NA	Tony
EMEA_Peripheral	Domestic_Sales_EMEA	Natasha, Clint
	Global_Sales_EMEA	Natasha
	Domestic_Support_EMEA	Loki
	Global_Support_EMEA	Loki, Phil

- Configured scripts route to the appropriate skill groups as follows.

Script	Skill Group
Global_Sales_Email_Script	Global_Sales_NA
	Global_Sales_EMEA
Global_Support_Email_Script	Global_Support_NA
	Global_Support_EMEA
Domestic_Sales_Email_Script	Domestic_Sales_NA
	Domestic_Sales_EMEA
Domestic_Support_Email_Script	Domestic_Support_NA
	Domestic_Support_EMEA

Configuration in ECE

- Two departments are configured in ECE, one each for **Sales** and **Support**. Both use the same single MRD, **Email_MRD**.
- The Sales department has two queues, one each for each of the script selectors for sales (global and domestic). Similarly, the Support department also has two queues. Note that by default, when an MRD is imported into ECE, queues are created using the following naming convention *<MRDName_ScriptSelectorName>*. All queue names are editable. For details about queues, see the *Enterprise Chat and Email Administrator's Guide to Routing and Workflows*.

In this scenario, the department to queue to skill group mapping, and the agents who service activities is as follows:

Department	Queue	Mapped Skill Groups	Agents Who Service Activities
Sales	Email_MRD_Domestic_Sales_Email_Script	Domestic_Sales_NA Domestic_Sales_EMEA	Clark, Bruce, Natasha, Clint
	Email_MRD_Global_Sales_Email_Script	Global_Sales_NA Global_Sales_EMEA	Clark, Natasha
Support	Email_MRD_Domestic_Support_Email_Script	Domestic_Support_NA Domestic_Support_EMEA	Peter, Loki
	Email_MRD_Global_Support_Email_Script	Global_Support_NA Global_Support_EMEA	Peter, Loki, Phil

Although this example does not use precision queues, you can extend the design to include them. Ensure that the appropriate precision queues are associated with the correct queues in ECE to make sure that activities are routed to the correct agent.

To keep the design simple yet effective, here are some general guidelines to follow while designing a system for unified routing:

1. Create one MRD per channel, for example, one for incoming email, one for outgoing email, one for chat and one for callback. Note that email MRDs are always interruptible in ECE.
2. Do not share skill groups and precision queues across departments.
3. While associating skill groups with queues in ECE, ensure that the association matches the skill groups selected in the queue to skill group node of the script selector.
4. While associating precision queues with queues in ECE, ensure that the association matches the PQ association with the script selector.



Sizing Guidelines

- ▶ [Sizing Inputs](#)
- ▶ [Planning for Database Growth](#)
- ▶ [About Deploying on Virtual Servers](#)
- ▶ [Sizing for Combined Email, Web, and Voice Scenarios](#)
- ▶ [Support for 401 to 2500 Standardized Concurrent Agents](#)

Sizing Inputs

Configurations presented in the following sections provide sizing for standardized agents who handle up to 6 email messages per hour, or one chat session at a time. If agents are expected to handle more than 6 email messages per hour, on average, or more than one chat session at a time, convert the agent count into a standardized agent count using the following formula:

Email:

$$\text{Standardized agent count for email} = \text{Actual agent count} * \text{Average number of messages handled per hour by each agent} / 6$$

Chat:

$$\text{Standardized agent count for chat} = \text{Actual agent count} * \text{Average number of concurrent chat sessions handled by each agent}$$

Email and Chat:

$$\text{Standardized agent count} = \text{Standardized agent count for email} + \text{Standardized agent count for chat}$$

Use the standardized agent count to find the appropriate configuration to fit your needs. For sizing for Combined Email, Chat, and Voice Scenarios, see [page 17](#).



Important: The number of concurrent agents per application server cannot exceed 400, as this is the maximum number of concurrent agents that can be supported for email and chat by one ECE server in a deployment of ECE.

Planning for Database Growth

The following factors are considered for calculating the rate of growth of database.

- ▶ Incoming and outgoing email volume per month.
- ▶ Number of email attachments
- ▶ Average size of each email (KB).

The following formula can be used to compute the approximate rate of growth of the database server (MB) per month for activities of type email:

$$((\text{Number of incoming and outgoing emails per month} * 2) * (6 + (\text{Average size of each email message in KB} * 2))) / 1024$$

If your system receives emails with attachments, use the following formula and add it to the value for emails.

$$((\text{Number of emails per month with attachments} * \text{Average size of attachments (K)}) / 1024)$$

For example, if average volume of incoming and outgoing emails with attachments is 50,000, and average size of each attachment is 5 KB, monthly rate of growth can be computed as:

$$((50,000 * 5) / 1024) = 245 \text{ MB per month}$$

The following formula can be used to compute the approximate rate of growth of the database server (MB) per month for activities of type chat or callback:

$((\text{Number of incoming and outgoing chat messages per month}) * (6 + (\text{Average size of each chat message in KB} * 3 * 2))) / 1024$



Important: These formulas are meant to be used to plan for database growth. Values arrived at using computation may not be an exact match to actual sizes

About Deploying on Virtual Servers

All components in the ECE deployment must be installed on virtual servers using the OVA templates available on cisco.com to ensure proper sizing and resource utilization. For details see the [Virtualization for Enterprise Chat and Email](#) guide and refer to the [Version 12.0](#) section.

Note that database server disk space can be expanded to accommodate data growth over time. By default, data is not automatically archived or purged from the databases. However, users can choose to manage data retention by setting up a purge job through the Administration Console. For details about increasing the disk size, see [“Expanding the Disk Size on Database Server Machines” on page 19](#). For details about implementing a deployment on virtual servers, see [Cisco Collaboration Virtualization](#).

Sizing for Combined Email, Web, and Voice Scenarios

ECE can support multiple media, namely, email, chat, and callback. The following combinations of users can be supported on respective configurations described here.

Important Information About Sizing

- ▶ For a distributed deployment, the concurrent load must be spread evenly across all the web-application servers in the cluster.
- ▶ In the sizing configurations described here, dual CPU can optionally be replaced by 2 single core CPUs and a quad CPU can optionally be replaced by 4 single core CPUs.
- ▶ Sizing is not affected by the existence of a firewall between the web server and the application server, and by whether the web and application servers are collocated or not.
- ▶ Data growth can be managed either by increasing the database disk size or by configuring purge jobs.

Support for up to 400 Standardized Concurrent Agents

Support for up to 400 concurrent agents handling email, chat, or callback, where each agent can work on emails at the rate of 6 emails per hour, or work on a single active chat or callback session, at the rate of 6 chat sessions per hour. Configuration supports an incoming email rate of up to 120,000 emails per month.

In ECE, any combination of agent-customer chat sessions, callback sessions, and email activities agents totaling to 400, can be supported on a two-server configuration consisting of one web server, and another server with the web, application, file, messaging, services, and database components.

This configuration also requires two workflow processes and instances to be configured in the application. For details see the *Enterprise Chat and Email Administrator's Guide to System Resources*.

To deploy this configuration, use the following OVA template:

▶ `ECE_12.6_400_Win2019_vmv13_v1.0.ova`

Support for 401 to 1500 Standardized Concurrent Agents

Support for 401 to 1500 concurrent agents handling email, chat, or callback, where each agent can work on emails at the rate of 6 emails per hour, or work on a single active chat or callback session, at the rate of 6 chat sessions per hour.

This is supported on a configuration consisting of five web servers, five application servers, one file server, one messaging server, one services server, and one database server. Load must be evenly distributed across the web servers.

This configuration also requires two workflow processes and instances to be configured in the application. For details see the *Enterprise Chat and Email Administrator's Guide to System Resources*.

To deploy this configuration, use the following OVA template:

▶ `ECE_12.6_1500_Win2019_vmv13_v1.0.ova`

Support for 401 to 2500 Standardized Concurrent Agents

Support for 401 to 2500 concurrent agents handling email, chat, or callback, where each agent can work on emails at the rate of 6 emails per hour, or work on a single active chat or callback session, at the rate of 6 chat sessions per hour.

This is supported on a configuration consisting of three web servers, three application servers, one file server, one messaging server, one services server, and one database server. Load must be evenly distributed across the web servers. Please note that the support for 2500 agents is available from ECE 12.6 ES3 onwards.

This configuration also requires two workflow processes and instances to be configured in the application. For details see the *Enterprise Chat and Email Administrator's Guide to System Resources*.

To deploy this configuration, use the following OVA template:

▶ `ECE_12.6_2500_Win2019_vmv13_v1.0.ova`

Additional Configurations for Sizing

Expanding the Disk Size on Database Server Machines

Over time as the data in your databases grows, you will need to increase the disk space on the database server machines. Follow the instructions provided by VMWare to increase the size of the disk on all the machines on which the data files for the active, master, and reports databases are installed.

Before you increase the disk space, complete the following tasks

- ▶ Stop the application
- ▶ Stop SQL Server

Now expand the disk space on the servers where the data files are installed following the VMWare guidelines for your particular operating system. Additional information is available here:

http://kb.vmware.com/selfservice/search.do?cmd=displayKC&docType=kc&docTypeID=DT_KB_1_1&externalId=1004071

After you expand the size, complete the following tasks

- ▶ Restart the Windows server
- ▶ Log into the database server where the data files reside and ensure that the disk space has been expanded to the new and increased size
- ▶ Restart SQL Server
- ▶ Restart the application on all the servers in the deployment

Fault Tolerance and Redundancy

- ▶ [Load Balancing Considerations](#)
- ▶ [High Availability Options](#)
- ▶ [Managing Failover](#)

To optimize resource utilization and enhance performance, access to the application can be set up for fault tolerance and redundancy.

Load Balancing Considerations


The web service component of a ECE deployment can be load-balanced to serve large number of agents accessing the application at the same time. The web (or web-application) servers can be configured behind the load balancer with a virtual IP address, and a user can access ECE through this IP address. Depending on the load balancing algorithm set, the load balancer will send a request to one of the web-application server behind it and send a response back to the user. This way, from a security perspective, the load balancer serves as a reverse proxy server too.

One of the most essential parameters while configuring a load balancer is to configure it to support sticky sessions with cookie based persistence. After every scheduled maintenance task, before access is opened for users, it is advised to verify that all web-application servers are available to share the load. In absence of this, the first web-application server could be overloaded, due to the sticky connection feature. With other configurable parameters, one can define load balancing algorithms to meet various objectives such as equal load balance, isolation of a web-application server, or sending lesser requests to a low powered web-application servers.

The load balancer monitors the health of all web-application servers in the cluster, and if a problem is observed, the load balancer removes the given web-application server from the available pool of servers, thus preventing new web requests from being directed to the problematic web-application servers.

High Availability Options

Based on typical customer deployment scenarios, the following recommendations apply towards achieving a high-available system deployment.



Aspect	Description	Advantages
Load Balancer	The load balancer is used for distributing web requests across different web servers. Various types of load balancers are available in the industry. Each of these could be configured with different options for work distribution, handling failures, or increased activity.	<ul style="list-style-type: none"> ▶ Helps distribute load across different servers. ▶ Helps configure load distribution based on server capacity and current server resources. ▶ Helps handle failures by alternate means of routing a web request.
High-Speed Dedicated LAN	The network is a key ingredient to a successful and highly available application. When agents connect to the application remotely, the permissible network latency is 300 milliseconds (one way). Higher latency between the agents and the applications servers could lead to slower performance on the agent interface.	<ul style="list-style-type: none"> ▶ Stable network connections for distributed components. ▶ Helps serve web requests in a more predictable and reliable manner. ▶ Less delay in responses and thereby increases user experience.
Configuring more than one web-application server	<p>It is highly recommended that more than 1 web-application server be configured.</p> <p>The load balancer detects web server failures and redirects requests to other available web servers, after which, users will have to re-login to the application and a new user session will be created on the target web server.</p>	<ul style="list-style-type: none"> ▶ More than one web-application server helps to load balance web requests to multiple servers based on both system load and availability of servers. ▶ Helps the system scale better to meet growing needs of the enterprise.
Configuring clustering for Messaging Servers	<p>A cluster of messaging servers can be installed to achieve failover. When the application starts, it connects to one messaging server in the cluster. If this server goes down the application automatically connects to the next available messaging server in the cluster.</p> <p>Once a server in the cluster becomes unavailable it has to be restarted manually.</p>	<ul style="list-style-type: none"> ▶ Application continues to be available even when one messaging server goes down. ▶ If the messaging server is installed on a separate machine, it can be restarted without affecting any other components in the deployment.
Configuring Primary and Secondary Services servers	Two services server components, one primary and one secondary can be installed to achieve failover. Both components stay active, and when the primary goes down, the secondary services server automatically takes over.	<ul style="list-style-type: none"> ▶ The application continues to process activities even when a services server component goes down ▶ Removes single points of failure from the deployment.
Configuring Always On for databases	MS SQL Server Always On is supported and can be configured to achieve high availability for the ECE databases.	

Recommendations for high availability needs

In addition to these recommendations, if a load balancer is configured to monitor the health of web-application servers, it also serves the purpose of high availability.

Managing Failover

Some of the key methods of handling failure conditions within an ECE and Unified CCE or Packaged CCE integrated deployment are listed here.

- ▶ **Web and Application Servers:** Multiple web-application servers can be deployed in any distributed server deployment. Please note that the web and application servers are installed on separate VMs. If any of the web or application servers go down, a load balancer can help handle the failure through routing requests to alternate web-application servers. The load balancer detects application server failure and redirects requests to another application server. Users can log into the application without experiencing any significant loss of productivity.

Also note that in deployments without a loadbalancer, if one or more application servers crash, the entire application does not need to be restarted. Only the affected application servers have to be restarted. The rest of the application continues to function normally.

- ▶ **File, Services and Messaging Servers:** ECE is certified with VMware versions that support VMware High Availability. When VMware HA is configured, automatic failover is managed by VMware.

For the Messaging server, a cluster of messaging servers can also be configured for fail-over across geographies. For details on configuring failover see the installation guides.

For the Services server, the active-active mode is supported. Two Services server components are installed to achieve fail-over.

- ▶ **Unified CCE components:** The deployment can allow application services to failover with duplex Unified CCE components (e.g., MR PIM of MR PG and CTI Server of CTI Gateway) to eliminate downtime of the application in failure circumstances.



Network Latency

- ▶ [Network Latency](#)
- ▶ [Bandwidth Requirements](#)
- ▶ [Geographic Server Distribution](#)

Network Latency

Like any web-based application, set up Enterprise Chat and Email in a high-performance network environment that has sufficient bandwidth with low latency. If the network conditions degrade, it could have an undesirable impact on application performance. Listed here are mandatory guidelines to reduce network latency:

- ▶ When agents connect to the application remotely, the permissible network latency is 300 milliseconds (one way). Higher latency between the agents and the applications servers could lead to slower performance on the agent interface.
- ▶ The maximum permissible one-way network delay between the ECE servers and the Unified CCE servers is 300 milliseconds.

It is also important to note that bandwidth is also related to what the user perceives as good performance. For example, one typical “operation” within the application may take n seconds to complete with certain bandwidth, and it may take $n - m$ seconds to complete, if the available bandwidth is more. In both cases, application is usable, although one user perceives it to be faster than the other.

Bandwidth Requirements

The minimum required network bandwidth for an agent connecting to the ECE servers at login is 384 kilobits/second or higher. After login, at a steady state, an average bandwidth of 40 kilobits/second or higher is required.

An attachment size of up to 50 KB can be accommodated within this required bandwidth. For attachments of size greater than 50 KB, temporary slowness may be experienced in the agent user interface during download of the attachments.

Geographic Server Distribution

ECE provides support for geographical distribution of ECE application components. The network latency each way between the ECE servers and the Unified CCE components must be ≤ 300 milliseconds in order to ensure optimal communication between ECE and the geographically distributed Unified CCE components.



Firewall and Hardening

- ▶ [Firewall Considerations](#)
- ▶ [Server Hardening Considerations](#)

This chapter discusses some of the firewall and hardening considerations that are useful for Enterprise Chat and Email.

Firewall Considerations

- ▶ For agents to access Enterprise Chat and Email, the HTTPS (for secured connections) port needs to be opened at the firewall.

Considerations of applying firewall rules may vary depending on the security policies in effect. If a web server is configured within the firewall with access to the file server ports, Port 139 or 445 to the file server can be blocked from outside the firewall.

- ▶ In a typical installation where agents using Enterprise Chat and Email could be spread across multiple locations, the load balancer, along with the Enterprise Chat and Email web servers, may be deployed in a DMZ. This is a required deployment for installations where customers enter chat sessions from outside the intranet. However, having the web-application servers within the intranet is possible, too. The services and database server can reside in the network over the same or different VLAN.

If integration of these servers is implemented with Active Directory, then associated ports should be opened for communication with Domain Controllers.

Server Hardening Considerations

Dual strategies could be implemented towards securing the application. The first includes implementing standard best practices for physical and software level access controls. These steps could typically be at the corporate level. The other measure is hardening of the server OS and its service components.

Default Windows and IIS Service Requirements for ECE

- ▶ *In Accessories*, No Document Templates, No Mouse Pointers.
- ▶ *In Communications*, No Hyper Terminal.
- ▶ *In Application Server*, No Application Server Console, No ASP.NET, No Enable network DTC access, No Message Queuing, IN IIS, No BITS, NO FTP, No FrontPage, No Internet Printing, No NNTP, No, SMTP, In WWW, only WWW Services.
- ▶ No Certificate Services
- ▶ No Email and Fax Services
- ▶ No Indexing Services
- ▶ No Networking Services
- ▶ No Other Network Files & Print Services
- ▶ No Security Configuration Wizard
- ▶ No Terminal Server

- ▶ No Terminal Server Licensing
- ▶ No UDDI
- ▶ No Windows Deployment
- ▶ No Windows Media Services
- ▶ *In Management & Monitoring Tools, Only SNMP*

Guidelines for Microsoft SQL Server

- ▶ Restrict windows authentication user to access .mdf and .ldf files and assign read/write access to appropriate users.
- ▶ Use NTFS file system as it provides advanced security and recovery features.
- ▶ Rename the Windows Administrator account on the SQL Server server to discourage hackers from guessing the administrator password.
- ▶ Hide SQL Server service from appearing in the server enumeration box in Query Analyzer, using the /HIDDEN: YES switch of NET CONFIG SERVER command.
- ▶ Disable Windows guest user account on production servers.
- ▶ Setup roles in SQL Server and configure permissions for windows authentication. Take advantage of the fixed server and database roles by assigning users to the appropriate roles.
- ▶ Restrict access to the SQL logs directory.
- ▶ Secure registry by restricting access to SQL Server registry keys like `HKEY_LOCAL_MACHINE\Software\Microsoft\MSSQLServer`.
- ▶ Encrypt User Views, Stored procedure, Functions, and triggers while going live.
- ▶ If the audit feature is enabled, you can examine the audit for login failure events and look for trends to detect possible intrusion. Note that the auditing feature must be disabled before applying CU 3 MS SQL Server.

Appendix: Failover Behavior

- ▶ [Customer and Agent Behavior During Failover](#)

Customer and Agent Behavior During Failover

The following table describes the customer-side and agent-side observations for different actions during a failover. Please note that the observations are valid for both collocated and distributed ECE environments.

Action	During Failover	
	Customer-side Observations	Agent-side Observations
Customer launches chat entry point URL	<ul style="list-style-type: none"> ▶ Off-hour message is displayed. ▶ Note that after the failover is completed, the chat launches successfully. 	N/A
Customer fills out the pre-chat form and clicks the Submit button	Off-hour message is displayed.	N/A
Chat is placed in a queue and is waiting to be assigned to an agent	The customer waits for the agent assignment and the estimated wait time is displayed.	N/A
Chat is assigned to an agent and the agent has not picked the chat yet	The customer waits for the agent assignment and the estimated wait time is displayed.	Chat gets assigned to an agent.
Chat is assigned to an agent and the agent picks the chat	The customer is notified that the chat has been assigned to an agent.	Chat gets assigned to an agent.
Customer is actively engaged in a chat with an agent	The customer is able to exchange messages seamlessly with the agent.	The agent is able to chat and send messages to the customer without any interruption.
Customer initiates a chat using either Google Chrome or Microsoft Edge browser but the chat session remains idle (chat browser is not in focus or minimized) for a duration of 90 seconds to 300 seconds	<ul style="list-style-type: none"> ▶ If the Extend Session Timeout for Customers (minutes) setting is configured, the application attempts to reconnect the agent and the customer for the amount of time specified in this field and the time resets upon a successful connection. ▶ If this setting is not configured, the chat session terminates. 	<ul style="list-style-type: none"> ▶ If the Extend Session Timeout for Customers (minutes) setting is configured, the application attempts to reconnect the agent and the customer for the amount of time specified in this field and the time resets upon a successful connection. ▶ If this setting is not configured, the agent receives a message that the customer has exited the chat.
Customer abandons the chat	The chat survey page is displayed to the customer.	The agent receives a message that the customer has exited the chat.
Customer submits the chat survey	The chat survey is submitted successfully.	N/A
Customer downloads the chat transcript	The chat transcript is downloaded successfully.	N/A

Action	During Failover	
	Customer-side Observations	Agent-side Observations
Customer initiates chat through API	The following error message is displayed: Chat conversation could not be started because of none of agents are available to serve the chat.	N/A
Agent Login (Finesse/ Lite Agent)	N/A	<ul style="list-style-type: none"> ▶ The agents trying to log in during a failover see the following error message: Login timed out with Unified CCE. Please contact your administrator for verification of CTI server status, network connectivity or latency for Unified CCE connections, and running state of EAMS process and instance in Cisco. Also ensure there are email and chat licenses available for this agent to use for login. Please note that agents can successfully log in after the failover is completed. ▶ Already logged in agents do not experience any issues.
Agent Availability for Chat and Email (Finesse/ Lite Agent)	N/A	Agents are able to mark themselves as available or unavailable for Chat and Email as agent availability status is not interrupted during failover.
Agent to agent transfer of chat activity	N/A	<ul style="list-style-type: none"> ▶ The following error message is displayed: Connection to the MR PIM is no longer active or the connection is timed out. This can occur in the following situations: 1. The MR PIM component is not able to connect to EAAS instance. 2. EAAS process or instance is not in Running state. ▶ The console may fail to load the list of agents. If loaded, agents might encounter problems when selecting them. Please note that EAAS and EAMS services may take a maximum of 180 seconds post the failover to start on the secondary node, after which agents should not experience any issues.
Queue to Queue transfer of chat activity	N/A	<ul style="list-style-type: none"> ▶ The console may fail to load the list of queues. If loaded, agents might encounter problems when selecting them. Please note that EAAS and EAMS services may take a maximum of 180 seconds post the failover to start on the secondary node, after which agents should not experience any issues.

Action	During Failover	
	Customer-side Observations	Agent-side Observations
Agent picks an email activity	N/A	Agent is unable to pick the activity as queues and users are not visible. Please note that EAAS and EAMS services may take a maximum of 180 seconds post the failover to start on the secondary node, after which agents should not experience any issues.
Agent pulls an email activity	N/A	Agent is unable to pull an activity. Please note that EAAS and EAMS services may take a maximum of 180 seconds post the failover to start on the secondary node, after which agents should not experience any issues.
Agent leaves the chat before the customer	<ul style="list-style-type: none"> ▶ The customer is notified that the chat has been ended by the agent. ▶ The Submit Survey button is displayed in the Customer Chat Console. 	The chat activity gets completed.
Agent completes the chat activity	N/A	<ul style="list-style-type: none"> ▶ The console remains in a loading state for an elongated duration and the following error message is displayed: <code>Could not process the request</code>. When the agent clicks the OK button, the browser refreshes and the agent is able to complete the chat activity. ▶ The console remains in a loading state for an elongated duration and then the chat is completed after successful failover.
Agent does not pick the activity and the activity is auto-pushed	The following message is displayed to the customer: An agent will be with you shortly.	<ul style="list-style-type: none"> ▶ If Enable Auto-pushback for Chats setting is enabled, then the chat activity gets assigned to another agent and the first agent is marked as inactive. ▶ If Enable Auto-pushback for Chats setting is disabled, then the chat activity is assigned or reassigned to the agent. ▶ The activity enters the 4100 state, meaning it stays in the queue and is not assigned to any agent.
Maintenance mode is executed on the UCCE side or CTI failover is triggered on the UCCE side	The chat is abandoned from the agent's side.	<ul style="list-style-type: none"> ▶ The activities are not assigned to agents as the activity assignment does not work. ▶ Please note that restarting the EAMS service can fix this issue.