



This PDF is generated from authoritative online content, and is provided for convenience only. This PDF cannot be used for legal purposes. For authoritative understanding of what is and is not supported, always use the online content. To copy code samples, always use the online content.

Workspace Desktop Edition Deployment Guide

eServices Business Continuity with UCS 8.5

4/4/2026

eServices Business Continuity with UCS 8.5

[**Added:** 8.5.106.19] [**Modified:** 8.5.109.16]

Important

This topic covers deployments that use UCS 8.5. For deployments that use UCS 9.1, refer to [eServices Business Continuity with UCS 9.1](#).

Contents

- [1 eServices Business Continuity with UCS 8.5](#)
 - [1.1 Setting up your system](#)
 - [1.2 Provisioning eServices Disaster Recovery](#)
 - [1.3 Provisioning](#)

Workspace enables you to configure and provision a Business Continuity model by defining Load Balancing rules in a single Workspace Application object in the Configuration Layer so that a single software package can be distributed to all agents, and to define a Disaster Recovery model that switches to a peer data center without requiring agents to re-start their application.

Setting up your system

This section describes how to set up your environment to implement disaster recovery and load balancing for eServices in a high availability (HA) environment.

Physical connections

You must create an eServices server environment that is divided between two data centers that are set up to support disaster recovery.

In a standard operational situation, the load is equally distributed between the two data centers, based on static distribution rules: Each agent is provisioned to connect preferably to one data center but can also connect to the other one if the preferred data center is no longer available. Therefore, each data center must be able to support the full agent load.

Set up your data centers following this model:


- **A Proxy Tier**
 - An Application Cluster containing a pool of Interaction Server Proxies and UCS Proxies running in a N+1 failover schema (there is no primary-backup definition)
- **A Core Application Tier**
 - Contains a primary-backup Interaction Server pair
 - Contains a primary-backup UCS pair
- **A Database Tier**
 - Contains the database node which includes UCS and Interaction Server data running in *primary* or *replica* mode

Data centers run in an active/passive mode with respect to the eServices components:

- **Active**
 - The database is in *Primary* mode
 - The Interaction Server pair and the UCS pair are running and connected to the primary database
 - The Interaction Server Proxies and UCS Proxies are running and connected to the core Interaction Server and core UCS, and each is listening on its configured connection port
- **Passive**
 - The database is in *Replica* mode

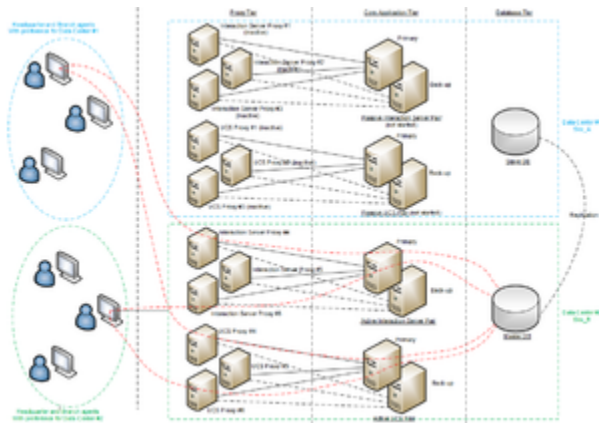
- The Interaction Server pair and the UCS pair are stopped
- The Interaction Server Proxies and UCS Proxies are running or not, connected or not, to the core Interaction Server and core UCS, and each is not listening on its configured connection port. This means that it is not possible to open any network connection to them

The following is an example of the physical network connections when the environment is running in regular operation, and data centers #1 and #2 are up and running:



Example of the physical network connections in regular operation when data centers #1 and #2 are up and running

The following is an example of the physical network connections when the environment is running in regular operation, and data center #1 is down:



Example of the physical network connections in regular operation when data center #1 is down

Provisioning eServices Disaster Recovery

The Database Tier is critical for eServices components. The disaster recovery site model that is used for SIP Server, Stat Server, and Configuration Server Proxy (refer to [Configuration of Voice and Statistic Business Continuity](#)) cannot be applied to eServices components. The advanced proxy

architecture that enables the $N+1$ model of High Availability and load balancing also does not apply to the Primary/Backup model.

eServices Disaster recovery takes advantage of the Application Cluster (a pool of Proxy servers of different types; for example, UCS Proxies and Interaction Server Proxies) in **Genesys Management Framework**. This model provisions *server pools* for the Interaction Server Proxy/UCS Proxy connections types. The Application Cluster model also enables the support of various scale factors that depend on the application type (for example, three Interaction Server Proxy nodes for the eServices agents and six UCS Proxy nodes for the eServices agents and the Voice agents). The following is an example of how an application cluster might be configured in Genesys Administrator Extension:



Application Cluster configuration in Genesys Administrator Extension

An Application Cluster can contain a set of optional Workspace specific key-value pairs which are defined in the `interaction-workspace` section. The `disaster-recovery.site` option defines the data center where the Cluster is located. The `disaster-recovery.site` option is a Key that you add in the options of another application to which Workspace connects (for example: SIP Server, Stat Server, and Config Server) to recognize it as a "peer" or a "preferred" site. See the examples below to see how this option is used.

The `disaster-recovery.eservices-site` Workspace option, which can be defined in the Application options or in the Tenant, Group, or User annex, is used to define the cluster assignment. This option defines the data center where a user must connect to enable eServices links (it is optional and taken into account dynamically).

In the example above, the Workspace Application has connections to one or several clusters. The list of connections is taken into account when the Application is restarted.

- If the Workspace `disaster-recovery.eservices-site` option is not defined, or is left empty, Workspace selects the first Cluster in the list that is available independent of the value specified for the `disaster-recovery.site` option, and uses the pool of proxies that it contains.
- If Workspace option `disaster-recovery.eservices-site` is assigned a site name, Workspace selects the first Cluster in the list that has an equivalent value specified in the `disaster-recovery.site`, and uses the pool of proxies that it contains.

This approach enables you to define different kinds of deployment models that enable Load Balancing and Disaster Recovery.

Provisioning Examples

The following tables provide administrators with an example of how an eServices Business Continuity/ Disaster Recovery environment might be set up.

Example 1: Load Balancing in one data center

Cluster Applications

	Connections	Options
Cluster eService Site_A	IxnProxy_A_1, IxnProxy_A_2, IxnProxy_A_3 UCS_Proxy_A_1, UCS_Proxy_A_2, UCS_Proxy_A_3	interaction-workspace/disaster-recovery.site=<not defined>

Notes:

- UCS_Proxy_x_y stands for UCS Proxy number y on site x.
- IxnProxy_x_y stands for Interaction Server Proxy number y on site x.
- For example, IxnProxy_A_2 is the Interaction Server Proxy #2 on Site A

Workspace Application

	Connections	Options
Workspace	Cluster eService Site_A	interaction-workspace/disaster-recovery.eservices-site=<not defined>

Example 2: Two data centers for Business Continuity, Load Balancing inside each data center

Cluster Applications

	Connections	Options
Cluster eService Site_A	IxnProxy_A_1, IxnProxy_A_2, IxnProxy_A_3 UCS_Proxy_A_1, UCS_Proxy_A_2, UCS_Proxy_A_3	interaction-workspace/disaster-recovery.site=Site_A
Cluster eService Site_B	IxnProxy_B_1, IxnProxy_B_2, IxnProxy_B_3 UCS_Proxy_B_1, UCS_Proxy_B_2, UCS_Proxy_B_3	interaction-workspace/disaster-recovery.site=Site_B

Notes:

- UCS_Proxy_x_y stands for UCS Proxy number y on site x.
- IxnProxy_x_y stands for Interaction Server Proxy number y on site x.
- For example, IxnProxy_A_2 is the Interaction Server Proxy #2 on Site A

Workspace Application

	Connections	Options before site switch-over	Options after site switch-over
Workspace	Cluster eService Site_A, Cluster eService Site_B	interaction-workspace/ disaster-recovery.eservices-site=Site_A	interaction-workspace/ disaster-recovery.eservices-site=Site_B

Runtime connection logic

[Modified: 8.5.109.16]

At runtime, Workspace selects the initial Proxy node within the appropriate pool defined for the active data center.

When Workspace needs to establish a connection to a node (UCS Proxy or Interaction Server Proxy), Workspace has the following behavior:

1. Isolate the list of Server objects of the appropriate type from the Cluster Application that has the site name that corresponds to the current agent.
2. Create a random list of those servers to load balance connections
3. Select the first server from the list
4. Attempt to connect to the selected server
5. If there is a connection failure, take the next server from the list (go to step 4)

If Workspace fails over from one proxy node to another proxy node within the same data center because the connection to a node (UCS Proxy or Interaction Server Proxy) is lost, Workspace tries to find another active node within the pool of nodes that were identified during initial node selection.

If Workspace fails over from the current active data center to the other other data center in a scenario that defines the failover by a manual update of the value of the `disaster-recovery.eservices-site` option, which is assigned to the agent through the configuration hierarchy (refer to *Administrator operations in case of data center failover*, below):

1. For each proxy type, Workspace refreshes the list of Server objects of the appropriate type from the Cluster Application that has the new site name
2. If Workspace is currently connected to UCS Proxy and/or Interaction Server Proxy node, it forces the disconnection
3. For each proxy type, Workspace applies the same random selection/round robin connection rules that it uses during initial connection.

Important

In a scenario where the Configuration Server is down during the modification of site, the following mechanism is applied when the connection of Configuration Server is back in service, Workspace:

1. retrieves the object of hierarchy defined by the `disaster-recovery.eservices-site` option
2. refreshes this object by using a request to Configuration Server
3. performs a *move site* if necessary

You can specify which eServices media Workspace should try to reconnect after reconnection to Interaction Server by using the following two options:

- `eservices.session-restore-mediatype`—Specifies which media types Workspace should attempt to reconnect.
- `eservices.session-restore-timeout`—Specifies the time, in seconds, after reconnection to Interaction Server to retrieve the ownership of interactions of media types that are specified by the value of the `eservices.session-restore-mediatype` option. The value of this option is dependent on the value that is specified for the Interaction Server `agent-session-restore-timeout` option.

When Workspace loses connection to Interaction Server or Interaction Server Proxy due to network disconnection, shutdown of Interaction Server Proxy instance, and so on, eServices interactions remain open in Workspace, but with a restricted set of controls until the connection to same or a new instance of Interaction Server or Interaction Server Proxy is restored.

The following general restrictions apply:

- Case Information is read only
- It is not possible to assign a contact to the interaction
- Disposition code is read only
- Record Information is read only

For **SMS** interactions, the following restrictions apply:

- The Transfer, Done, and Delete controls are disabled
- Send message is active only for SMS in session mode

For **email** interactions, the following restrictions apply:

- All actions except for print preview are disabled
- Agent can edit To, From, Cc, Bcc, subject, and email body in outbound email only are active

For **chat** interactions, the following restrictions apply:

- The End, Transfer, Conference, and Done (if the chat session is already ended) controls are disabled
- The Send message control is active if the chat session is still active and the Chat Server connection is still up

Administrator operations in case of data center failover

The key operation in a data center failover scenario is the transfer of mastership of the database from Site_A to Site_B.

Important

You should note that this is an intensive operation that can take a significant amount of time. It involves working with the **Database Administrator** who must initiate the procedure. Since a failover scenario can happen at any time, the Database Administrator might not be available immediately.

The following is a typical failover operation scenario:

1. A disaster is detected
2. The decision to switch over to data center 2 (Site_B) is taken
3. Stop all the *active* servers (core and proxies) that might still be running on Site_A
4. eServices agent activity is suspended
5. The database administrator must make the DB of Site_B the *primary*
6. Start the core UCS pair and Interaction Server pair of Site_B
7. Start the Interaction Server Proxy and UCS Proxy instances of Site_B
8. Change the value of the `disaster-recovery.eservices-site` option of the agent at the appropriate hierarchy level (this can be at the Application level if everyone shares the same data center at the same time) so that it points to Site_B
9. Workspace instances automatically reconnect to UCS Proxies and Interaction Server Proxies of Site_B and agents can resume their eServices activity

Important

The transition from passive to active of the *peer* eServices data-center takes a significant amount of time, during which the eServices capabilities are not usable by agents.

For blended agents (voice + at least one eServices channel)

The data center failover logic between data centers is not the same for SIP Voice/IM channels and eServices channels, nor is the login model the same. However, Workspace and Stat Server maintain a composite connection between these two solutions known as "blended agents". In particular, the Stat Server composite representation requires a single agent-place virtual runtime link.

In SIP Business Continuity:

- The configuration model is based on a pair of DNs (one for each data center)
- In typical deployment model, each DN belongs to its dedicated Place (preferred DN in Place 1, peer DN in Place 1_DR). To have the two DNs configured in one place, they must be of type ACD Position.
- Agents logon one DN at a time, according to the site currently defined as the *active preferred*, according to the agent configuration; this can change dynamically at runtime according to site failure detection.
- At runtime, Stat Server detects that an agent dynamically moves from the default Place to peer Place.

In Interaction Server:

- An agent logs in to a Place at the beginning of a session. This place remains the same during the session, even if Workspace has to connect to another Interaction Server Proxy, from the local data center or the remote data center.

Therefore, there are situations when an agent might be virtually logged in to two distinct Places, one

for voice and one for eServices, which is not a supported situation for Stat Server.

One possible approach to resolve this situation is to use the Single Place configuration model for the Voice DNs.

Provisioning

You can configure eService Cluster and eServices Business Continuity together or independently, depending on the design of your eServices architecture. Use the following options to enable Workspace to connect correctly to your environment.

Cluster provisioning

The following options enable you to specify how the connection to any node of your cluster should be set-up and how the transition between nodes should behave.

Warm standby

These two options enable you to specify the time interval for reconnection. Refer to the configuration option reference for details about how to configure these options:

- warm-standby.reconnection-random-delay-range
- warm-standby.retry-delay

Addp

You can define the addp parameters for all applications of the cluster from the Workspace application connections table. These parameters can be overridden in each server application.



The screenshot shows a table with columns: Name, Enabled, Connection Protocol, Local, Remote, and Transferable. The table lists several applications with their respective configurations.

Name	Enabled	Connection Protocol	Local	Remote	Transferable
addp	<input checked="" type="checkbox"/>	addp	Y	Y	Transfer/Transfer Off
addp	<input checked="" type="checkbox"/>	addp	Y	Y	Transfer/Transfer Off
addp	<input checked="" type="checkbox"/>	addp	Y	Y	Transfer/Transfer Off
addp	<input checked="" type="checkbox"/>	addp	Y	Y	Transfer/Transfer Off
addp	<input checked="" type="checkbox"/>	addp	Y	Y	Transfer/Transfer Off
addp	<input checked="" type="checkbox"/>	addp	Y	Y	Transfer/Transfer Off
addp	<input checked="" type="checkbox"/>	addp	Y	Y	Transfer/Transfer Off

Example of Workspace in a cluster configuration for Disaster Recovery

ip-version

You can define the ip-version option for all applications of the cluster from the Workspace application connections table. ip-version can be overridden in each server application options. In this example, the ip-version of all applications of the cluster is 4,6:

The screenshot shows a 'Connection' dialog box with the following fields and values:

- Server: Cluster eService Site_A
- Port ID: default
- Connection Protocol: addp
- Local Timeout: 5
- Remote Timeout: 9
- Trace Mode: Trace Is Turned Off
- Transport Protocol Parameters: ip-version=4,6
- Application Parameters: (empty)

Setting ip-version preferences for Workspace connections for Disaster Recovery

eServices Disaster Recovery provisioning

These two options enable you to configure the data center infrastructure that is applicable to eServices components:

- disaster-recovery.eservices-site
- disaster-recovery.eservices-random-delay-range