FORESCOUT

Forescout

Resiliency and Recovery Solutions

User Guide

Version 8.1

Contact Information

Forescout Technologies, Inc.

190 West Tasman Drive

San Jose, CA 95134 USA

https://www.Forescout.com/support/

Toll-Free (US): 1.866.377.8771

Tel (Intl): 1.408.213.3191

Support: 1.708.237.6591

About the Documentation

- Refer to the Resources page on the Forescout website for additional technical documentation: https://www.Forescout.com/company/resources/
- Have feedback or questions? Write to us at documentation@Forescout.com

Legal Notice

2019-03-17 18:34

Table of Contents

About Forescout Resiliency and Recovery Solutions

Forescout resiliency and recovery solutions provide support for availability of Forescout services to minimize down-time in cases of system failure.

Solutions for Appliances

The following resiliency solutions are available for Forescout Appliances:

- Failover Clustering
  - Implemented with clusters of Appliances.
  - Redundancy is achieved by defining clusters of Appliances that can automatically take over discovery, assessment and control in case of single or multiple Appliance failure within the cluster/s.
  - Workload is balanced among the Appliances in the cluster/s after failover.
- High Availability Pairing
  - Implemented in pairs of two Appliances.
  - Redundancy is achieved by assigning an Active node and a Standby node. The Standby node automatically takes over discovery, assessment and control in case the Active node fails.
  - The two nodes are synchronized by a redundant pair of directly interconnecting cables.

For a comparison of the above solutions, and to learn more about which solution is appropriate for your deployment, see Comparison of Resiliency Solutions for Appliances.

Solutions for the Enterprise Manager

The following resiliency and recovery solutions are available for the Enterprise Manager:

- High Availability Pairing
  - Implemented in pairs of two Enterprise Managers.
  - Redundancy is achieved by assigning an Active node and a Standby node. The Standby node automatically takes over discovery, assessment and control in case the Active node fails.
  - The two nodes are synchronized by a redundant pair of directly interconnecting cables.
- Disaster Recovery for Enterprise Manager
  - Implemented in pairs of two Enterprise Managers.
  - Redundancy is achieved by defining a Recovery Enterprise Manager that is manually triggered to take over from an Enterprise Manager that is no longer functioning as a result of, for example, a disaster.
  - The Enterprise Manager and the Recovery Enterprise Manager are synchronized by communication performed on port 13000/TCP.

The Failover Clustering solution does not support Enterprise Manager failure scenarios.

See the Glossary for descriptions of terms related to Forescout resiliency and recovery solutions.

Comparison of Resiliency Solutions for Appliances

To learn about how to choose the right solution for your deployment, and to see how these solutions answer the needs of a typical deployment, read the following sections:

- Choosing the Right Solution for Your Deployment
- Deployment Example: Network Architecture
- Deployment Example: Addressing Your Needs

The table below compares the High Availability Pairing and Failover Clustering solutions.

	High Availability Pairing	Failover Clustering
Node configuration model	Active/Standby	All Appliances are active
Allows distribution across geographic locations	No	Yes
Failover detection mechanism	Active and Standby nodes are regularly monitored to detect their operational state.	Enterprise Manager cannot reach the Appliance through port 13000.
Failover triggers	Active Appliance outage: power off, reboot, etc. *** Service down does NOT trigger failover	Appliance outage at the TCP stack level *** Service down does NOT trigger failover
Failover triggers Forescout service restart	Yes. Affects overall failover time (until Standby becomes Active)	No
Total failover time	From a few to over ten minutes (due to Forescout service start).	Configured Failover Detection Time (default and minimum of 3 minutes) plus up to 30 additional seconds.

	High Availability Pairing	Failover Clustering
Continuity for all endpoint types?	Yes	Endpoint, switch and wireless device visibility assessment and enforcement are supported. See Continuity of Endpoint, Switch and Wireless Device Visibility, Information and Enforcement. For endpoints and features that are not currently supported, see Features Not Currently Supported.
Network requirements	Requires redundant pair of interconnecting cables to sync the 2 Appliances.	Yes. See Network Deployment Requirements.
Requires dedicated setup and configuration during Appliance installation	Yes	No
Requires available capacity from active Appliance/s	No	Yes. See Distribution of Workload upon Failover.
Allows using a combination of physical and virtual Appliances	No	Yes
Allows using Appliances with different models	No	Yes

Choosing the Right Solution for Your Deployment

The table below describes when to use each solution for a Standalone or group of Appliances within your deployment, based on the details of your deployment and/or business needs. You can also use both solutions in parallel.

To see the network architecture of a sample deployment, and how each solution can address the needs of this deployment, read the following sections:

- Deployment Example: Network Architecture
- Deployment Example: Addressing Your Needs

High Availability Pairing	Failover Clustering
Dedicated Appliances: ▪ Appliances with a Connecting CounterACT Device (Focal Appliance) serving as a proxy to a 3rd party server. ▪ Appliances that need to be configured on a 3rd party server. For example: - Appliances that receive SNMP traps from switches and controllers - Appliances that monitor a network SPAN port to see network traffic ▪ Appliances where plugins or features are configured individually per Appliance instead of globally for all Appliances.	Appliances can access network devices and endpoints managed by other Appliances: ▪ Appliances at the same location (e.g. the data center) ▪ Appliances from larger site can access network devices and endpoints from smaller site ▪ Appliances behind load balancers
Isolated Appliances: ▪ Appliance at a remote site with no access to other parts of the network that it is not handling.	

Deployment Example: Network Architecture

The diagram below shows an example of a Forescout deployment with the following components:

- One data center and two sites, Site A and Site B
- ***Data center***.
  - Focal Appliance connected to, for example, a SIEM server
- ***Site A***. Appliances in the data center can access some of the site's traffic. This site includes, for example:
  - A switch that the data center can access
  - A firewall that sends traffic to a load balancer in site A and is also configured to send traffic to a load balancer in the data center as a secondary connection in case the primary fails
  - An 802.1X NAS device configured with secondary/tertiary RADIUS servers
  - An Appliance monitoring a network SPAN port
- ***Site B***. Appliances in the data center cannot access the site's endpoint traffic or its network devices. This site includes, for example:
  - An Appliance receiving SNMP traps or DHCP helpers from a switch
  - An Appliance managing endpoints and switch/wireless devices/access points

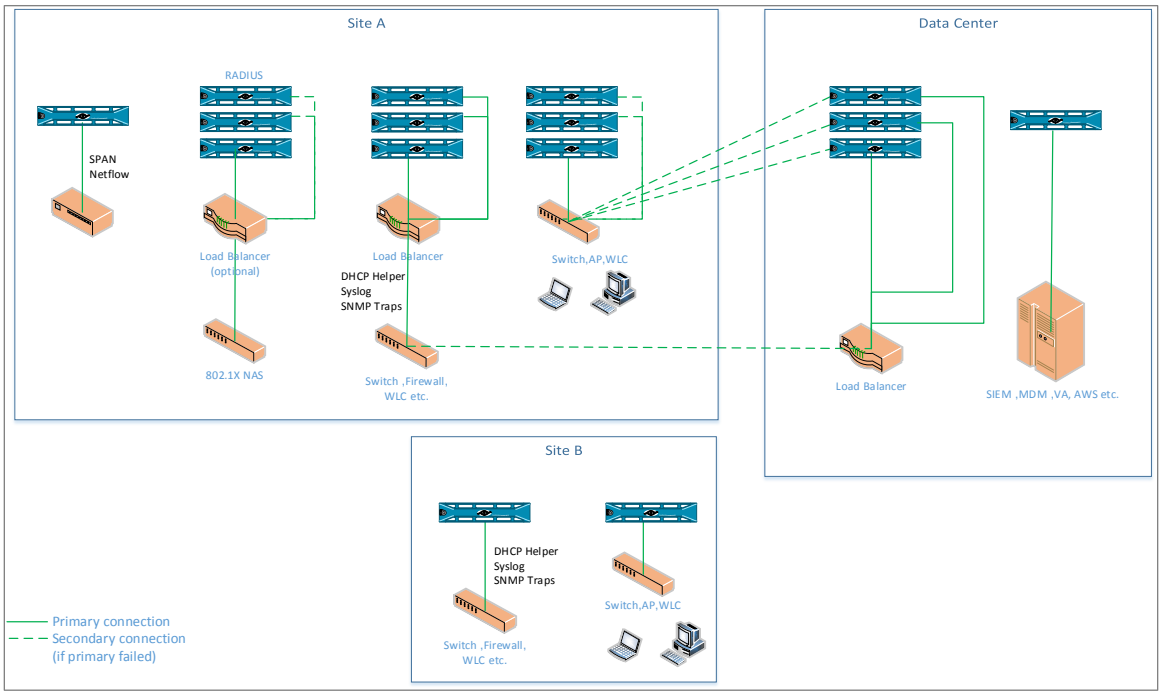

Deployment Example: Addressing Your Needs

The diagram below shows how each Forescout Resiliency solution addresses the various needs of the example deployment described above.

- All dedicated Appliances require a High Availability pair
  - The Focal Appliance in the data center
  - The Appliance monitoring a network SPAN port in Site A
- Isolated Appliances require a High Availability pair
  - The Appliances in Site B, where other Appliances cannot access the traffic
- All other Appliances in the data center and Site A should use Failover Clustering, since these Appliances can take over for one another in case of failure.

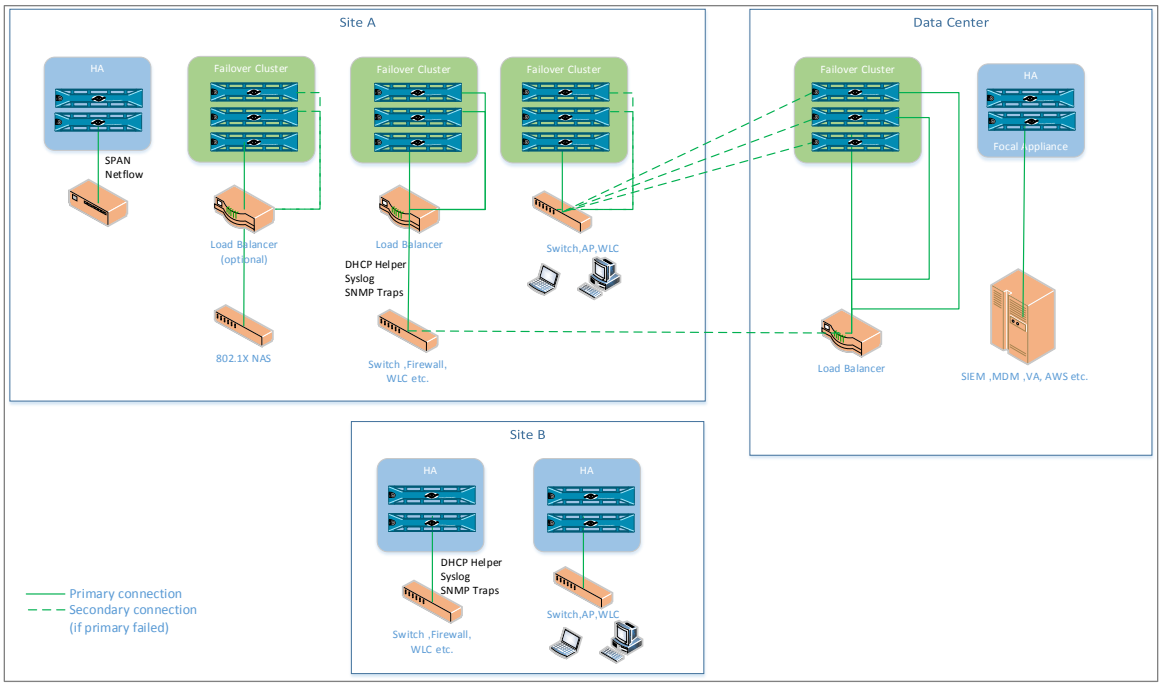
Site A
HA
SPAN
Netflow
Failover Cluster
Load Balancer
(optional)
802.1X NAS
Failover Cluster
Load Balancer
DHCP Helper
Syslog
SNMP Traps
Switch ,Firewall,
WLC etc.
Failover Cluster
Switch,AP,WLC
Data Center
Failover Cluster
Load Balancer
HA
Focal Appliance
SIEM ,MDM ,VA, AWS etc.
Site B
HA
DHCP Helper
Syslog
SNMP Traps
Switch ,Firewall,
WLC etc.
HA
Switch,AP,WLC
Primary connection
Secondary connection
(if primary failed)

Failover Clustering

Forescout deployments (e.g. centralized or hybrid) require resiliency within and across sites and geographic locations. Each site (e.g. data center) can contain many Appliances. *Failover Clustering* provides a resiliency solution in the event that an Appliance, a number of Appliances or an entire site fails. This means that the workload handled by the failed Appliance/s will automatically be transferred to other functioning Appliances with free capacity.

Failover Clustering enhances reliability by ensuring that you can:

- Withstand either single or multiple points of failure in the event of resource failure, for example, in case of a power outage.
- Maintain Forescout service continuity with no need for manual intervention.

The Failover Clustering solution offers a deployment architecture that uses fewer idle standby resources, resulting in reduced cost, complexity and power consumption. Additionally, in the event of an Appliance failure, the workload is balanced among Appliances so as to avoid overloading the recipient Appliances.

What to Do

To work with Failover Clustering, perform the following:

- Verify that you have met requirements. See Requirements.
- Understand how failover works. See About Failover and Failback.
- Acquire a license to enable viewing and using the feature in the Console. See Begin Working with Failover Clustering.
- Define clusters of logically connected or geographically close Appliances that will provide failover support to other Appliances. See Configuring Failover.

In addition, you can also perform the following to better understand and work with failover clusters:

- Understand how Appliance failover assignments are distributed upon failover. See Distribution of Workload upon Failover.
- Identify Appliance capacity and understand what happens to endpoints that exceed this capacity. See Handling Endpoints that Exceed Capacity.
- View endpoint, switch device, WLAN device and Appliance information related to failover. See Viewing Failover Information in the Console.
- Track failover activity in the Audit Trail or Event Viewer. See Tracking Failover Activity.

Requirements

Verify that the following software and network requirements are met and that Console users have the necessary permissions.

- Forescout Software Requirements

- Network Deployment Requirements
- Console User Permissions

Forescout Software Requirements

- A license to work with Failover Clustering. The required license depends on which licensing mode your deployment is using. See Begin Working with Failover Clustering for more information.
- Forescout version 8.1
- Network Module 1.1 with the following components running:
  - Wireless Plugin
  - Switch Plugin
- If you are working with Windows, Linux or macOS/OS X endpoints, you need Endpoint Module 1.1 with the following components running:
  - HPS Inspection Engine
  - Linux Plugin
  - OS X Plugin
- All Appliances participating in failover must have uniform configuration settings applied. Plugins or features cannot be configured individually per Appliance. See Choosing the Right Solution for Your Deployment for more information.

Network Deployment Requirements

- A network infrastructure that enables rerouting or load balancing traffic (e.g. SPAN, SNMP traps from network devices) so that potential recipient Appliances can see and handle endpoints and network devices, in case of an Appliance or site failure.
- This feature can be enabled on *Appliances* that are connected to the same Enterprise Manager, but ***not*** on the following CounterACT devices:
  - A single, standalone Appliance that is not connected to an Enterprise Manager.
  - An Enterprise Manager.
- Appliances that can potentially manage network devices as a result of a failover must have network access to those switches.
- Make sure that all Active Directory services can be accessed by all Appliances in the failover cluster:
  - If the DNS Detection checkbox is selected for Appliances in the failover cluster, this happens automatically.

    This checkbox is configured when defining User Directory servers in the User Directory Plugin.
  - Otherwise, make sure that this is configured manually by adding the relevant server IP addresses in the User Directory Plugin.

Refer to the *Forescout User Directory Plugin Configuration Guide* for more information.

Console User Permissions

In order to use and configure Failover Clustering in the Console, users must have *Update* access to the following Console permissions:

- CounterACT Appliance Control
- Multiple CounterACT Appliance Management

To access Permission and Scope options:

1. Select **Options** from the Tools menu and then select **Console User Profiles**.
2. Select **Add** to create a new user/group or **Edit** to change an existing one.
3. In the **Permissions** section, select the relevant permissions listed above.

Refer to the section on permission and scope options in the *Forescout Administration Guide* for more information. See Additional Forescout Documentation.

About Failover and Failback

Failover occurs after an Appliance fails, and the range of endpoints and network devices handled by the Appliance is distributed to one or more *recipient Appliances* for continued handling. A recipient Appliance is a functioning Appliance managed by the same Enterprise Manager as the failed Appliance. The range of transferred endpoints and network devices is the *Failover Assignment*, which will now be handled by the recipient Appliance in addition to the *Original Assignment*, which is the range of endpoints and network devices originally configured for the recipient Appliance.

The Forescout platform considers an Appliance failed when attempts by the Enterprise Manager to connect to it through port 13000 result in connection time outs for a defined period of time. This means that the Enterprise Manager cannot connect to the Appliance at the operating system level for this period of time. This may occur for a variety of reasons, for example, if the Appliance is unplugged or in the event of a natural disaster. See Configuring Failover for more information about the failover detection time period.

If the Forescout *service* is down on the Appliance, for example, during an upgrade, the Appliance is *not* considered failed and failover will not occur, even if the service is down for longer than the detection time period.

Endpoints and devices generally fail over to recipient Appliances immediately after the detection time has passed. The total failover time (failover detection time + time to failover to recipient Appliances) may depend on several additional factors, including the performance capabilities of the Appliances participating in failover.

Failback

Failback occurs after a failed Appliance reconnects to the Enterprise Manager, and regains control of the endpoints and network devices previously distributed to recipient Appliances.

If a switch device is transferred to a recipient Appliance due to failover, and then the recipient Appliance auto-discovers additional switches, only the original switch devices will be transferred back to the original Appliance after failback. The auto-discovered switch devices will remain handled by the recipient Appliance.

Continuity of Endpoint, Switch and Wireless Device Visibility, Information and Enforcement

When transferred to a recipient Appliance, endpoints (both wired and wireless) and switch and wireless devices continue to be visible in the Console and most previously discovered data is preserved. Properties resolved on endpoints, and actions previously performed on endpoints by the original Appliance are transferred with the endpoints and will continue to apply.

In addition, endpoints are rechecked for policy evaluation after being transferred to the recipient Appliance. This ensures that endpoint information is preserved and that enforcement continues functioning as in a regular, non-failover scenario.

Data transferred with the endpoint includes:

- Matched policies and sub-rules
- Manual continuous actions applied on the endpoint and information relevant to them
- Events and properties that cannot be relearned and are used in policies

The Authentication, Signed In Status *property is not transferred.*

There might be occurrences in which an applied, switch restrict action is temporarily cancelled by the Switch Plugin. However, as soon as the Forescout platform re-discovers and re-evaluates the affected endpoints, the Switch Plugin re-applies the action that was in effect on the endpoints, as necessary.

Switch Plugin Auto-Discovery

During a failover scenario, the Switch Plugin auto-discovery feature is affected as follows:

- ***Managed Switch Failover to a Recipient Appliance***. All neighboring switch devices that the Switch Plugin, on the *recipient* Appliance, learns about from the auto-discovery efforts of a failed-over managed switch device remain managed by the Switch Plugin on the *recipient* Appliance.
- ***Managed Switch Failback to the Original Appliance***. When failback occurs, only the failed-over managed switch device is re-assigned back to the *original* Appliance.

Failover Clusters

A *failover fluster* defines a group of geographically/logically connected Appliances, such as those at a data center. A failover cluster is the basic unit of failover. Appliances within a failover cluster that fail, can fail over to other Appliances in the cluster. For example, if you create failover cluster *New York*, an Appliance within the cluster that fails can fail over to the other Appliances in the cluster.

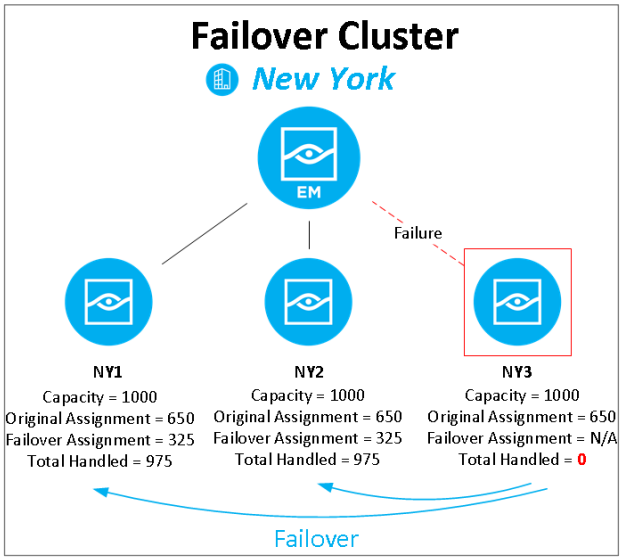

Failover Scope

In addition to supporting failover within a single failover cluster, you can configure a *failover scope* for a cluster, consisting of multiple clusters, to allow failover across clusters. This allows distribution of endpoint, and switch and wireless device assignments to a wider range of Appliances in the case of Appliance failure, and allows you to protect and back up an entire site to other sites in a different geographic location.

When an Appliance fails, its handled endpoints and network devices are first distributed to other Appliances that have free capacity within the same failover cluster. When no Appliances in the cluster have free capacity, assignments are distributed to Appliances in other clusters in the failover scope.

For example, you can add failover cluster *New Jersey* to the failover scope of the *New York* cluster. The scope will consist of both *New Jersey* and *New York*. If an Appliance from New York fails, its handled assignments are distributed across Appliances from New York and New Jersey, with a preference to local, intra-cluster Appliances from New York. If the entire New York site fails, its assignments are distributed across Appliances from New Jersey.

When a failover occurs, the endpoint capacity of the recipient Appliance/s may be exceeded. See Handling Endpoints that Exceed Capacity for more information.

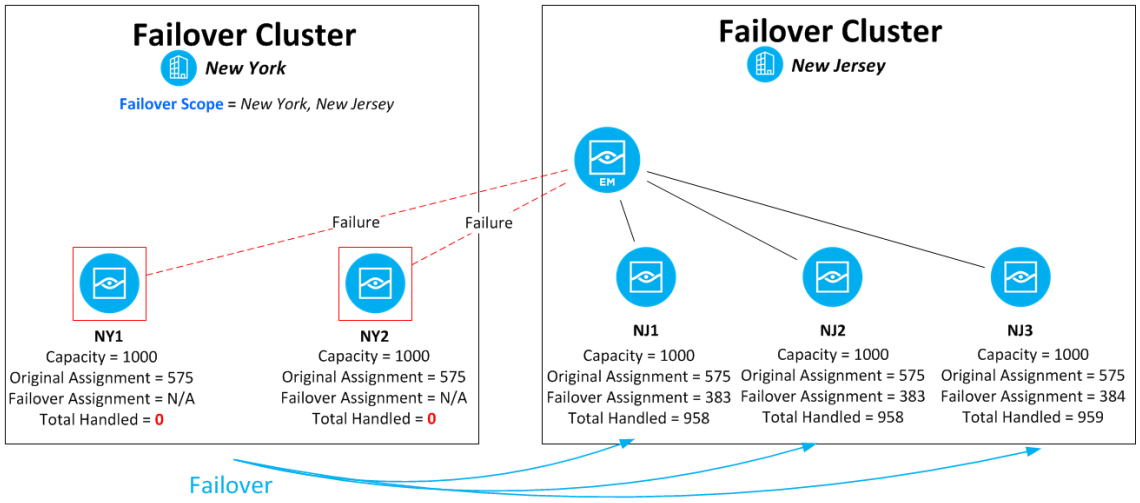

A single cluster can be in the scope of more than one other cluster. Scope definitions are per-cluster and are not bidirectional. So, for example, adding the failover cluster *New Jersey* to the failover scope of the *New York* cluster does not automatically add *New York* to the scope of the *New Jersey* cluster.

See Configuring Failover for information about configuring failover clusters.

Failover Cluster Folder Type

A *failover cluster* folder type in the CounterACT Devices > IP Assignment and Failover pane allows users to configure failover.

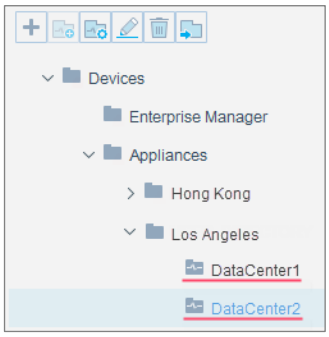

Failover clusters integrate with the existing Appliance folders feature, which lets you organize your network Appliances into logical groups in a tree structure, helping you to create a visual representation of your network Appliance deployment. Refer to the chapter on Managing Groups of Appliances in the *Forescout Administration Guide* for information about organizing Appliances using Appliance folders. See Additional Forescout Documentation for information on how to access the guide.

Begin Working with Failover Clustering

To begin working with Failover Clustering, you need a license for the feature. The license required depends on which licensing mode your deployment is using:

- Per-Appliance Licensing Requirements - Requires installation of Failover Clustering Module, version 1.0.0.
- Flexx Licensing Requirements – Requires a Forescout eyeRecover license.

To identify your licensing mode:

- From the Console, select **Help > About Forescout**.

Per-Appliance Licensing Requirements

If your deployment is using per-Appliance licensing, you must install the *Failover Clustering Module*, which enables viewing and usage of the feature in the Console. *Until this module is installed with a valid license, interface elements related to the feature will not appear in the Console.* See Install the Module.

When installing the module you are provided with a 90-day demo license. If you would like to continue exploring the feature before purchasing a permanent license, you can request a demo license extension. Consult with your Forescout representative before requesting the extension. You will receive email notification and alerts at the Console before the demo period expires.

When the demo period expires, you will be required to purchase a permanent license. *In order to continue working with the feature, you must purchase the license*.

Demo license extension requests and permanent license requests are made from the Forescout Console.

Requesting a License

When requesting a demo license extension or permanent license, you are asked to provide the device *capacity* requirements. This is the number of endpoints/network devices that you want this license to handle. You must define the number of endpoints/network devices that you want covered by failover clusters. You can request a license that handles more to ensure that you are licensed for support on additional endpoints/network devices as your deployment grows.

Enter this number in the **Devices** pane of the Module License Request wizard, in the Modules pane of the Forescout Console.

Refer to the *Forescout Administration Guide* for information on requesting a permanent license or a demo license extension. You can also contact your Forescout sales representative for more information.

Install the Module

The installation package is in the form of a Forescout module.

To install the module:

1. Navigate to the Product Updates Portal, Forescout Modules page and download the module **`.fpi`** file.
2. Save the file to the machine where the Forescout Console is installed.
3. Log into the Forescout Console and select **Options** from the **Tools** menu.
4. Select **Modules**. The Modules pane opens.

5. Select **Install**. The Open dialog box opens.
6. Browse to and select the saved module **`.fpi`** file.
7. Select **Install**.
8. If you have not yet purchased a permanent license, a message appears indicating that the module will be installed with a demo license. Select **Yes** and then select **Install**.
9. An installation or upgrade information dialog box and an End User License Agreement will open. Accept the agreement to proceed with the installation.
10. When the installation completes, select **Close**. The module is displayed in the Modules pane. The **Module Status** column indicates the status of your license. See License Requirements or the *Forescout Administration Guide* for details on requesting a permanent license or a demo license extension.

Flexx Licensing Requirements

If your deployment is using Flexx licensing, you must acquire a valid *Forescout eyeRecover* license.

The Forescout eyeRecover license supports:

- Failover Clustering
- High Availability Pairing for Appliances

  High Availability Pairing for Enterprise Manager is supported by the Forescout eyeSight license.

If you do not have a valid eyeRecover license, Console users will not be able to:

- Enable failover clusters
- Configure failover detection time for failover clusters
- Define a failover scope for failover clusters

Refer to the section on License Enforcement in the *Forescout Administration Guide* for more information.

Activating the License

When you set up your Forescout deployment, you must activate a license file containing valid licenses for each feature you want to work with in your deployment, including the Forescout eyeRecover license. After the initial license file has been activated, you can update the file to include the eyeRecover license if it was not included. For more information on obtaining licenses, contact your Forescout representative. Refer to the *Forescout Flexx Licensing How-to Guide* for information on activating a new license file or updating an existing one. See Additional Forescout Documentation for information on how to access the guide.

No demo license is automatically installed during system installation.

License entitlements are managed in the Forescout Customer Portal. After an entitlement has been allocated to a deployment, you can activate or update the relevant licenses for the deployment in the Console.

An eyeRecover license has an associated capacity, indicating the number of endpoints the license can handle. This number should include the total number of unique endpoints that are handled by either failover clusters or High Availability pairs. An endpoint handled by both of the above-mentioned resiliency solutions is counted as a single endpoint for licensing purposes.

Configuring Failover

To configure failover in your environment, perform the following tasks:

- Define a Failover Cluster
- View Information about Failover Status

Define a Failover Cluster

Forescout uses a folder tree to organize Appliances. A failover cluster is a folder in this tree to which specific settings have been applied. For details about working with the Appliance folder tree, refer to the *Forescout Administration Guide*.

To create a failover cluster, follow this general procedure:

1. Create folder in the tree, or select an existing folder.
2. Identify segments of the Forescout Internal Network that should participate in failover, and assign these segments to the folder.
3. Identify Appliances that you want to support failover between these segments, and place these Appliances in the folder.
4. Define the folder as a failover cluster.

To define a Failover Cluster:

1. Select **Options** from the Tools menu.
2. Select **CounterACT Devices > IP Assignment and Failover**.
3. Do one of the following:
   - Select a folder in the Appliances tree that you want to configure as a failover cluster. Rename the folder to indicate it is a failover cluster.
   - Right-click a node and create a new folder where you want the failover cluster. The folder name should indicate it is a failover cluster.

   If you define a failover cluster folder under a parent folder that is itself a failover cluster, the two failover clusters will function independently.

4. Populate the failover cluster folder with Appliances that will support failover:
   a. Select the root-level Appliances node of the folder tree. The table displays all Appliances defined in the Console, and shows the path to their location in the tree.

   *Verify that the **Show child folder information** option is enabled.*

b. Use table search and sorting tools to find the Appliances you want to place in the failover cluster. You can select folders in the tree to work with sub-sections of the tree.

c. Select the Appliances and select **Move**.

d. In the dialog box, select the failover cluster folder and select **OK**. The Appliances are placed in the failover cluster folder.

If you are adapting an existing folder to become a failover cluster, remove from the folder any Appliances that should not participate in failover.

5. Clear statically assigned segments from Appliances in the failover cluster folder. Appliances in the failover cluster support only the network segments assigned to the folder. They cannot support individually assigned segments.

Segments listed in italic font are assigned to the parent folder; segments listed in plain font are statically assigned to the Appliance itself.

a. In the Appliance folders tree, select the failover cluster folder. The table shows the Appliances you placed in the folder.

b. Verify that no Internal Network segments are assigned to an Appliance: the Assigned Segments and IP Addresses fields of the table should only list segments assigned to the folder (shown in grey italic text).

6. Assign Internal Network segments that will participate in failover:

a. Select the folder and select **Assign Segments**  .

b. Specify the segments you want to participate in failover.

c. If you are adapting an existing folder to become a failover cluster, remove segments that should not participate in failover.

d. Select **OK**. The specified segments are assigned to the folder.

7. Select the prepared folder and then select **Configure Failover**  . The Configure Failover Cluster dialog box opens.

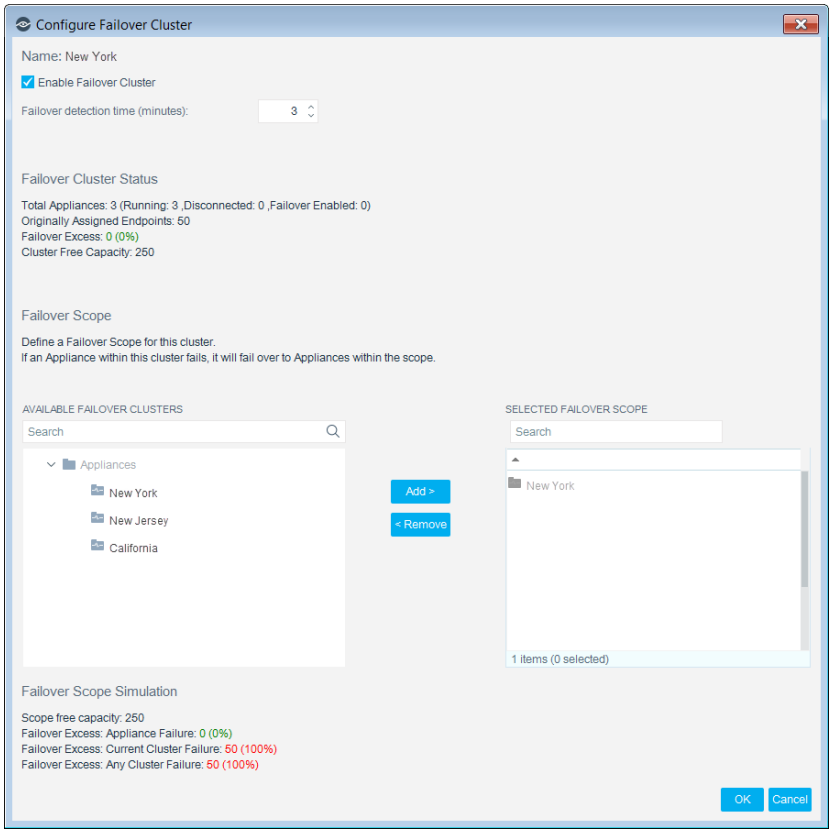

8. Select **Enable Failover Cluster** to enable failover capabilities for all Appliances in the cluster. Disabling a cluster that is in the failover scope of one or more other clusters will remove it from the scope of these clusters.
9. Select a **Failover detection time**, in minutes. The Failover detection time is the time it takes to detect that an Appliance has failed. The minimum and default time period is 3 minutes, to ensure that any disconnection from the Enterprise Manager is not temporary and to reduce false positive failovers.

   Failure detection starts once the configured number of minutes has passed, and can take up to an additional 30 seconds until failover occurs. For example, if you configured a failover detection time of 5 minutes, it may take up to 5 minutes and 30 seconds before the Appliance is considered failed.

   Endpoints and network devices generally fail over to recipient Appliances immediately after the detection time has passed. The total failover time (failover detection time + time to failover to recipient Appliances) may depend on several additional factors, including the performance capabilities of the Appliances participating in failover.
10. Define a failover scope for this cluster by assigning one or more clusters whose Appliances will serve as failover recipients. Select a folder from the Available Failover Clusters list and select **Add** to move it to the Selected failover scope list. You can add multiple clusters. If you add an Appliance folder to the scope, it is automatically converted to a failover cluster.

If an Appliance within this failover cluster fails, the Appliance will fail over to Appliances within the failover scope, with a preference for local, intra-cluster Appliances.

11. Select **OK**. The folder is defined as a failover cluster.

If you make other configuration changes to this folder in the IP Assignment and Failover pane, these changes are implemented without violating the rules of the failover cluster. For example, Appliances with statically assigned segments cannot be part of a failover cluster. If you move an Appliance with a statically assigned segment to the failover cluster folder, the Appliance's new folder location is accepted, but the Appliance is excluded from failover behavior.

View Failover Cluster Information

You can view the following information about the cluster in the Configure Failover Cluster dialog box. Information is updated as you change the cluster configuration.

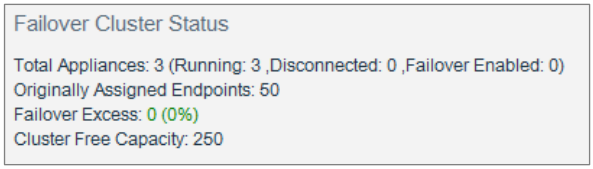

Total Appliances. Number of Appliances in the selected cluster, including how many are running (not failed), temporarily disconnected (not failed) and have failover enabled.

Originally Assigned Endpoints. Number of endpoints originally assigned to Appliances in the cluster, excluding any endpoints that were assigned later as a result of failover.

Failover Excess. Number of endpoints that exceed Appliance capacity, and the percentage of such endpoints out of the number of originally assigned endpoints.

Cluster Free Capacity. Number of endpoints that are currently available for the failover cluster to handle.

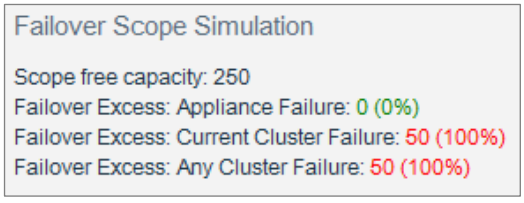

Scope Free Capacity. Number of endpoints that are currently available for Appliances in the failover scope to handle.

Failover Excess: Appliance Failure. Simulation of the highest number of endpoints that will exceed Appliance capacity if any single Appliance in the scope fails.

Failover Excess: Current Cluster Failure. Simulation of the number of endpoints that will exceed Appliance capacity if the entire current cluster fails.

Failover Excess: Any Cluster Failure. Simulation of the highest number of endpoints that will exceed Appliance capacity if an entire cluster in the scope fails.

View Information about Failover Status

You can view updated information about the status of failover clusters in the *Failover Status* table.

To view failover status information:

1. Select **Options** from the Tools menu.
2. In the Options list, expand **CounterACT Devices** and select **Failover Status**. The Failover Status Table opens.

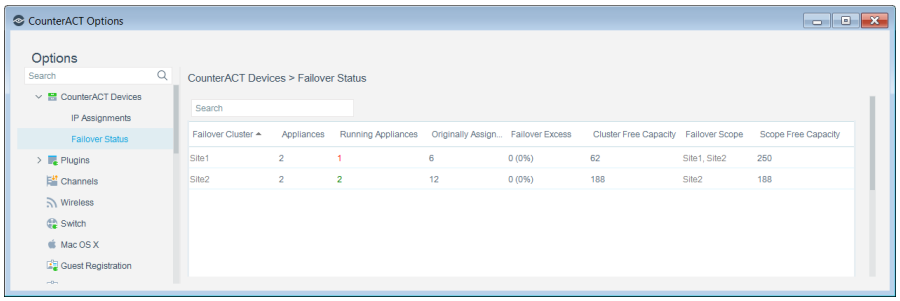

You can add or remove columns to view specific information. Right-click a column and select **Remove Column**.

Failover Status Table

The Failover Status table contains the following information:

Information will not be available for Appliances that are failed or do not have failover enabled.

- ***Failover Cluster.*** Name of the cluster.
- ***Appliances.*** Number of Appliances in the selected cluster.
- ***Running Appliances.*** Number of Appliances in the selected cluster that are running and not in a state of failover. If there are any temporarily disconnected (but not failed) Appliances in the failover cluster, the number is indicated in red.

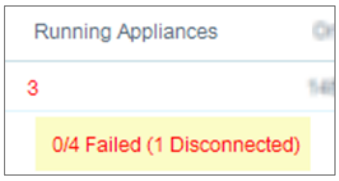

- ***Originally Assigned Endpoints***. Number of endpoints originally assigned to Appliances in the cluster, excluding any endpoints that were assigned later as a result of failover.

- ***Failover Excess***. Number of endpoints that exceed Appliance capacity, and the percentage of such endpoints out of the number of originally assigned endpoints.
- ***Cluster Free Capacity***. Number of endpoints that are currently available for the failover cluster to handle.
- ***Failover Scope***. Scope of clusters that back up failed Appliances within the selected cluster.
- ***Scope Free Capacity***. Number of endpoints that are currently available for Appliances in the failover scope to handle.

The following information, not shown by default, can be displayed in the table.

- ***Failover-Enabled Appliances.*** Number of Appliances in the selected failover cluster that have failover enabled.
- ***Failover Excess: Appliance Failure***. Simulation of the highest number of endpoints that will exceed Appliance capacity if any single Appliance in the scope fails.
- ***Failover Excess: Current Cluster Failure***. Simulation of the number of endpoints that will exceed Appliance capacity if the entire current cluster fails.
- ***Failover Excess: Any Cluster Failure***. Simulation of the highest number of endpoints that will exceed Appliance capacity if an entire cluster in the scope fails.

Performing Manual Failover

You can manually fail over one or more Appliances. You may want to perform manual failover if, for example, you want to continue to handle endpoints while maintenance work is being performed on an Appliance. When finished, you can manually fail back the Appliance.

The following limitations apply to an Appliance that is manually failed over:

- Failover cannot be disabled for the Appliance
- The Appliance cannot be moved to a different folder
- Plugin functionality is stopped, and plugins cannot be manually started

To manually fail over an Appliance:

1. Select **Options** from the Tools menu.
2. Select **CounterACT Devices**.
3. Select the Appliance/s you want to manually failover.
4. Right-click and select **Failover > Manually Fail Over**.
5. Select **Yes** to confirm.

   To manually fail back the Appliance, right-click the Appliance/s and select **Failover > Manually Fail Back**.

Distribution of Workload upon Failover

When Appliance/s fail, the workload of the Appliance is distributed using an internal mechanism across the failover cluster/scope such that the recipient Appliance/s can best handle the newly acquired failover assignment. This is done so that the recipient Appliance/s are not overburdened or unable to handle the additional devices. A single IP Range assignment may be distributed across multiple Appliances.

Endpoints and network devices are first distributed to other Appliances that have free capacity within the same failover cluster. When no Appliances in the cluster have free capacity, assignments are distributed to Appliances in other clusters in the failover scope.

> *The new IP assignment of recipient Appliances is not reflected in the Console, in, for example, the Assigned IPs column in the CounterACT Devices pane. You can, however, view which recipient Appliance is handling a specific endpoint. See View Indication of Recipient Appliances Handling Endpoints .*

Recalculation of Failover Assignments

In an environment where one or more Appliances failed, the following events trigger a recalculation of the failover assignments, and may cause these assignments to be redistributed:

- An Appliance reconnects after failure
- An Appliance is removed from a failover cluster
- An Appliance is added to a failover cluster
- A failover scope assignment is changed
- A failover cluster is enabled/disabled
- Failover is disabled/enabled for an Appliance

Handling Endpoints that Exceed Capacity

Each Appliance has a set number of endpoints allotted to the Appliance. This number is set automatically based on default values assigned to the hardware model of your Appliance. This is the Appliance Capacity. Refer to the section on Appliance Endpoint Performance Capacity in the *Forescout Administration Guide*. See Additional Forescout Documentation for information on how to access the guide.

When a failover occurs, and an Appliance receives a Failover Assignment, the Appliance Capacity may be exceeded. In such a case, any endpoints exceeding the capacity are not fully handled by the Appliance. Such *Failover Excess* endpoints are displayed in the Console Detections pane, but not all of their host properties are fully resolved. This means that these endpoints will not match policies that depend on such unresolved properties, and as a result, the relevant actions will not be applied to the endpoints.

You can track the number of excess endpoints by viewing the *Failover Excess* column in the *Failover Status* table. When configuring failover clusters, verify that you have sufficient capacity to handle all endpoints so that no endpoints exceed capacity if a failure occurs. See View Information about Failover Status for more information. You

can also view excess endpoints using a filter in the Detections pane. See Filter Endpoint Failover Information for more information.

Viewing Failover Information in the Console

View endpoint, switch device, WLAN device and Appliance information related to failover in the Console.

- View Endpoint Failover Information
- View Switch Device Failover Information
- View WLAN Device Failover Information
- View Appliance Failover Information

View Endpoint Failover Information

Several interface elements in the Forescout Detections pane support the Failover feature.

- View Indication of Recipient Appliances Handling Endpoints
- Filter Endpoint Failover Information

View Indication of Recipient Appliances Handling Endpoints

During a failover scenario, the **Appliance** column in the Detections pane displays the Appliance name in orange text and with an asterisk (e.g. 10.10.1.22*) for handled endpoints that are currently handled by a *recipient* Appliance due to failover.

An **Appliance** column tooltip is displayed for handled endpoints that are currently failed over to a *recipient* Appliance, containing information about the following Appliances:

- ***Current.*** Current Appliance (recipient) handling the endpoint, after failover.
- ***Original.*** Original Appliance that handled the endpoint, prior to failover.

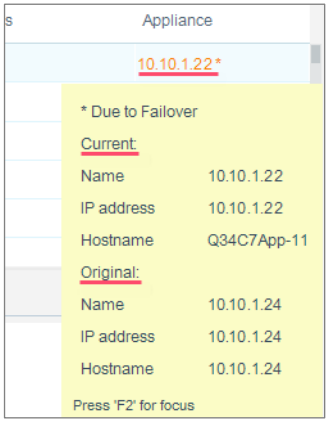

Filter Endpoint Failover Information

If a failover occurs in your deployment, you can filter the view of endpoints at the Detections pane, based on the following:

- ***Failover\Non-Failover***. All endpoints, including both those handled by their originally assigned Appliance and those handled by a failover recipient Appliance.
- ***Failover\Failover Excess***. All endpoints handled by a failover recipient Appliance, including endpoints that exceeded the capacity of the failover recipient Appliance and that are not fully handled by the Appliance.
- ***Failover Excess***. Endpoints that exceeded the capacity of the failover recipient Appliance and that are not fully handled by the Appliance. See Handling Endpoints that Exceed Capacity for more information.

View Switch Device Failover Information

During a failover scenario, the **Managed By** column in the Switch pane displays the Appliance name in orange text and with an asterisk (e.g. 10.10.1.22*) for managed switch devices that are currently failed over to a *recipient* Appliance due to failover. The current status of the Appliance is also displayed.

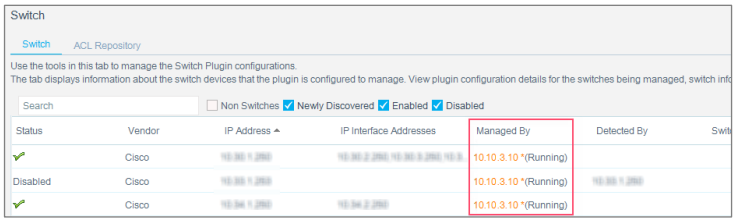

A **Managed By** column tooltip is displayed for managed switch devices that are currently failed over to a *recipient* Appliance, containing the following information:

- ***Current***. Current managing Appliance, after failover.
- ***Original***. Original managing Appliance, prior to failover.
- ***Plugin status***. *The plugin status on the current Appliance is {plugin status}*. Possible statuses are Running, Not Running and Initializing.

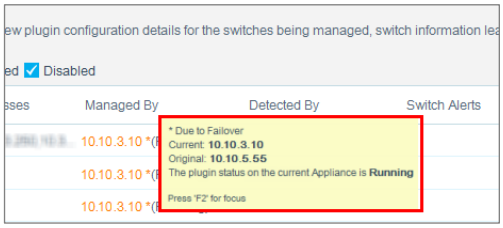

View WLAN Device Failover Information

During a failover scenario, the Wireless pane displays the following information in the **Managed By** column for managed WLAN devices that are currently failed over to a *recipient* Appliance:

- {*current managing Appliance, after failover*} * ({*current managing Appliance status*})

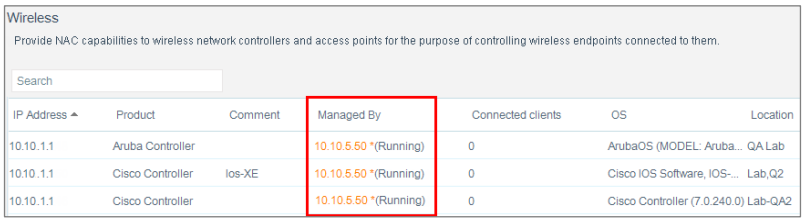

A **Managed By** column tooltip is displayed for managed WLAN devices that are currently failed over to a *recipient* Appliance, containing the following information:

- ***Current***. Current managing Appliance, after failover.
- ***Original***. Original managing Appliance, prior to failover.
- ***Plugin status***. *The plugin status on the current Appliance is {plugin status}*. Possible statuses are Running, Not Running and Initializing.

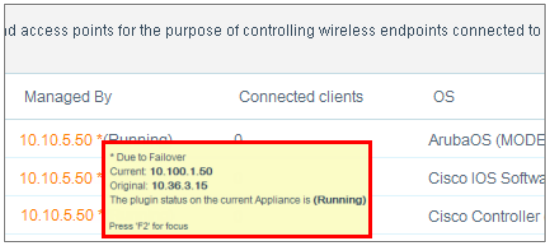

View Appliance Failover Information

Appliance failover details can be viewed in the CounterACT Devices pane.

- Appliance and Failover Cluster Failure Warning Icons
- Appliance Health Information
- Failover Scope Column

Appliance and Failover Cluster Failure Warning Icons

The following icons appear in the CounterACT Devices pane in certain circumstances:

- Appears in the Devices folder list when a failover cluster contains one or more failed Appliances.

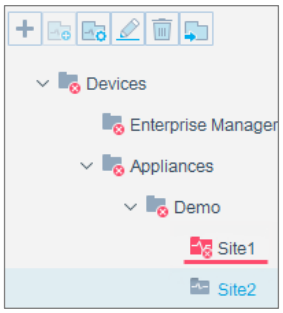

- Appears in the Status column of the CounterACT Devices table when the given Appliance has failed.

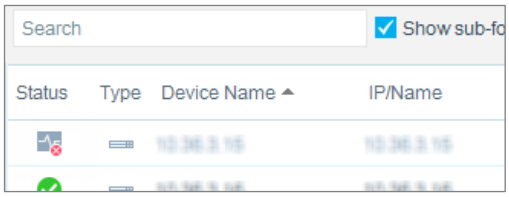

Appliance Health Information

The Status section of the CounterACT Devices pane, containing Appliance health details, includes an entry related to failover. The entry, *Endpoints (Failover)*, contains the following information:

- Total number of endpoints handled by the Appliance:
  - Endpoints that are part of the original IP assignment of the Appliance.
  - Endpoints that are now handled by the recipient Appliance due to a failover assignment from another Appliance.
  - Number of *Failover Excess* endpoints that have exceeded CounterACT device capacity and are not fully handled by any CounterACT device.
- CounterACT device capacity. See Handling Endpoints that Exceed Capacity.

This information is only displayed if failover is enabled for the Appliance and an Appliance failure occurs. See Configuring Failover for more information.

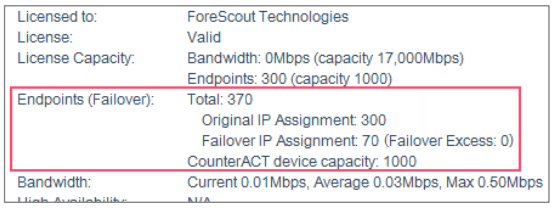

Failover Scope Column

View the failover clusters that the selected Appliance will fail over to, including the failover cluster that the selected Appliance belongs to. This information appears in the CounterACT Devices pane, *Failover Scope* column. This column appears when a failover cluster is selected. It also appears when a regular Appliance folder containing a failover cluster as a sub-folder is selected, if the *Show sub-folders* option is selected.

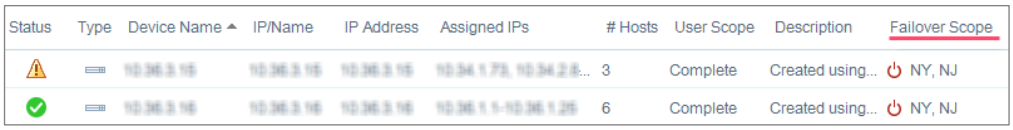

If the selected Appliance is in a failover cluster, an icon appears if failover is disabled. If the selected Appliance is not in a failover cluster, the status *No Failover* is displayed.

Tracking Failover Activity

You can display selected information about failover activity in the Forescout Audit Trails log and Event Viewer. Email alerts are also sent for selected activities. A MIB Table for CounterACT Appliances provides objects that report information about failover, and SNMP trap notifications are issued when the value of a configuration, status, or performance attribute of the MIB Table changes on an Appliance.

- Audit Trail
- Event Viewer
- Email Alerts
- MIB Table Objects
- SNMP Trap Notifications

Refer to the Forescout Administration Guide *for more information about any of these tracking methods. See Additional Forescout Documentation for information on how to access the guide.*

Audit Trail

The Audit Trail reports indicate the following information about users who have modified failover configuration settings:

- When a user adds/removes an Appliance to/from a failover cluster.
- When a user assigns/removes a failover cluster to/from a failover scope.
- When a user enables or disables a failover cluster.
- When a user performs manual failover/failback on an Appliance.

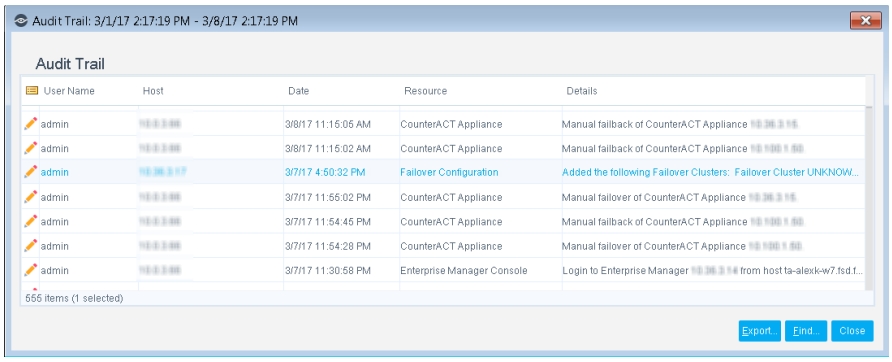

To access the Audit Trail:

1. Select **Audit Trails** from the **Log** menu.
2. Enter a time period and select **OK**. The Audit Trails log opens.

Event Viewer

The Event Viewer indicates:

- When an Appliance fails.
- When an event occurred which caused a recalculation of failover assignments within a failover cluster, a reassignment of the failover assignment to other Appliances, and potentially an excess of endpoints that will not be fully handled by any Appliance.
- When an Appliance fails back and resumes managing assignments that were transferred during a failover.

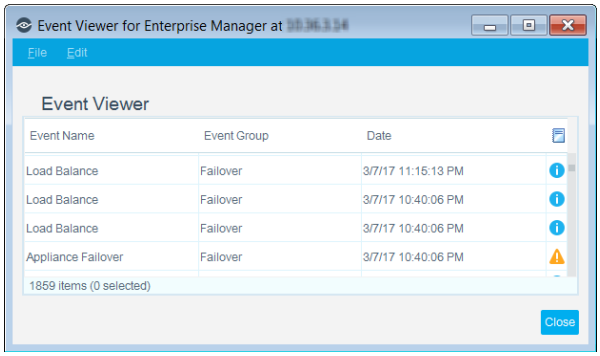

To access the Event Viewer:

1. Select **Event Viewer** from the **Log** menu. The Event Viewer opens.

Email Alerts

An email alert is sent to the Forescout operator for each event reported in the Event Viewer. Refer to the section on managing email notification addresses in the *Forescout Administration Guide* for information on configuring the address to which email alerts are sent. See Additional Forescout Documentation for information on how to access the guide.

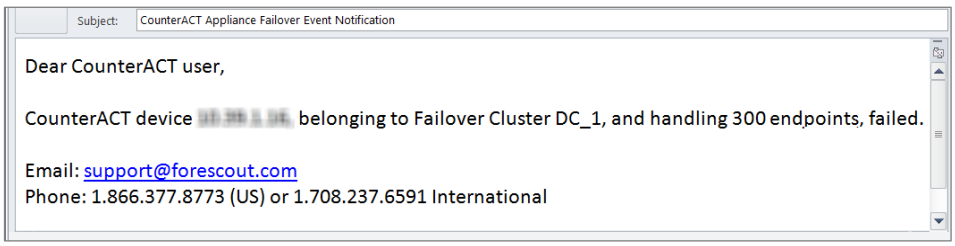

MIB Table Objects

The following are CounterACT Appliance MIB objects associated with failover:

ctDeviceFailoverStatus

OID: .1.3.6.1.4.1.11789.4.3.1.16

Indicates the status of the CounterACT device and whether or not it has experienced failover.

Possible states are:

Ok (1) Indicates that the CounterACT device is running.

Failed (2) Indicates that the CounterACT device has failed to connect to the Enterprise Manager for the configured timeout period.

The following related trap notifications are provided:

- ctDeviceFailoverStatusChangedTrap
- ctConfigurationChangedTrap

ctDeviceFailoverConfig

OID: .1.3.6.1.4.1.11789.4.3.1.17

Indicates the status of the failover feature on the CounterACT Appliance. Possible states are:

Enabled (1) Indicates that failover is enabled on the CounterACTdevice.

Disabled (2) Indicates that failover is disabled on the CounterACT device.

ctFailoverClusterName

OID: .1.3.6.1.4.1.11789.4.3.1.18

Indicates the name of the failover cluster that the CounterACT device belongs to.

The following related trap notification is provided:

- ctConfigurationChangedTrap

ctNumberOfFailoverEndpoints

OID: .1.3.6.1.4.1.11789.4.3.1.19

Indicates the number of endpoints handled by this CounterACT device due to a failover assignment.

ctNumberOfFailoverExcessEndpoints

OID: .1.3.6.1.4.1.11789.4.3.1.20

Indicates the number of endpoints that exceed the CounterACT device capacity as a result of a failover.

The following related trap notifications are provided:

- ctDeviceFailoverCapacityExceeded
- ctDeviceFailoverCapacityNormal

SNMP Trap Notifications

CounterACT Appliance trap notifications associated with the failover feature can contain the following MIB objects:

ctDeviceFailoverStatusChangedTrap

OID: .1.3.6.1.4.1.11789.0.34

This trap notification is sent when there is a change in the ctDeviceFailoverStatus MIB attribute.

ctDeviceFailoverCapacityExceeded

OID: .1.3.6.1.4.1.11789.0.35

This trap notification is sent when the CounterACT device has exceeded its configured capacity as a result of a failover.

This alarm trap is cleared by the **ctDeviceFailoverCapacityNormal** trap notification.

ctDeviceFailoverCapacityNormal

OID: .1.3.6.1.4.1.11789.0.36

This trap notification is sent when the CounterACT device no longer exceeds its configured capacity.

This trap notification is only sent after a **ctDeviceFailoverCapacityExceeded** alarm trap was sent.

System Backup

Performing a system backup will store configuration changes made in the Console related to the failover feature.

Failover on High Availability Systems

Failover within a failover cluster or scope occurs on high availability systems if the Active node fails and the Standby node does not take over within the failover detection time. See Configuring Failover for more information about the failover detection time period.

Features Not Currently Supported

Continuity of visibility, information and enforcement is not currently available for:

- Endpoints connecting via VPN
- Forescout extended modules

See Continuity of Endpoint, Switch and Wireless Device Visibility, Information and Enforcement for more information.

High Availability Pairing

A high availability system provides you with standby support in the event of system malfunction or failure. It is implemented in pairs of two CounterACT devices; either two Appliances or two Enterprise Managers. Redundancy is achieved by assigning an Active node to manage activities required for effective NAC, and a Standby node to take over in case the Active node fails. The two nodes are synchronized by a redundant pair of interconnecting cables.

- ***Primary Node***: the node you use to set up high availability. It is initially the Active node.
- ***Secondary Node***: the node that you set up to take over if the primary node fails. It is initially the Standby node.
- ***Active Node***: The active node manages your environment. Only one node is active at a time.
- ***Standby Node***: The operating system on the passive node is active and ready to be used as the failover system.

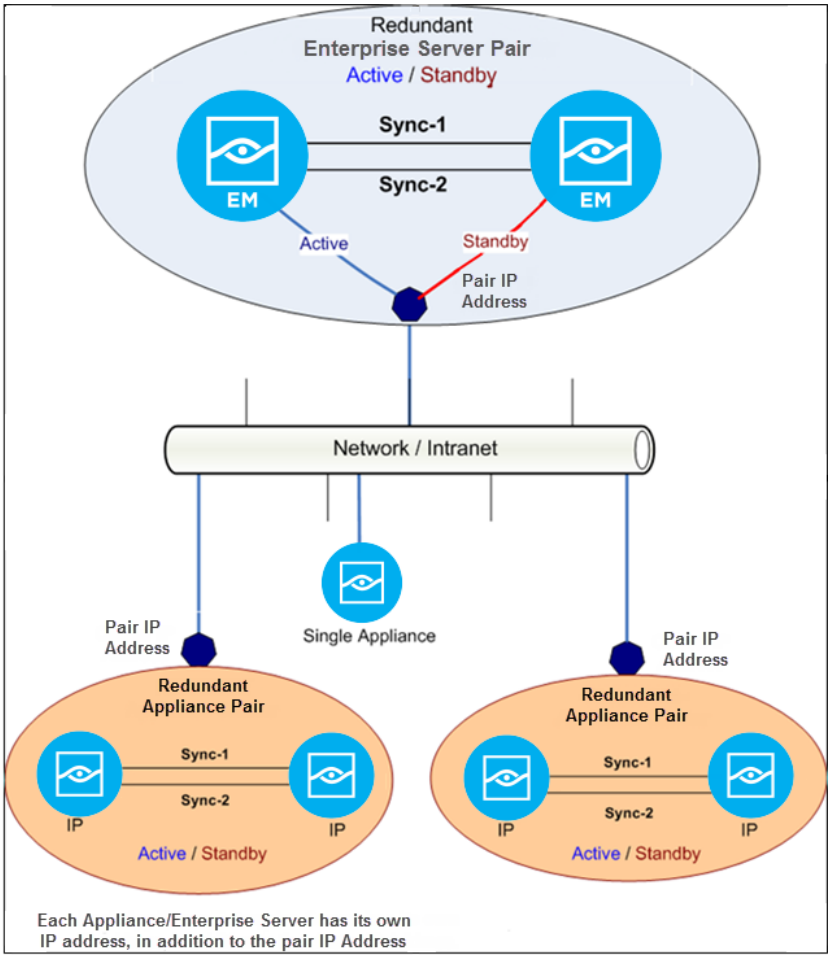

Redundant Enterprise Server Pair

License Setup Requirements

If your deployment is using Flexx Licensing mode, see Flexx Licensing Requirements for license requirements.

To identify your licensing mode:

- From the Console, select **Help > About Forescout**.

Per-Appliance Licensing Mode

The demo license for your high availability system is valid for 30 days. You must install a permanent license before this period expires. You will be contacted via email regarding the expiration date of the license.

Forescout virtual systems do not come with a demo license.

You should use the IP address of the high availability pair when requesting a high availability license. If a license is only issued to the Active node in a high availability pair, the system may not operate after failover to the Standby node.

Both nodes must be up when requesting a license.

Network Access Requirements

Deploying the Forescout platform requires TCP/IP communication. This section details Forescout platform connectivity requirements. Check your security policy (Router ACLs etc.), and modify it, if required, to allow for this communication.

Each Appliance requires a single management connection to the network. This connection requires an IP address on the local LAN and port 13000/TCP access from machines that run the Forescout Console. The connectivity listed in the following table is required.

This section details minimum connectivity requirements needed for Appliances in high availability systems. For a complete list of requirements for CounterACT devices, refer to the Forescout Installation Guide. See Additional Forescout Documentation for more information on how to access the guide.

Port	Service	To or From Forescout	Function
22/TCP	SSH	From	Allows remote inspection of OS X and Linux endpoints. Allows the Forescout platform to communicate with network switches and routers.
		To	Allows access to the Forescout platform command line interface.

Port	Service	To or From Forescout	Function
2222/TCP	SSH	To	(High Availability) Allows access to the physical CounterACT devices that are part of the high availability pair. Use 22/TCP to access the shared (virtual) IP address of the pair.
53/UDP	DNS	From	Allows the Forescout platform to resolve internal IP addresses.
80/TCP	HTTP	To	Allows HTTP redirection.
123/UDP	NTP	From	Allows the Forescout platform access to a local time server or ntp.Forescout.net By default the Forescout platform accesses ntp.Forescout.net
161/UDP	SNMP	From	Allows the Forescout platform to communicate with network switches and routers. For information about configuring SNMP, refer to the *Forescout Administration Guide*. See Additional Forescout Documentation for information on how to access the guide.
162/UDP	SNMP	To	Allows the Forescout platform to receive SNMP traps from network switches and routers. For information about configuring SNMP, refer to the *Forescout Administration Guide*. See Additional Forescout Documentation for information on how to access the guide.
443/TCP	HTTPS	To	Allows HTTP redirection using TLS.
13000/TCP	Forescout	From/To	For systems with only one Appliance – from the Console to the Appliance. For systems with more than one CounterACT Device – from the Console to the CounterACT Device and from one CounterACT Device to another. CounterACT Device communication includes communication with the Enterprise Manager and the Recovery Enterprise Manager, using TLS.

Communication with the High Availability System

You will usually want to communicate with the high availability pair (and not the individual CounterACT devices that make up the pair). Use 22/TCP to access the Forescout command line interface of the shared (virtual) IP address of the pair.

If you do need to access a physical CounterACT device that is part of the high availability pair, use 2222/TCP. Note that the CLI prompt in this case begins with **`Miniroot`**.

Many commands (in particular, **`fstool service status`**) work on the shared IP address only. Running them at the **`Miniroot`** prompt does not work.

Switch Connectivity

Below are examples of high availability pair-switch connections. In the relevant examples, the switch must support Dual Span sessions.

Dual cross cables must be connected for redundancy.

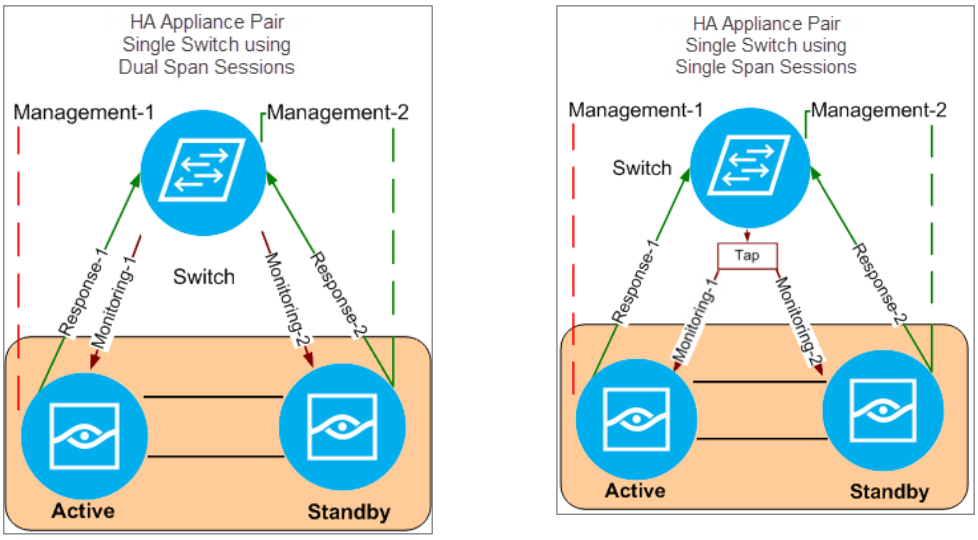

High Availability Pair – Switch Connection Example (1/3)

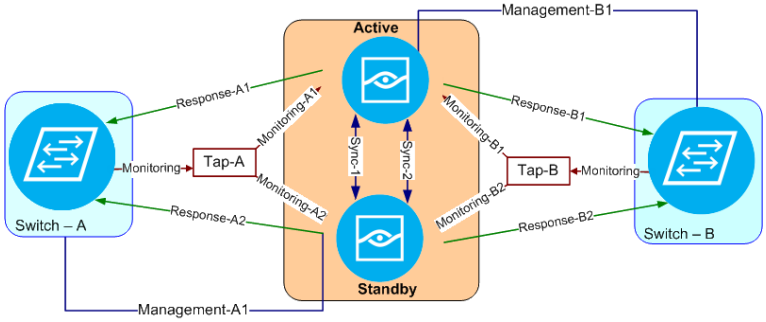

High Availability Pair – Switch Connection Example (2/3)

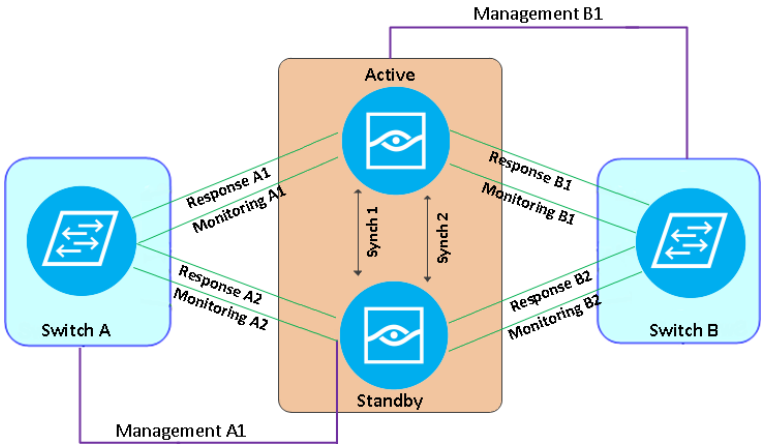

High Availability Pair – Switch Connection Example (3/3)

Connecting to the Network

This section shows sample ports and connectors for a single CounterACT device.

CT-R Appliances are not supported in high availability pairs.

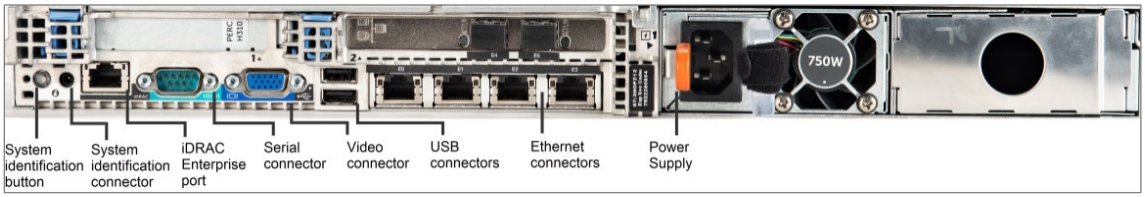

Sample Appliance Rear Panel –CT-xxxx Appliance

CT-1000 – Sample Connections				
Interface	**Cable**		**Interface**	**Cable**
eth0	Management-1		eth4*	Monitoring-2*
eth2	Monitoring-1		eth5*	Response-2*
eth3	Sync-1		eth7	Sync-2
eth1	Response-1			

*Only for redundant switch configuration.

It is recommended to use two sync cables whenever possible. In addition, you can attach the sync>management cables to sockets on different NICs to improve handling of NIC failure with all attached sockets.

Failover

The Active and Standby nodes are monitored every second for operational updates. By default, failover from the Active node to the Standby node occurs 60 seconds after the Standby node detects that the Active node is down.

The Standby node becomes the Active node typically within 10 minutes or more after the Active node fails.

Failover Triggers

Full high availability mode requires that both nodes are up and that the Standby node is synchronized with the Active node and is fully up-to-date.

When full high availability mode is in effect, any of the following cause the Standby node to become the Active node:

- System failure: Active node outage
- System failure: Hardware RAID array breakdown; that is, some disks are not functioning
- System maintenance: Active node powered off or cold boot occurred
- Management interface failure: A management interface hardware failure on the Active node. Failover occurs in this case only if you defined pingable hosts during the high availability setup.
- Triggering failover with a manual command. On the Active node, use the command:
  `fstool ha stop -f`

Node Status

The status of the Active and Standby nodes is affected by restarts as follows:

- ***Restart Active node***: In case the Active node fails, the Standby node becomes the Active node (swapping roles). After restart, the switchover remains in effect; that is, the Active node that originally failed remains the Standby node, and the newly appointed Active node continues with that role.
- ***Restart Standby node***: After restarting the Standby node, the Active/Standby roles do not change.
- ***Both nodes are restarted***: Depending on which node restarts first, the nodes can remain as originally designated or assume reverse roles; the first node to restart becomes the Active node.

Installing High Availability Software

During the installation procedure, the nodes are referred to as the *Primary node* and the *Secondary node*. The nodes are referred to as the *Active node* and the *Standby node* after installation and during operation, according to their current status.

The installation and configuration procedure is performed in three main stages:

1. Set up high availability for the Primary node.
2. Configure the Primary node.
3. Set up high availability for the Secondary node. The Secondary node automatically obtains the high availability configuration file from the other high availability pair during the setup process.

- *Reboots may occur during these stages. This does not indicate any type of failure or problem.*
- *Following the second reboot in the high availability setup, allow time for data synchronization.*

Identifying Ethernet Ports

If you do not know the Ethernet port layout of either the Primary or Secondary node rear panel, perform the following procedure before configuring the Forescout platform.

To identify Ethernet ports:

1. Power on the CounterACT device.

```
CounterACT <version>-<build> options:

1) Configure CounterACT
2) Restore saved CounterACT configuration
3) Identify and renumber network interfaces
4) Configure keyboard layout
5) Turn machine off
6) Reboot the machine

Choice (1-6) :
```

2. To identify the ports on the rear panel of the Appliance, type **3** and press **Enter**.

   Text is displayed indicating which interface has been detected. The associated port LED blinks on the rear panel.

3. Label the port on the panel so that it is easily identifiable and press **Enter**.

   More text is displayed indicating the next detected interface. The associated port LED now blinks.

4. Label this port as well and press **Enter**. This process continues until all active interfaces are detected and you have labeled the associated port for each active interface.

5. Once all interfaces have been detected, press **Enter**.

Primary Node CounterACT Device Setup

To set up a CounterACT device as a High Availability Primary node:

1. Power on the CounterACT device.

```
CounterACT <version>-<build> options:

1) Configure CounterACT
2) Restore saved CounterACT configuration
3) Identify and renumber network interfaces
4) Configure keyboard layout
5) Turn machine off
6) Reboot the machine

Choice (1-6) :
```

2. If necessary, identify the device interfaces as described in the preceding section.

   When you have finished, the menu reopens.

3. Type **1** and press **Enter**.

```
Select High Availability Mode:

1) Standard Installation
2) High Availability – Primary Node
3) Add node to existing Active Node (Primary or Secondary)

Choice (1-3) [1] :
```

4. Type **2** and press **Enter**.

```
>>>>>> CounterACT Initial Setup (High Availability – Primary
node) <<<<<<
You are about to setup a CounterACT Primary High Availability
node.
Note that High Availability mode requires 2 CounterACT machines.
During the initial setup process you will be prompted for basic
parameters used to connect the High Availability pair to the
network.
When this phase is complete, you will be instructed to continue
the setup from the CounterACT Console.
Note: If you would like to attach this node to an existing High
Availability node, please return to the previous menu and select
option 3.
Continue? (yes/no) :
```

5. Type **yes** and press **Enter** to continue.

The following prompt appears when running a clean installation of 8.1.

```
Certification Compliance Mode? (yes/no) [no] :
```

6. Unless your organization needs to comply with Common Criteria and DoDIN APL certification, type **No** and press **Enter**. Refer to the *Forescout Installation Guide* for information on Certification Compliance mode.

```
>>>>>> Select CounterACT Installation Type <<<<<<

1) CounterACT Appliance
2) CounterACT Enterprise Manager

Choice (1-2) :
```

7. Type **1** or **2** and press **Enter**.

```
>>>>>> Select Licensing Mode <<<<<<

1) Per Appliance licensing mode
2) Flexx licensing mode

Choice (1-2) :
```

8. Select the licensing mode that your deployment uses. The licensing mode is determined during purchase. ***Do not type a value until you have verified what licensing mode your deployment uses***. Contact your Forescout representative to verify your licensing mode or if you entered the wrong mode.

This option does not appear on Forescout 51xx Appliances.

Type **1** or **2** and press **Enter**.

```
>>>>>> Enter Machine Description <<<<<<

Enter a short description of this machine (e.g. New York
office).

Description [<Enterprise Manager|Appliance>] :
```

9. Type a description and press **Enter**.

```
>>>>>> Set Administrator Password <<<<<<

This password is used to log in as 'cliadmin' to the machine
Operating System and as 'admin' to the CounterACT Console.
The password should be between 6 and 15 characters long and
should contain at least one non-alphabetic character.

Administrator password :
```

10. Type the string that is to be your password (the string is not echoed to the screen) and press **Enter**. You are asked to confirm the password:

```
Administrator password (confirm) :
```

11. Retype the password and press **Enter**.

```
>>>>>> Set Host Name <<<<<<

It is recommended to choose a unique host name.

Host name :
```

12. Type a host name and press **Enter**.

Suggestion: When upgrading, use the previous host name.

- The host name can be used when logging into the Console. In addition, it is displayed on the Console to help you identify the CounterACT Appliance that you are viewing. The hostname should not exceed 13 characters.
- When you enter a pair host name, for example, High_Availability_pair, the system will automatically assign the name High_Availability_pair-1 to the Primary node and the name High_Availability_pair-2 to the Secondary node. You should add these in the DNS server.

The **HA Pair management interface** prompt is displayed (subsequent prompts are displayed after you enter a value for the preceding prompt):

```
>>>>>> Configure Network Settings <<<<<<

HA pair management interface (one of: <interface_list>) :
HA pair management IP address :
Network mask [255.255.255.0] :
Default gateway :
Domain name :
DNS server addresses :
```

13. Type a value at the **HA Pair management interface** prompt and press **Enter**. The number of pair management interfaces listed depends on the Appliance model.

14. Type a value at each subsequent prompt and press **Enter**.

- The gateway address should be within the network of the pair address.

Suggestion: When upgrading, use the previous IP address of the pair.

- If there is more than one **DNS server address**, separate each address with a space—Most internal DNS servers resolve external addresses as well but you may need to include an external-resolving DNS server. As nearly all DNS queries performed by the Appliance will be for internal addresses, the external DNS server should be listed last.

  The **Enter the Fully Qualified Domain Name** prompt is displayed:

```
>>>>>> Set Fully Qualified Domain Name <<<<<<

Enter the Fully Qualified Domain Name (or ‘none’) [<FQDN>] :
```

15. Type the fully qualified domain name of the pair management interface.

The **This node private IP Address** prompt is displayed:

```
>>>>>> High Availability Parameters <<<<<<
This node private IP Address [<IP address>]:
Other node private IP Address [<IP address>]:
Management IPv6 address or ’none’ :
Assign an Out-Of-Band IP address? (yes/no)
Out-Of-Band IP Address (one of: eth1, eth2, eth3, eth4, eth5) [
]:
HA Pair Out-Of-Band IP Address [<IP address>]:
This node Out-Of-Band IP Address [<IP address>]:
Other node Out-Of-Band IP Address [<IP address>]:
Netmask size of the HA Pair IP address [ ]:
Configure management network keepalive ping test? (yes/no)
IP address for management network keepalive ping tests [<IP
address>]:
```

```
Select the primary Ethernet sync interface (one of: eth1, eth2,
eth3, eth4, eth5) [ ]:
Configure secondary sync interface (one of: eth1, eth2, eth3,
eth4, eth5) [ ]:
Secondary Ethernet sync interface (one of: eth1, eth2, eth3,
eth4, eth5) [ ]:
Sync subnet [172.17.2.0] :
Would you like to configure bonding for the sync interfaces?
(yes/no)
```

Unless the nodes are synchronized via cross cable, you should configure different sync subnets for different pairs.

It is recommended to configure bonding for the HA sync interfaces.

Suggestion: When upgrading, use the previous sync interface, otherwise you may lose connections in the Control screen.

16. Type a value at each prompt, and press **Enter**. After answering **yes** at the last prompt, the setup summary is displayed:

```
>>>>>> Setup Summary <<<<<<
Role:                          <CounterACT_device_type>
HA Pair host name:             <user_entered_value>
Description:                   <user_entered_value>
HA Pair management Interface:  IP: < user_entered_value >,

Interface: eth<n>,

Netmask: <user_value>
This node physical IP address:  <user_entered_value>
Other node physical IP address: <user_entered_value>
Primary sync interface:        eth<n>
Secondary sync interface:      eth<n>
Default gateway:               <user_entered_value>
DNS server:                    <user_entered_value>
Domain name:                   <user_entered_value>
(T)est,(R)econfigure,(D)one :
```

17. To test the configuration, type **T** and press **Enter**.

The test verifies the following:

- Storage I/O performance (Virtual systems only)
- Connected interfaces
- Connectivity of the default gateway
- DNS resolution

Results indicate if any test failed so that you can reconfigure if necessary.

If there are no failures, the following is displayed:

```
Checking eth0...OK. (100Mb/s Full duplex)
Checking default gateway...OK.
Checking DNS resolution...OK.

Press ENTER to review configuration summary
```

18. Press **Enter**. The setup summary is displayed again.

19. To complete the installation, type **D** and press **Enter**.

```
Finalizing CounterACT setup, this will take a few minutes
```

After setup is complete, the following is displayed:

```
CounterACT <version-build> is booting...
```

And then:

```
Setting up HA primary node...
```

The system reboots. Following the reboot, this notification is displayed:

20. Type **C** and press **Enter** to continue.

21. Continue by setting up the Secondary node.

Secondary Node CounterACT Device Setup

Before you begin setting up the Secondary node CounterACT device, verify that the Primary node CounterACT device is powered on, set up, and successfully configured.

When setting up the Secondary node CounterACT device, use the same sync interfaces and netmask settings used in the Primary node CounterACT device.

To set up a CounterACT device as a High Availability Secondary node:

1. Power on the CounterACT device.

```
CounterACT <version>-<build> options:

1) Configure CounterACT
2) Restore saved CounterACT configuration
3) Identify and renumber network interfaces
4) Configure keyboard layout
5) Turn machine off
6) Reboot the machine

Choice (1-6) :
```

2. If necessary, identify the device interfaces as described in Identifying Ethernet Ports.

When you have finished, the preceding menu reopens.

3. Type **1** and press **Enter**.

```
Select High Availability Mode:

1) Standard Installation
2) High Availability – Primary Node
3) Add node to existing Active Node (Primary or Secondary)

Choice (1-3) [1] :
```

4. Type **3** and press **Enter**.

```
>>>> CounterACT Initial High Availability Setup <<<<

You are about to set up a CounterACT High Availability node.
Note that secondary High Availability configuration for this
node will be fetched from the Active Node (Primary or
Secondary).
Before continuing verify that the Active node is reachable.


Continue? (yes/no) :
```

5. Type **yes** and press **Enter** to continue.

   The **Select the dedicated primary sync interface** prompt is displayed (subsequent prompts are displayed after you enter a value for the preceding prompt):

```
>>>>>> Secondary High Availability Parameters <<<<<<

Select the primary Ethernet sync interface (one of:
<interface_list>) :
Sync subnet [172.17.2.0] :
Enter the administrator password of the active node:
```

   – The root password of the Primary node is the password that you entered in step 10 of Primary Node CounterACT Device Setup.

6. Type a value at the **Select the primary Ethernet sync interface** prompt and press **Enter**.

7. Type a value at each subsequent prompt and press **Enter**.

   After pressing **Enter** at the last prompt, the setup summary is displayed:

```
>>>>>> Setup Summary <<<<<<

Primary sync interface:        eth<n>
Sync interface subnet:         <user_entered_value>

(T)est,(R)econfigure,(D)one :
```

8. To test the configuration, type **T** and press **Enter**.

```
* Attempting to retrieve the parameters from <address> *
```

The address is 172.17.2.171 (based on the default subnet, 172.17.2.0).

   The test verifies the following:

   – Storage I/O performance (Virtual systems only)
   – Connected interfaces
   – Connectivity of the default gateway
   – DNS resolution

   Results indicate if any test failed so that you can reconfigure if necessary.

   If there are no failures, the following is displayed:

```
Secondary config tested successfully.
Press ENTER to review setup summary
```

9. Press **Enter**.

   The setup summary is displayed again.

10. To complete the installation, type **D** and press **Enter**.

```
Preparing the High Availability file system
```

After setup is complete, the following is displayed:

```
>>>>>> Rebooting the system <<<<<<

The first phase of the High Availability setup is complete
After the reboot, the final High Availability setup phase will
be performed automatically.
Rebooting in <n> seconds
```

Shutting Down and Rebooting High Availability Devices

You should shut down and reboot high availability devices using fstool commands, and not the usual Linux commands.

To shut down a High Availability CounterACT device:

1. Run **`fstool reboot -s`**

To reboot a High Availability CounterACT device:

1. Run **`fstool reboot`**

Moving a High Availability Pair

To move a High Availability pair from one network to another:

1. Shut down the Standby node by running the command **`fstool reboot -s`**
2. Shut down the Active node by running the command **`fstool reboot -s`**
3. Relocate both CounterACT devices and connect them as described in Connecting to the Network.
4. Restart the Active node.
5. Run **`fstool ha_setup`** on the Active node, making sure to use the new network settings.
6. Restart the Standby node.
7. Run **`fstool ha_setup`** on the Standby node. It is recommended to do this from the Linux Console since the management IP address will probably be different.

8. If the new DNS settings are different, run **`fstool dns setup <new_DNS>`** on both machines and reconfigure them.
9. If the new NTP settings are different, run **`fstool ntp setup <new_NTP> (<Optional_second_NTP_IP>)`** on both machines and reconfigure them.
10. Verify that the pair is up and running.

High Availability Backup and Restore

The backup and restore procedures for high availability differ from the standard backup and restore procedures. If one of the nodes crashes, the other node takes over. If necessary, you can then install a new CounterACT device as a Standby node.

To protect your system from a situation where both nodes crash or for some other reason you must back up the system while the high availability pair is operational.

You will require an external storage device to restore the configuration file.

To back up a High Availability pair:

1. Connect the two CounterACT devices with redundant cross cables.
2. Perform a backup of the system settings and copy the configuration backup files to an external storage medium. Refer to the section on backing up system settings in the *ForescoutCounterACT Administration Guide*. See Additional Forescout Documentation for information on how to access the guide.

To restore a High Availability pair:

1. Format the disk on each CounterACT device.
2. Reinstall the Forescout software on the Primary node. Use the latest image file (*.ISO) that is available on the Product Downloads page of the Customer Portal.
3. Power on the first CounterACT device.

```
CounterACT <version>-<build> options:

1) Configure CounterACT
2) Restore saved CounterACT configuration
3) Identify and renumber network interfaces
4) Configure keyboard layout
5) Turn machine off
6) Reboot the machine

Choice (1-6) :
```

4. Type **2** and press **Enter**.

```
Restore options:

  1) Restore from USB storage device
  2) Restore from CD-ROM
  3) Get shell prompt
  4) Reset to factory setup
  5) Cancel

Choice (1-5) :
```

5. Type the number of the relevant restore option and press **Enter**.

```
The restore process will now search for backup files in the
selected media. Note that backup file names must have a ".fsb"
extension. Insert the media where the backup file reside and
press ENTER to continue
```

6. Insert the media where the backup file resides and press **Enter**.

   All FSB files found on the media are displayed.

```
Searching for backup files in <selected_storage_type>...

Choose backup file:
1) <backup_file1_name>.fsb
2) <backup_file2_name>.fsb
3) Cancel

Choice (1-3) :
```

7. Type the number of the relevant backup file and press **Enter**.

```
Verifying <full_path_and_file_name>.fsb...
-------------------------
Backup Volume Information
-------------------------

Product     : CounterACT
Host-name   : <host_name>
Address     : <IP_address>
Backup date : <date_and_time_stamp>

Verifying Backup volume,please wait.

Restore? (yes/no) :
```

8. Type **yes** and press **Enter**.

```
Setup the restored machine in High Availability mode? (yes/no)
[no]
```

9. Type **yes** and press **Enter**. The original high availability values were saved along with the backup and are presented as default values, which should be accepted. See Primary Node CounterACT Device Setup for more information about these values.

10. Install the Forescout software on the Secondary node.

11. Set up high availability on the Secondary node, as described in Secondary Node CounterACT Device Setup.

Do not restore a backup to the Secondary node.

12. Connect the redundant (dual) physical cables to the management, monitor and response ports between the CounterACT device and the switches after the configuration determines the layout of the Ethernet interfaces on the rear panel.

When you backup and restore system settings using two different CounterACT devices, the interface numbering may change. To correlate the new interface numbering with the correct interfaces you must run ***`fstool ethtest`*** *and reassign the interfaces accordingly.*

If you are restoring after upgrading the Forescout version on the devices:

1. Log in to the Primary node and run the following command:

   ```
   fstool upgrade
   ```

Restoring as a Standard Device

Note that you can select to restore system settings to a standard device, even if the backup was taken from a high availability device.

Upgrading High Availability Systems to the Latest Version

You can upgrade the CounterACT device version from the Console. For details about this procedure, refer to the section on *Upgrading Appliance Software* in the *Forescout Administration Guide*. See Additional Forescout Documentation for information on how to access the guide.

It is recommended to back up your system before performing the upgrade. See High Availability Backup and Restore for details on how to proceed.

Converting a High Availability node to a standalone CounterACT Device

Use this procedure to convert a node with high availability to standard mode.

Until you perform the final step, you can re-activate the Active node and continue using high availability mode.

1. Back up the configuration of the Active node.

   If necessary, perform an rSite backup as well.

2. Make a clean, standard installation of the Forescout software on the Standby node. (during this time the Forescout services will continue to be provided by the Active node)
3. Disconnect the high availability cables.

4. Shut down the Active node (it can be reactivated later if necessary) by running the command **`fstool reboot -s`**
5. Restore the configuration data (and rSite data, if backed up previously) to the Standby node.
6. Verify that the Standby node is configured in standard mode (not high availability) and is running.

   The Standby node is now a standalone CounterACT device.
7. Run a clean, standard installation of the Forescout software on the Active node and reactivate it if necessary.

Converting CounterACT Devices to High Availability

This section details how to convert CounterACT devices to a high availability system. The conversion makes the CounterACT devices suitable for use in a high availability pair.

This section includes:

- Requirements for High Availability Conversion
- Conversion Procedures

Requirements for High Availability Conversion

- Both CounterACT devices must have the same number of Ethernet interfaces.
- The capacity of the Secondary node disk must be at least that of the Primary node disk.
- Both CounterACT devices must be running the same version of Forescout.

Conversion Procedures

To convert the first CounterACT device:

1. Perform one of the following:
   - Back up Forescout system settings and restore in high availability mode. The restore procedure requires that you reinstall Forescout. See High Availability Backup and Restore. Refer to the *Forescout Administration Guide.* See Additional Forescout Documentation for information on how to access the guide.
   - Reinstall Forescout.
     You will lose configuration settings during the reinstallation process, so first record all relevant settings, including interface setup, interface assignment, and IP address details.

   Proceed with high availability configuration of the CounterACT device as detailed in: Identifying Ethernet Ports and Primary Node CounterACT Device Setup.

To convert the second CounterACT device:

1. ***For an Appliance***: Connect the redundant (dual) physical cables to the management, monitor and response ports between the Appliance and the switches.

   For an Enterprise Manager: Connect the redundant (dual) physical cables to the management port between the Enterprise Manager and the switches.

2. Configure the high availability settings, as detailed in Secondary Node CounterACT Device Setup.

Tracking High Availability Activity

The Forescout platform provides the following capabilities to track and notify Forescout operators about high availability-related events occurring in their Forescout deployment:

- Event Viewer and Email Tracking
- High Availability Indicators on the Console

Event Viewer and Email Tracking

For every high availability failover or status change that occurs, the Forescout platform tracks each event by recording an entry in the Event Viewer. These entries can be viewed in the Console.

To access the Event Viewer:

1. From the Console **Log** menu, select *Event Viewer*. The **Time Period** dialog opens.
2. Enter a time period and select **OK**. The **Event Viewer** opens.

For every high availability failback that occurs, the Forescout platform sends an email alert to the Forescout operator. Refer to the section on managing email notification addresses in the *Forescout Administration Guide* for information on configuring the address to which email alerts are sent. See Additional Forescout Documentation for information on how to access the guide.

High Availability Indicators on the Console

The Forescout Console indicates the status of your high availability pair. The relevant icon from the following table appears on the status bar of the Console.

	Indicates that high availability functionality is up on a Forescout high availability pair.
	Indicates that high availability functionality is down on a Forescout high availability pair.

A tooltip provides additional information.

In addition, the CounterACT Devices pane in the Options window provides information on the high availability status of each CounterACT device in the enterprise. (To display this information, right-click a column header, select **Add/Remove Columns** and then add the **HA** column.)

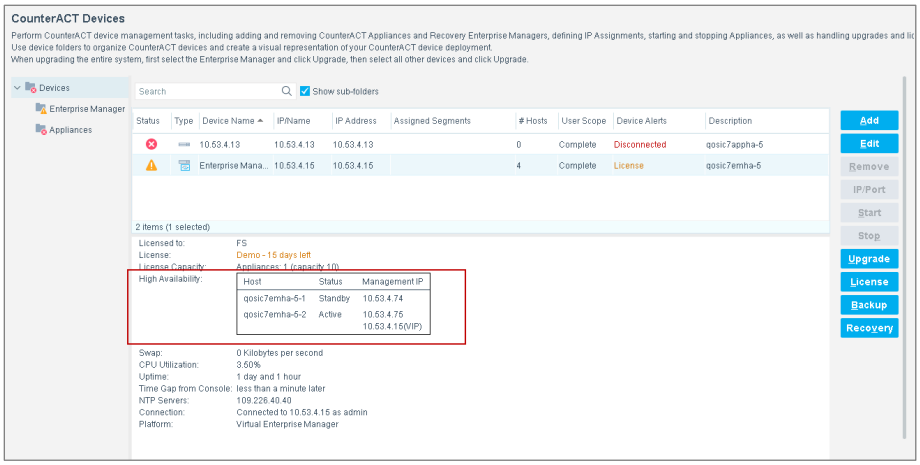

The following high availability status indications are available:

- **Upgrade**: Forescout is in the process of upgrading.
- **Setup**: Forescout is in the process of configuring.
- **Active**: The Active node
- **Standby**: The Standby node

Disaster Recovery for Enterprise Manager

A Forescout recovery tool provides a comprehensive recovery system for an Enterprise Manager that is no longer functioning, for example, if it failed as a result of an earthquake or fire. This feature provides complete and continued management of Appliances from a Recovery Enterprise Manager after the crisis.

Not all users have access to this feature. Refer to the section about access to Console tools in the *Forescout Administration Guide* for more information about limiting and allowing user access to this feature. See Additional Forescout Documentation for information on how to access the guide.

How It Works

A Recovery Enterprise Manager registered at the Console maintains a lightweight TCP connection with all CounterACT devices in the corporate network. The purpose of this connection is to:

- Verify that the Recovery device can connect to other Forescout components.
- Transmit primary Enterprise Manager system settings to the Recovery device.

This connection is used to manage network Appliances when the Recovery Enterprise Manager is switched over as the primary Enterprise Manager.

Communication between the Enterprise Manager and the Recovery Enterprise Manager is performed on port 13000/TCP using standard TLS encryption.

You may set up one Recovery Enterprise Manager in your enterprise.

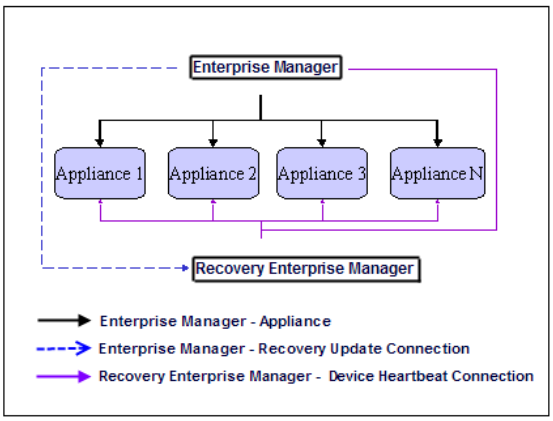

Requirements

- Disaster recovery may only be deployed on networks that have installed Forescout components running version 6.3.X or above.
- The Recovery Enterprise Manager may not manage Forescout managed Appliances

Configuring a Disaster Recovery System

To set up and initiate a disaster recovery system:

1. Verify that you installed the Recovery Enterprise Manager at the Data Center.
2. Log in to the Enterprise Manager via the Console.
3. Select **Options** from the **Tools** menu.
4. Select **CounterACT Devices**.

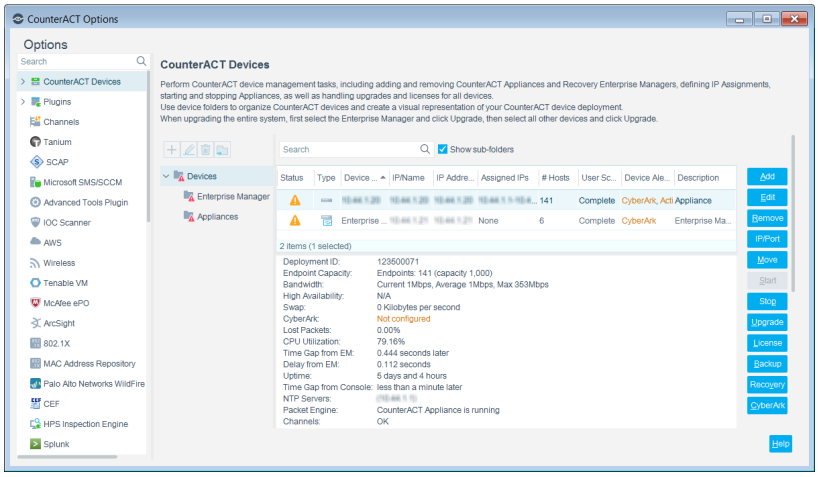

5. In the CounterACT Devices pane, select **Recovery**. The Define Recovery Enterprise Manager dialog box opens.

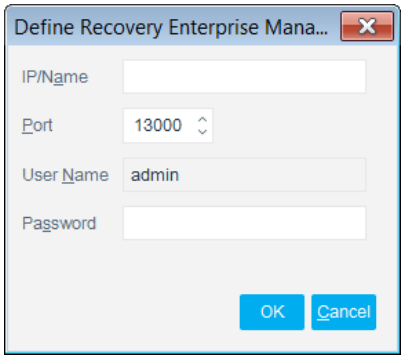

6. Type the IP address of the Recovery device.
7. The default port should be 13000. This enables network connection. It is not recommended to change this value.
8. Type a password.
9. Select **OK**. The Recovery Enterprise Manager is registered. You are notified that the current configuration and system settings for the Recovery Enterprise Manager (defined during installation at the Data Center) are replaced with the configuration and settings of the current Enterprise Manager. Refer to the section on CounterACT Device Management Overview in the *Forescout Administration Guide* for a list of these settings. See Additional Forescout Documentation for information on how to access the guide.

Activating the Switchover

To activate the switchover:

1. Log in to the Recovery Enterprise Manager device via the Console. You are notified that the machine is operating as the Recovery Enterprise Manager, and are asked to confirm the switchover.

2. Select **OK**.

   The devices are switched, i.e., the Recovery device acts as the primary and the original primary acts as the secondary.

   The status of the Enterprise Manager (now the Recovery Enterprise Manager) is disconnected in the case of an actual disaster or for any other reason it is no longer functioning.

Tracking Recovery Enterprise Manager Activity

Information about recovery activity is automatically shown in the Forescout Audit Trails log and Events Viewer.

Audit Trails Log

The Audit Trails log indicates:

- When a Recovery Enterprise Manager was defined
- When the switchover is made to the secondary Enterprise Manager

To access the Audit Trails log:

1. In the Console **Log** menu, select **Audit Trails**.
2. Enter a time period and select **OK**. The Audit Trails log opens.

Event Viewer

The Event Viewer indicates:

- When a switchover is made between the primary and secondary
- When the connection status of the Recovery Enterprise Manager changes

To access the Event Viewer:

1. Select **Event Viewer** from the **Log** menu. The Event Viewer opens.

Viewing Recovery Enterprise Manager Connection Status and Details

Recovery Enterprise Manager details are displayed in the CounterACT Devices pane. Refer to the section on viewing information about CounterACT Devices in the *Forescout Administration Guide* for information about how to work in the pane. See Additional Forescout Documentation for information on how to access the guide.

Glossary

This glossary provides a brief description of terms related to Forescout resiliency and recovery solutions. References to relevant sections in the *Forescout Administration Guide* are also included.

Failover Clustering Terms

Term	Description	See also
Failback	Occurs after a failed Appliance reconnects to the Enterprise Manager, and regains control of the endpoints and network devices previously distributed to recipient Appliances.	Failback
Failed Appliance	An Appliance that has a malfunction. This occurs when the Enterprise Manager fails to connect to the Appliance at the operating system level for a defined period of time. Attempts by the Enterprise Manager to connect to it through port 13000 result in connection time outs. For example, if the Appliance is unplugged or in the event of a natural disaster.	About Failover and Failback
Failover Assignment	The range of endpoints and network devices transferred to a recipient Appliance after failover.	About Failover and Failback
Failover Cluster	A group of geographically/logically connected Appliances, such as those at a data center. Appliances within a failover cluster that fail, can fail over to other Appliances in the cluster.	Failover Clusters
Failover Clustering	A resiliency solution implemented with clusters of Appliances. Redundancy is achieved by defining clusters of Appliances that can automatically take over discovery, assessment and control in case of single or multiple Appliance failure within the cluster/s.	About Forescout Resiliency and Recovery Solutions
Failover Excess	Number of endpoints that exceed Appliance capacity after a failover. These endpoints are not fully handled by the Appliance, meaning that not all of their host properties are fully resolved, and they will not match policies that depend on such unresolved properties. As a result, relevant actions will not be applied to these endpoints.	Handling Endpoints that Exceed Capacity
Failover Scope	A group of one or more failover clusters. Allows failover across clusters.	Failover Scope
Original Assignment	The range of endpoints and network devices originally handled by the recipient Appliance, before failover.	About Failover and Failback

Term	Description	See also
Recipient Appliance	A functioning Appliance managed by the same Enterprise Manager as the failed Appliance which takes over handling of the range of transferred endpoints and network devices after a failover.	About Failover and Failback

High Availability Terms

Term	Description	See also
Active Node	The CounterACT device in a high availability pair that is currently handling discovery, assessment and control of endpoints.	High Availability Pairing
High Availability Pairing	A resiliency solution implemented in pairs of two CounterACT devices. Redundancy is achieved by assigning an Active node and a Standby node. The Standby node automatically takes over discovery, assessment and control in case the Active node fails.	About Forescout Resiliency and Recovery Solutions
Primary Node	The node you use to set up high availability. It is initially the Active node.	High Availability Pairing
Secondary Node	The node that you set up to take over if the primary node fails. It is initially the Standby node.	High Availability Pairing
Standby Node	The backup CounterACT device in a high availability pair that takes over if the Active Node fails.	High Availability Pairing
Pingable Host	IP address of a host that is tested to check if the management interface is Reachable or Not Reachable.	
Out of Band Configuration	This feature defines an additional out of band management interface.	

Disaster Recovery Terms

Term	Description	See also
Disaster Recovery for Enterprise Manager	A recovery solution implemented in pairs of two Enterprise Managers. Redundancy is achieved by defining a Recovery Enterprise Manager that is manually triggered to take over from an Enterprise Manager that is no longer functioning as a result of a disaster.	About Forescout Resiliency and Recovery Solutions
Recovery Enterprise Manager	A backup Enterprise Manager that is manually triggered to take over from an Enterprise Manager that is no longer functioning as a result of a disaster.	About Forescout Resiliency and Recovery Solutions

Additional Forescout Documentation

For information about other Forescout features and modules, refer to the following resources:

- Documentation Downloads
- Documentation Portal
- Forescout Help Tools

Documentation Downloads

Documentation downloads can be accessed from the Forescout Resources Page, or one of two Forescout portals, depending on which licensing mode your deployment is using.

- ***Per-Appliance Licensing Mode*** – Product Updates Portal
- ***Flexx Licensing Mode*** – Customer Portal

Software downloads are also available from these portals.

To identify your licensing mode:

- From the Console, select **Help > About Forescout**.

Forescout Resources Page

The Forescout Resources Page provides links to the full range of technical documentation.

To access the Forescout Resources Page:

- Go to https://www.Forescout.com/company/resources/, select **Technical Documentation** and search for documents.

Product Updates Portal

The Product Updates Portal provides links to Forescout version releases, Base and Content Modules, and eyeExtend products, as well as related documentation. The portal also provides a variety of additional documentation.

To access the Product Updates Portal:

- Go to https://updates.Forescout.com/support/index.php?url=counteract and select the version you want to discover.

Customer Portal

The Downloads page on the Forescout Customer Portal provides links to purchased Forescout version releases, Base and Content Modules, and eyeExtend products, as well as related documentation. Software and related documentation will only appear on the Downloads page if you have a license entitlement for the software.

To access documentation on the Forescout Customer Portal:

- Go to https://Forescout.force.com/support/ and select **Downloads**.

Documentation Portal

The Forescout Documentation Portal is a searchable, web-based library containing information about Forescout tools, features, functionality, and integrations.

If your deployment is using Flexx Licensing Mode, you may not have received credentials to access this portal.

To access the Documentation Portal:

- Go to https://updates.Forescout.com/support/files/counteract/docs_portal/ and use your customer support credentials to log in.

Forescout Help Tools

Access information directly from the Console.

Console Help Buttons

Use context sensitive *Help* buttons to quickly access information about the tasks and topics you are working with.

Forescout Administration Guide

- Select **Forescout Help** from the **Help** menu.

Plugin Help Files

- After the plugin is installed, select **Tools** > **Options** > **Modules**, select the plugin and then select **Help**.

Online Documentation

- Select **Online Documentation** from the **Help** menu to access either the Forescout Resources Page (Flexx licensing) or the Documentation Portal (Per-Appliance licensing).