



Using the Engine Management Service

Neverfail Engine

Notice

Neverfail, LLC has taken all reasonable care to ensure the information in this document is accurate at the date of publication. In relation to any information on third party products or services, Neverfail, LLC has relied on the best available information published by such parties. Neverfail, LLC is continually developing its products and services, therefore the functionality and technical specifications of Neverfail's products can change at any time. For the latest information on Neverfail's products and services, please contact us by email (info@neverfail.com) or visit our Web site (neverfail.com).

Neverfail is a registered trademark of Neverfail, LLC. All third party product names referred to in this document are acknowledged as the trade marks for their respective owner entities.

Copyright (c) 2026 Neverfail, LLC. All rights reserved.

Contents

EMS User Interface

Dashboard

Inventory

Servers

Server Details

Server Status

Server Events

Server Services

Server Data

Server Tasks

Server Rules

Server Shadows

Server Alerts

Server Monitoring

Server Actions

Events

Support

Settings

Compliance

Server Protection

Protect a New Server

Add a High Availability (HA) Instance

Add a Disaster Recovery (DR) Instance

Add both HA and DR instances

Reconstruct or Reconfigure the Protected Cluster

Reclone Secondary or Tertiary instances

Upgrade Applications

[Upgrade Engine on Protected Cluster](#)

[Uninstall Engine from Protected Servers](#)

[Add Already Protected Server to EMS](#)

[Control the Protected Server or Cluster State](#)

[Create VMware SRM Plan](#)

[Manage Server Monitoring](#)

[Manage Server Shadows](#)

[License Server](#)

Best Practices

[Engine Role Transitions](#)

[Configuring Switchover](#)

[Configuring Failover](#)

[Configuring Auto-Switchover](#)

[Public Network Connectivity Loss](#)

[Configuring Isolation](#)

[Failover and Auto-Switchover](#)

[HA Pair, Dedicated NICs](#)

[DR Pair, Dedicated NICs](#)

[HA Pair, Shared NIC](#)

[DR Pair, Shared NIC](#)

[HA + DR Trio](#)

About This Book

The Using Engine Management Service Guide provides information about the Engine Management Service user interface and about the operations available. It also explains a brief set of best practices recommended by Neverfail.

Intended Audience

This guide assumes the reader has a working knowledge of server management, cluster management and VMware vSphere and vCenter software.

Overview of Content

This guide is designed to provide guidance on the utilization of Neverfail Engine Management Service, and is organized into the following sections:

- **About This Book** (this chapter) provides an overview of this guide and the conventions used throughout.
- **EMS User Interface** presents an overview of Neverfail Engine Management Service graphical user interface.
- **Server Protection** describes the operational flows required to fully protect servers using Engine.
- **Best Practices** presents a set of recommendations on server protection.

Document Feedback

Neverfail welcomes your suggestions for improving our documentation and invites you to send your feedback to docfeedback@neverfail.com.

Abbreviations Used in Figures

| Abbreviation | Description |
|--------------|-------------------|
| Channel | Neverfail Channel |

| Abbreviation | Description |
|--------------|---------------------------|
| EMS | Engine Management Service |
| CE | Neverfail Engine |
| NIC | Network Interface Card |
| P2V | Physical to Virtual |
| V2V | Virtual to Virtual |

Technical Support and Education Resources

The following sections describe technical support resources available to you. To access the current version of this book and other related books, go to <https://www.neverfail.com/services-and-support/>.

Online and Telephone Support

Use online support located at <https://www.neverfail.com/services-and-support/> to view your product and contract information, and to submit technical support requests.

Support Offerings

To find out how Neverfail Support offerings can help meet your business needs, go to <https://www.neverfail.com/services-and-support/> .

Neverfail Professional Services

Neverfail Professional Services courses offer extensive hands-on labs, case study examples, and course materials designed for use as on-the-job reference tools. Courses are available on site, in the classroom, and live online. For the day-to-day operations of Neverfail Engine, Neverfail Professional Services provides offerings to help you optimize and manage your Neverfail Engine servers. To access information about education classes, certification programs, and consulting services, go to <https://www.neverfail.com/services-and-support/> .

Neverfail Engine Documentation Library

The following documents are included in the Neverfail Engine documentation library:

| Document | Purpose |
|------------------------------------|---|
| Installation Guide | Provides detailed setup information. |
| Using Neverfail EMS | Provides detailed usage instructions for Engine Management Service. |
| Administrator's Guide | Provides detailed configuration and conceptual information. |
| Deploying to AWS Cloud Environment | Deploying Neverfail Engine in Amazon Web Services Cloud Environment. |
| SCOPE Data Collector | SCOPE Data Collector Service Overview. |
| Release Notes | Provides late-breaking information, known issues, and updates. The latest Release Notes can be found at https://www.neverfail.com/services-and-support/ . |

Conventions

The documentation uses consistent conventions to help you identify items throughout the printed and online library.

| Convention | Specifying |
|----------------------------------|--|
| Bold | Window items including buttons. |
| <i>Italics</i> | Book and CD titles, variable names, new terms, and field names. |
| Fixed font | File and directory names, commands and code examples, text typed by you. |
| Straight brackets, as in [value] | Optional command parameters. |
| Curly braces, as in `` | Required command parameters. |
| Logical OR, as in value1 value2 | Exclusive command parameters where only one of the options can be specified. |

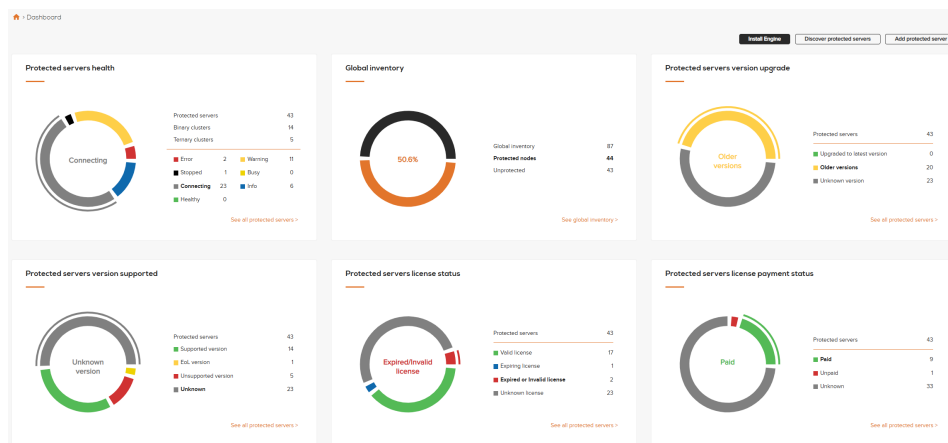
EMS User Interface

The EMS User Interface chapter describes in detail the graphical user interface of Neverfail Engine Management Service.

- **Dashboard**
- **Inventory**
- **Servers**
- **Server Details**
- **Server Status**
- **Server Events**
- **Server Services**
- **Server Data**
- **Server Tasks**
- **Server Rules**
- **Server Shadows**
- **Server Alerts**
- **Server Monitoring**
- **Server Actions**
- **Events**
- **Support**
- **Settings**
- **Compliance**

Dashboard

The Engine Management Service (EMS) Dashboard provides a synthetic aggregation of the most relevant information about your protected servers (dashboard panels) and quick ways to access EMS core functionality (protect servers).



Server protection menu

The top-right buttons available in the Dashboard page provide a quick way to access the core functionality of the EMS: **server protection**. Each of the tree buttons starts a specific protection flow, described in detail in the **Server Protection** chapter of this Help section.

- **Install Engine:** starts the process of protecting (installing Engine) on a new server.
- **Discover protected servers:** starts the process of discovering servers that run Engine in your network, and add them to EMS.
- **Add protected server:** starts the process of adding a known server, running Engine, to EMS.

Protected server health

The **Protected server health** panel aggregates configuration information and basic status for all of your currently protected servers:

- the number of protected servers, binary clusters and ternary clusters;
- the number of servers outputting error, warning or info messages;
- the number of servers in either healthy, busy, connecting or stopped states;

Hovering over any displayed status will update the chart representation, highlighting the hovered information.

Clicking on any status will take you to the **Servers page**, showing you the actual server(s) affected by the selected state.

The **See all protected servers** link will also take you to the **Servers page**, displaying all currently protected servers, regardless of their state.

Global inventory

The **Global inventory** panel provides a quick overview of the total number of servers in your inventory, further divided in protected and unprotected servers.

The **See global inventory** link will take you to the Inventory page, listing all the machines available in your vCenter Inventory.

Protected servers version upgrade

The **Protected servers version upgrade** panel displays the Engine upgrade information across all your protected servers. The version information is aggregated in three version types for which the number of servers is shown:

- Older versions
- Unknown versions
- Latest version

Hovering over any version type will update the chart representation, highlighting the hovered information.

Clicking on any version type will take you to the **Servers page**, displaying the servers running the selected version of Engine.

The **See all protected servers** link will also take you to the **Servers page**, displaying all currently protected servers, regardless of their version.

Protected servers license status

The **Protected servers license status** panel shows the license status of Engine running on your protected servers. The license states can fall in one of four categories, each one displaying the number of servers in it:

- Expired or Invalid license
- Unknown license
- Expiring license
- Valid license

Hovering over any license category will update the chart representation, highlighting the hovered information.

Clicking on any license category will take you to the **Servers page**, listing the servers with licenses in the selected category.

The **See all protected servers** link will also take you to the **Servers page**, displaying all currently protected servers, regardless of their license.

Recent activity

The **Recent activity** panel lists the most recent EMS events in a simple table view, while providing the option to open the **Events page** for more options.

The event entries can be sorted based on each column's data.

Inventory

The **Inventory** page integrates the contents of your **vCenter Inventory** in EMS, in an easy-to-use table format.

All machines, virtual or physical, available in your vCenter Inventory, are ready for exploration and **protection**, thanks to the easy-to-use tabular presentation.

| Host/Virtual Machines | Protection Virtual Machines | Status | Actions |
|-----------------------|-----------------------------|------------------------------|---------|
| 192.168.83.10 | 0 | 0 Protected / 0 Unprotected | ▼ |
| 192.168.83.11 | 16 | 8 Protected / 8 Unprotected | ▼ |
| 192.168.83.19 | 12 | 0 Protected / 12 Unprotected | ▼ |
| 192.168.83.20 | 17 | 1 Protected / 16 Unprotected | ▼ |
| 192.168.83.31 | 12 | 6 Protected / 7 Unprotected | ▼ |
| 192.168.83.32 | 15 | 9 Protected / 6 Unprotected | ▼ |
| 192.168.83.33 | 12 | 3 Protected / 7 Unprotected | ▼ |
| 192.168.83.34 | 12 | 6 Protected / 6 Unprotected | ▼ |

Per guest server protection

Each host machine, displayed as a table row, can be expanded to reveal its guest nodes, for which the you can start the protection process individually.

Only the powered-on machines, running **VMware Tools** are available for Engine protection.

Search, filtering and navigation

The dedicated **search box**, available at the top of the Inventory page, allows you to search your vCenter Inventory host and guest machines, using their name or IP address.

The list of host and guest machines can be filtered based on their categories, using the **category tags** available at the top of the table.

For example, this could enable you to quickly identify machines from a specific physical location, provided that the vCenter machines are categorized based on their location.

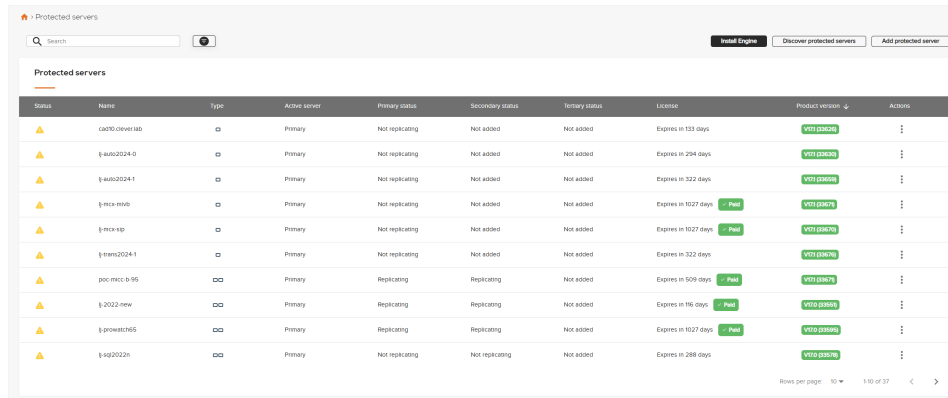
The presentation table also allows you to sort the contents based on the information available in any of its columns.

Navigation between multiple pages of the table is also possible using the dedicated controls available at the bottom of the table.

Servers

The **Servers** page provides an overview of the complete set of protected servers, or in other words, servers which run Engine and are managed by the Engine Management Service.

It also allows you to apply advanced filters to the list of protected servers, engage in common protected server actions, and, using the **Server protection menu**, to add new servers in this list.



| Status | Name | Type | Active server | Primary status | Secondary status | Tertiary status | License | Product version | Actions |
|--------|---------------|------|---------------|-----------------|------------------|-----------------|----------------------|-----------------|---------|
| ▲ | cm370-0b6e1ab | □ | Primary | Not replicating | Not added | Not added | Expires in 131 days | V10.03070 | ⋮ |
| ▲ | §-auto2024-0 | □ | Primary | Not replicating | Not added | Not added | Expires in 294 days | V10.03020 | ⋮ |
| ▲ | §-auto2024-1 | □ | Primary | Not replicating | Not added | Not added | Expires in 322 days | V10.03050 | ⋮ |
| ▲ | §-rckc-0b6e | □ | Primary | Not replicating | Not added | Not added | Expires in 1027 days | V10.03070 | ⋮ |
| ▲ | §-rckc-0b6e | □ | Primary | Not replicating | Not added | Not added | Expires in 1027 days | V10.03070 | ⋮ |
| ▲ | §-rckc-0b6e | □ | Primary | Not replicating | Not added | Not added | Expires in 322 days | V10.03070 | ⋮ |
| ▲ | §-rckc-0b6e | □ | Primary | Not replicating | Not added | Not added | Expires in 322 days | V10.03070 | ⋮ |
| ▲ | §-rckc-0b6e | □ | Primary | Not replicating | Not added | Not added | Expires in 509 days | V10.03070 | ⋮ |
| ▲ | §-2022-new | □ | Primary | Replicating | Replicating | Not added | Expires in 116 days | V10.03050 | ⋮ |
| ▲ | §-2022-new | □ | Primary | Replicating | Replicating | Not added | Expires in 116 days | V10.03050 | ⋮ |
| ▲ | §-2022-new | □ | Primary | Replicating | Replicating | Not added | Expires in 1027 days | V10.03050 | ⋮ |
| ▲ | §-2022-new | □ | Primary | Not replicating | Not replicating | Not added | Expires in 288 days | V10.03070 | ⋮ |

Each protected server row can be expanded, using a simple click on the corresponding row (or using the right-side expansion arrow), to show additional information like License and Product Version.

Clicking a server name opens the **Server Details** page, described in details in the next section.

Server protection menu

The top-right buttons available in the *Servers* page provide a quick way to access the core functionality of the EMS: **server protection**. Each of the tree buttons starts a specific protection flow, described in detail in the **Server Protection** chapter of this Help section.

- **Install Engine:** starts the process of protecting (installing Engine) on a new server.
- **Discover protected servers:** starts the process of discovering servers that run Engine in your network, and add them to EMS.
- **Add protected server:** starts the process of adding a known server, running Engine, to EMS.

Protected servers table

The Engine protected servers are displayed in a table format, showing the following server information in each column:

- **Status:** the graphical representation the state of the protected server. The displayed icons encompass the overall state of the server which aggregates different statuses (instance availability, replication status, license status) into an easy-to-spot state. The aggregation is done based on individual states severity and the resulting state is color coded as follows:
 - **Error** (red error icon): an error degrades the state of the protected server to a non-functioning state.
 - **Warning** (yellow warning icon): a warning degrades the state of the protected server to a functional but requiring attention state.
 - **Info** (blue info icon): the state of the protected server is functional but an information message is available.
 - **Healthy** (green check mark icon): the state of the protected server is functional with no errors, warnings or information messages.
 - **Stopped** (grey stop icon): the Engine service is stopped on the protected server.
 - **Connecting** (grey reconnect icon): the protected server is in an unknown state due to the communication between the EMS and the Primary instance is unavailable.
- **Name:** the name of the protected server
- **Type:** the graphical representation of the protected server configuration type. It can be:
 - **Pair:** the protected server has a Secondary instance. Data can replicate on the Secondary instance, which can be used either as a High Availability instance or a Disaster Recovery instance.
 - **Trio:** the protected server has both Secondary and Tertiary instances. Data can replicate to the Secondary instance (High Availability) and from the Secondary to the Tertiary instance (Disaster Recovery).
- **Active Server:** the server instance that is currently active, meaning that it serves its clients and is the source of replication for the other instances (if existing). Can be: Primary, Secondary, Tertiary or Unavailable.

The **Unavailable** status signals that active server identity cannot be determined (service stopped or instance disconnected).

- **Primary Status:** the status of the primary instance of the protected server. Can be: Replicating, Not Replicating or Unavailable.

The **Unavailable** status signals that instance status cannot be determined (service stopped or instance disconnected).

- **Secondary Status:** the status of the secondary instance of the protected server. Can be: Replicating, Not Replicating, Not Added or Unavailable.

The **Unavailable** status signals that instance status cannot be determined (service stopped or instance disconnected).

- **Tertiary Status:** the status of the tertiary instance of the protected server. Can be: Replicating, Not Replicating, Not Added or Unavailable.

The **Unavailable** status signals that instance status cannot be determined (service stopped or instance disconnected).

- **Actions:** the three dots button allows you to expand the **Quick server actions** menu, revealing the following possible actions:
 - **Upgrade:** starts the upgrade Engine procedure on the server.
 - **License server:** starts the licensing procedure for the selected server.
 - **Make Primary active:** turns the Primary instance to active, if the currently active instance is either the Secondary or Tertiary instance.
 - **Make Secondary active:** turns the Secondary instance to active, if the currently active instance is either the Primary or Tertiary instance.
 - **Make Tertiary active:** turns the Tertiary instance to active, if the currently active instance is either the Primary or Secondary instance.
 - **Start replication:** starts the data replication process between the available instances, if the replication is currently stopped.
 - **Stop replication:** stops the data replication process between the available instances, if the replication is currently started.
 - **Start applications:** starts the server applications managed by Engine application-specific plugins.

- **Stop applications:** stops the server applications managed by Engine application-specific plugins.
- **Clear application health:** clears the health statistics of the server applications managed by Engine application-specific plugins.
- **Cancel application start/stop:** force stops the starting or stopping procedures of server applications managed by Engine application-specific plugins.
- **Add standby servers:** starts the procedure of defining, reconstructing or reconfiguring a High Availability and / or Disaster Recovery instance.
- **Upgrade applications:** starts the procedure of upgrading the server applications managed by Engine.
- **Reclone Secondary or Tertiary:** starts the procedure of recloning a High Availability and / or Disaster Recovery instance.
- **Create VMware SRM plan step:** Create a script to initiate a switch-over of the current server as part of an SRM recovery plan.

Requires Powershell V2 on the SRM server and permission for powershell scripts to run locally without signing. For servers which are members of Business Application Groups, all members of a group will failover or switchover together. It is recommended to add only the 'First to switch' server of a group to the SRM plan.

- **Check filesystem:** starts the procedure of filesystem check on the selected server. Depending on the cluster configuration, the check can be done on the available instances (Primary, Secondary, Tertiary).
- **Check Primary registry:** starts the procedure of registry check on the selected server. Depending on the cluster configuration, the check can be done on the available instances (Primary, Secondary, Tertiary).
- **Check Primary for orphaned files:** verifies the selected server for presence of files that no longer exist on the replication source. Depending on the cluster configuration, the verification can be done on the available instances (Primary, Secondary, Tertiary).
- **Startup Engine service:** starts the Engine service on the selected instances of the current server.
- **Shutdown Engine service:** stops the Engine service on the selected instances of the current server.

- **Uninstall Engine:** removes the Engine service and its corresponding files as well as removes any existing standby (Secondary and/or Tertiary) instances.
- **Remove:** stops the server by being managed by EMS, but does not remove the Engine service or files from the server.

Hovering the mouse cursor over **any unavailable Action** will display the reason for which that specific action is not available.

Operations in progress

Whenever an operation like **Protect, Add standby, Upgrade, Uninstall** or **Reclone** is executed on a server, the **Operations in progress** table is displayed at the top of the **Servers** page. EMS also notifies you about the currently running operations in the global top menu, along with a badge indicating the number of concurrent running tasks. Clicking the notification will take you to the Servers page, where the operation in progress is displayed.



The table displays the following information about the currently running operation:

- **Name:** the name of the server on which the operation is running.
- **Progress:** the progress of the operation being executed.
- **Details:** the operation being executed on the server.
- **Actions:** acknowledge then delete unfinished/orphan/failed operations.

Search, filtering and navigation

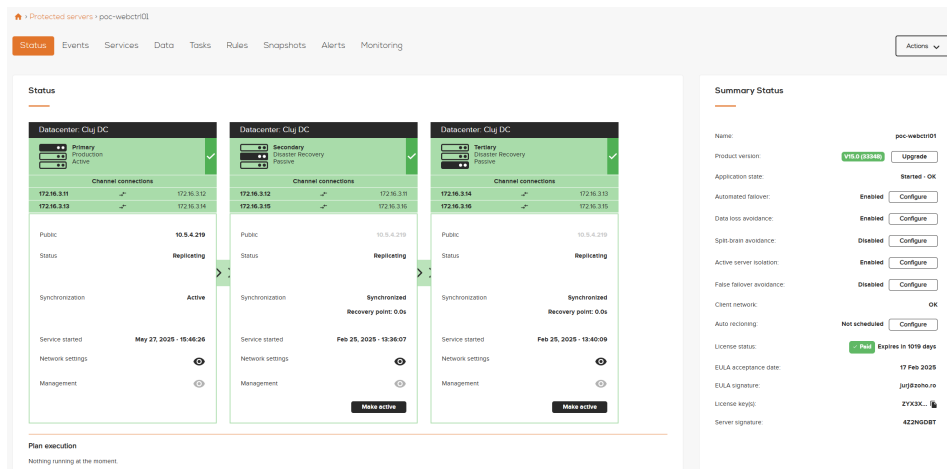
The protected server list is searchable using the Search bar in the top section of the Servers page. The search string is matched against all text fields in the protected servers table.

The Filter options, available next to the Search bar, allows the filtering of the listed servers based on the following server properties: - Active Server (instance) - Product Version (Engine version) - Type (cluster configuration type). - Status (server status) - License Status - Version Upgrade Status (unknown, old or latest version)

The bottom side of the protected servers table provides navigation and display controls for managing long lists of protected servers.

Server Details

The **Server Details** page aggregates all the information and management options of a protected server. It allows you to monitor the state of a server and of the applications managed by Engine application-specific plugins.



The Server Details page information is organized in views, which can be selected at the top of the page. Each view provides access to a different set of server information or options:

- **Status:** displays the protected server topology and detailed status along with the status of the protected applications. It also provides access to topology management actions for the represented server.
- **Events:** displays the particular events of the current server.
- **Services:** displays the protected applications services and their dependency services and provides access to the protected services management actions.
- **Data:** provides an overview of the data synchronization between the instances of the protected cluster.
- **Tasks:** displays the server tasks and allows you to define and manage tasks that run on your protected server.
- **Rules:** displays an overview of the Engine rules available for your configuration.
- **Shadows:** provides access to the protected cluster shadow functionality.
- **Alerts:** provides access to the protected server's alert (event notifications) settings.
- **Monitoring:** provides access to the server and network monitoring settings as well as the applications monitoring configuration

Server protection actions are available in the top-left **Actions** menu, across all the Server Details page's views (read more about the available actions at the bottom of this article).

Server Status

The **Status** view of the Server Details page aggregates the most important information about the state of the protected servers, the Engine running on it and various server applications managed by the Engine plugins. The information is structured in three panels: **Status**, **Summary Status** and **Applications and Platforms**.

Status panel

The **Status** panel provides a comprehensive overview of the protected server replication topology. It depicts, in a graphical manner, the currently defined instances in the replication chain. The order of displaying the server instances is based on the replication direction, from left to right (from the active instance to the first passive and from the first passive to the second passive, if available).

The graphical representations display detailed information about each available instance:

- Each instance is described in the top part of the graphical representation based on its identity (Primary, Secondary, Tertiary), role (Production, High Availability, Disaster Recovery) and state (active, passive).
- **Channel connections** are listed for each instance:
 - For the Primary instance, the first connection represents the **Primary-to-Secondary (Pri-Sec)** connection, depicted by the link and direction of the two Channel IPs dedicated to this particular Pri-Sec connection. The second connection represents the **Primary-to-Tertiary (Pri-Ter)** connection, depicted by the link and direction of the two Channel IPs dedicated to this particular Pri-Ter connection.
 - For the Secondary instance, the first connection represents the **Secondary-to-Primary (Sec-Pri)** connection, depicted by the link and direction of the two Channel IPs dedicated to this particular Sec-Pri connection (reverse direction of Pri-Sec connection). The second connection represents the **Secondary-to-Tertiary (Sec-Ter)** connection, depicted by the link and direction of the two Channel IPs dedicated to this particular Sec-Ter connection.
 - For the Tertiary instance, the first connection represents the **Tertiary-to-Primary (Ter-Pri)** connection, depicted by the link and direction of the two Channel IPs dedicated to this particular Ter-Pri connection (reverse direction of Pri-Ter connection). The second connection represents the **Tertiary-to-Secondary (Ter-Sec)** connection, de-

picted by the link and direction of the two Channel IPs dedicated to this particular Ter-Sec connection (reverse direction of Sec-Ter connection).

There can be multiple channel connections between any two instances. The Pri-Sec, Pri-Ter order is arbitrary. An IP address emphasized on a particular instance means that the IP address is defined on that instance.

- The **Public** IP shows the IP address of the network interface controller (NIC) used for public access. Hovering over the view (eye) icon reveals all the public IPs assigned to this server.
- The **Status** info shows the current instance status. It represents both the replication state and the actual instance activity (like stopped, connecting, cloning, upgrading etc.).
- The **Synchronization** info indicates if the data on the passive instances is synchronized with the data on the active instance. The Recovery Point indicates the queue of data to be synchronized on the passive server.
- The **Service started** info shows the date and time when the Engine service was started on this server instance.
- The **Network settings** info resumes the network configuration when hovering the mouse cursor over the view icon:
 - The public IP addresses of the instance.
 - The channel connections described above for each instance.
 - The available NICs of the instance
 - The IP addresses of the configured Gateway and DNS servers
- The **Management** info shows the Management IP address when hovering the mouse cursor over the view icon.

The inability to retrieve any of the above information is caused either by running Engine product versions older than 8.5 Update 6 or by the Management IP address or Management name being not configured.

Instance status color-coding

While the protected server status is described in detail in the Summary Status panel, the visual representation of each instance is color-coded in order to easily display its current state. The implemented color-coding uses the standard color representations (Green, Yellow and Red) to signal a state aggregation composed of the most relevant metrics, as follows:

- **Green** (working state): all monitored metrics are in good condition.
- **Yellow** (warning state): at least one of the following metrics present warnings.
- **Red** (critical state): at least one of the metrics are in critical condition, rendering the server unprotected.

The aggregated states are based on the metrics listed below.

- **Channel connection:** an unavailable channel connection would make data replication impossible. The active instance would also not be able to gather information about the Passive instance(s). The channel connection can also be unavailable if it is not defined, when no standby instance is added for the protected server.
 - **Management connection:** the lack of management connection would render any instance information unavailable.
 - **Engine product version:** older product versions might not be able to integrate with EMS. The servers should be updated to supported Engine versions.
 - **Applications status:** any application-specific warnings or errors will set this status to warning or error.
 - **Engine License status:** a license that is due to expire will set this status to warning. An expired license will set this status to error and disable the server protection.
 - **Client network:** any client networking errors on the protected server would render the server incapable of conducting its normal operation and would set this status to warning or error.
 - **Passive instance down/unknown state:** if a passive instance is either down or its state is unknown, it would be rendered as disconnected, thus the warning state would be triggered.
 - **Replication status:** when replication is stopped, no application data is replicated and synchronized on the passive instances, triggering the warning state.
 - **Synchronization status:** if instances are out-of-synch, the filesystem is unchecked or synchronization is in progress, the warning state would be triggered.
-

Replication status icons

The state of the replication process between two available instances of the protected server is represented by the icon between the instance representations.

- **Rolling arrows** represent the **replicating state**, where no connection or replication issues exist.
- **Blinking crossed link** represents a **disconnected state** where the channel connection is not available; the two instances are disconnected.
- **Blinking exclamation mark** represents the **not replicating state** where replication is stopped or not possible from other reasons than channel connection unavailability. For example, when the Engine license is expired.

Instance actions

The **Make active** option is present on every standby instance. This allows the user to manually turn any standby passive instance into an active instance.

Plan execution

The **Plan execution** section of the Status panel lists the plans being executed by Engine. **Plans** are sequences of actions required to perform functions such as switch-over or installing a new plug-in. Plans can be executed in response to user action (such as Make active) or automatically (such as failover). Once a displayed plan is complete, it is removed from the Plan Execution pane.

The executed plans are displayed in real time. If a plan execution is very short, it will be displayed briefly or even not displayed at all (if the execution time is shorter than the time required to render the plan in this section).

Summary Status panel

The **Summary Status** panel presents a condensate list of server, network and license statuses.

- **Name:** the name of the protected server.
- **Product version:** the Engine version running on the protected server.
- **Application state:** the status of the applications protected by Engine plugins.

-
- **Automated failover:** the status of the automated failover monitoring feature. The Configure button provides quick access to the automated failover configuration. Check out the **Automated Failover** chapter of the **Manage Server Monitoring** article for more details.
 - **Data loss avoidance:** the status of the data loss avoidance monitoring feature. The Configure button provides quick access to the data loss avoidance configuration. Check out the **Data Loss Avoidance** chapter of the **Manage Server Monitoring** article for more details.
 - **Split-brain avoidance:** the status of the split-brain avoidance feature. The Configure button provides quick access to the split-brain avoidance configuration. Check out the **Split-brain Avoidance** chapter of the **Manage Server Monitoring** article for more details.
 - **Active server isolation:** the status of the active server isolation feature. The Configure button provides quick access to the active server isolation configuration. Check out the **Active Server Isolation** chapter of the **Manage Server Monitoring** article for more details.
 - **False failover avoidance:** the status of the false failover avoidance feature. The Configure button provides quick access to the false failover avoidance configuration. Check out the **False Failover Avoidance** chapter of the **Manage Server Monitoring** article for more details.
 - **Client network:** the state of the network connection used by the server's applications.
 - **Auto recloning:** indicates the schedule of active instance automatic recloning, if enabled.
 - **License status:** the status of the Engine license.
 - **License activation date:** the date when the Engine license was activated.
 - **EULA signature:** the end user license agreement signature represented by the user who verified and accepted the license agreement upon Engine installation.
 - **License key(s):** the Engine license keys deployed on the server.
 - **Server signature:** the signature of the server unique per Engine cluster.

Any disruptive status that requires user action will display a **quick action** option next to the presented status, allowing the user to easily initiate the required action.

For example, if a new Engine version is available, the Product Version status will display the Upgrade quick action. The user can click the Upgrade button to initiate an automatic product upgrade on the server.

Applications and Platforms panel

The **Applications and Platforms** panel lists the server side running applications that are protected by Engine plugins.

Two quick actions are available above the Applications and Platforms panel:

- **Clear application health:** resets the health statuses for all protected applications.
- **Configure:** Provides quick access to the **Edit Application Configuration** panel. Learn more about applications configuration in the **Manage Application Monitoring** chapter of the **Manage Server Monitoring** article.

The table listing the protected applications exposes the following details:

- **Name:** the name of the application protected by a Engine plugin.
- **Health:** the health status of the application, indicated by a corresponding icon. As per the default color coding, a green icon represents good health, a yellow icon represents a warning health state while a red icon represents a critical application health issue.
- **Status:** the textual description of the application status.
- **Actions:** The actions available for the application plugins for which the Engine provides user-editable settings.

Clicking on any listed application (or on the expansion symbol on the far right) will expand the table row, displaying detailed information about the application being protected.

The bottom side of the protected applications table provides navigation and display controls for managing long lists of protected applications.

Server Events

The **Events** view of the Server Details page presents the logged events for the protected server.

| Name | Details | Date Time | Node |
|---|--|-------------------------------|----------|
| Standard compression activated | [*] Standard compression has been activated on this server. The compression system is initializin... | Wed Sep 23 17:00:39 EEST 2020 | TERTIARY |
| Reconciling configuration change | Using value from PRIMARY in resolving difference in configuration of AMF/network.Services On TERTIAR... | Wed Sep 23 17:00:38 EEST 2020 | TERTIARY |
| Reconciling configuration change | Using value from PRIMARY in resolving difference in configuration of Controller/ActivePartition On... | Wed Sep 23 17:00:38 EEST 2020 | TERTIARY |
| Reconciling configuration change | Using value from PRIMARY in resolving difference in configuration of Controller/Replication/Installed O... | Wed Sep 23 17:00:38 EEST 2020 | TERTIARY |
| A channel has connected | Connection made between TERTIARY and PRIMARY. For further details please see KB 1094. | Wed Sep 23 17:00:24 EEST 2020 | TERTIARY |
| Network Monitoring Ping Established | Connection from server to IP0:168.168.1 established. | Wed Sep 23 16:51:15 EEST 2020 | PRIMARY |
| Login | login by Administrator | Wed Sep 23 16:51:06 EEST 2020 | PRIMARY |
| Login | login by Administrator | Wed Sep 23 16:51:06 EEST 2020 | PRIMARY |
| The Full File System Check has finished | For more information please see KB 1088. | Wed Sep 23 16:51:03 EEST 2020 | PRIMARY |
| Protected Files are Synchronized | Files are synchronized between PRIMARY and SECONDARY | Wed Sep 23 16:51:03 EEST 2020 | PRIMARY |

The events available here are generated based on the alerts settings configured in the **Alerts** view.

The events are listed in a table format, displaying the following event details:

- **Name:** the name of the logged server event.
- **Details:** the details of the event.
- **Date Time:** the date and time when the event has occurred.
- **Node:** the instance on which the event has occurred.

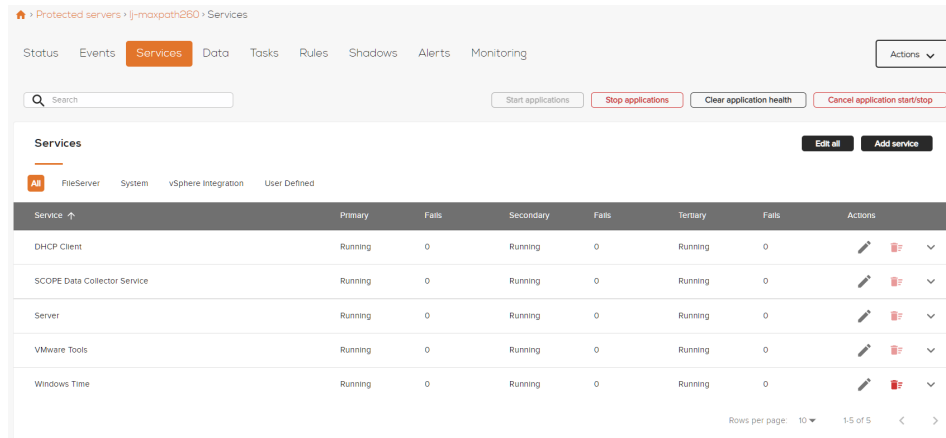
Clicking on any event in the table will expand the detailed information, displaying additional information when available.

The top search bar allows the user to filter the listed events by any textual information provided in the table.

The bottom side of the events table provides navigation and display controls for managing long lists of events.

Server Services

The **Services** view of the Server Details page presents the protected applications services and their dependencies, as well as provides the means of interacting with those services.



The services running behind the protected applications (application services) are displayed in the **Services panel** while the **Dependent Services panel** lists both the services that depend on the application services and the services that the application services depend on.

The search bar at the top of the Services view allows the user to filter the application services by their name.

The actions available on the right side of the search bar are global actions that affect all available application services listed in the Services panel:

- Start applications
- Stop applications
- Clear application health
- Cancel applications start/stop
- Add service
- Edit all

Services panel

The **Services** panel shows the background services of the protected applications running on the server. These can be both services discovered and protected automatically by the Engine installed plugins or services added manually by the user in the protected set (user defined).

The list of services can be filtered based on applications by selecting an application tag from the Services panel header. By default, the panel displays the services of all protected applications.

The following service information is available in the table:

- **Service:** the service display name.
- **Primary:** the status of the service on the Primary instance.
 - **Fails:** the number of failures of the service on the Primary instance.
- **Secondary:** the status of the service on the Secondary instance.
 - **Fails:** the number of failures of the service on the Secondary instance.
- **Tertiary:** the status of the service on the Tertiary instance.
 - **Fails:** the number of failures of the service on the Tertiary instance.
- **Actions:** the user actions available for the service.
 - **Edit:** allows the user to open the **Edit service protection** panel and manage the settings of the service protection (read more in the **Manage Server Application** article).
 - **Delete:** if the service is a user added service, allows the user to remove it and stop protection for the represented service.

Clicking any service row in the table or the expand icon in the far right side of a service row will expand the service, revealing detailed information about its protection setup:

- **Service:** the name of the service itself.
- **Target state on active server:** the expected (correct) state of the service on the active instance of the protected server.
- **Target state on passive server:** the expected (correct) state of the service on the passive instance of the protected server.
- **Original startup type:** the start-up type type of the service when installed on the operating system, before being managed by Engine.
- **Managed:** indicates if the applications service is managed (started, stopped) by Engine.

- **Monitored:** indicates if the applications service is monitored by Engine.
- **On first failure:** the action initiated by Engine on the first failure of the application service.
- **On second failure:** the action initiated by Engine on the second failure of the application service.
- **On third failure:** the action initiated by Engine on the third failure of the application service.
- **Allocate entire application time-out when recovering service:** indicates if the amount of time to wait for service recovery is the total configured recovery time-out for a protected application.

The bottom side of the Service panel provides navigation and display controls for managing long lists of application services.

Dependant Services panel

The **Dependant Services** panel lists the services that depend on the application services or the ones which the application services depend on.

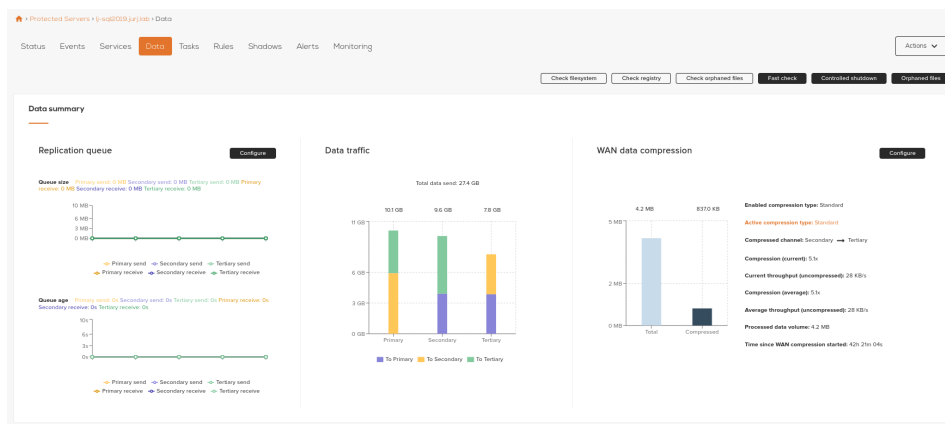
These dependency services **may be monitored but not managed** by Engine. This means that a service depending on a protected application service is not started if stopped by external agency. Nevertheless, a service that a protected application service depends on will be restarted if stopped by external agency, since Engine must ensure the running state of the application service. Most of these dependencies are services installed by the operating system by default or as additional features.

The displayed application services dependencies can be switched between **services that protected services depend on** and **services that depend on protected services** using the dependency tags in the panel's header.

The application services dependencies are listed in a table format, showing the same details as the Service panel table. Since the dependency services are not managed by Engine, no user actions (edit, delete) are available.

Server Data

The **Data** view of the Server Details page provides an overview of the data synchronization between the instances of the protected cluster. It also provides direct methods of checking file system and registry integrity as well as orphaned files.



Data Summary

The **Data Summary** panel shows the graphical representations of data traffic between instances: replication queue, data traffic and WAN compression.

The **Replication Queue** graph summarizes the state of the **send** and **receive** replication queues for all instances in the protected cluster. The summary is displayed in two graphs, the first one plotting the sent and received data for all instances based on the queue size. The second graph plots the send and received data based on the age of the queue.

The **Replication Queue** represents a buffer disk space for data queued for synchronization on all cluster instances. The queue disk size can be configured by the user using the **Configure** button next to the chart.

The **Data Traffic** graph plots the total traffic between protected cluster instances. The traffic is represented by stacked chart bars showing the source and destination instances.

The **WAN Data Compression** chart shows the transferred data compressed by the WAN compression feature of Engine. The graph plots the total and compressed data for the channel for which data compression is active. The chart legend shows detailed information about the WAN compression configuration and status.

The **WAN Compression** feature enables Engine to reduce the size of transferred data over the Wide Area Network (WAN, usually employed for a remote Tertiary instance in a protected cluster). It uses advanced compression mechanisms like deduplication to achieve the minimum data transfer over the WAN channel, in order to keep data in synch on a remote instance. WAN Compression can be configured by the user using the **Configure** button next to the chart. The **Configure WAN Compression** dialog allows the user to choose the enabled compression type, the active compression type and the location of the compression logs.

The WAN Data Compression type can be set to:

- **Auto:** Engine selects the level of WAN compression based upon current configuration without user intervention. This is the recommended setting.
- **Standard:** Engine uses compression on data before it is sent across the WAN to improve WAN data throughput speed.
- **Advanced:** Engine uses the WAN Deduplication feature in addition to compression to remove redundant data before transmitting across the WAN thereby increasing critical data throughput.
- **None:** Selected when deployed in a LAN or where WAN Compression is not required.

Inclusion and Exclusion Filters

The **Inclusion Filters** and **Exclusion Filters**, available in their dedicated panels below the Data Summary, allow you to control the data added to the replication queue and synchronized across the cluster instances.

Inclusion Filters are used to add data (using pattern matching for a file or multiple files) to the replication queue. This data is always replicated and protected in the cluster.

Exclusion Filters are used to exclude specific data from the replication queue (using pattern matching for a file or multiple files). The excluded data is never added to the replication queue thus never synchronized and protected in the cluster.

The Inclusion or Exclusion Filters panels list the filters defined automatically by Engine or user defined filters. The following filter properties are listed in the panel's columns:

- Path: the path to the directory where the pattern is applied for filtering.
- State: the state of the filter.
- Details: the details of the filter.
- Actions: the available actions for the filter
 - Edit filter: allows you to use the **Edit inclusion/exclusion filter** dialog to modify the filter's pattern or state (enabled or disabled).
 - Delete filter: removes the filter.

The **Add Inclusion/Exclusion Filter** dialog can be accessed using the **Add Inclusion/Exclusion** button on the top-right side of each filter panel. It allows you to define a file filter by setting a path and a pattern to match your desired file or files.

The filter can be specified either by giving their complete path, or by specifying a pattern containing wildcards. The two forms of wildcards available are '*', which matches all files in the current folder or '**', which matches all files, subfolders, and the files in the subfolders of the current folder.

Data Actions

The **Data** view provides access to several data protection actions, using the six buttons at the top of the view. The following actions are available:

- **Check filesystem:** opens the Check Filesystem dialog that allows you to initiate a filesystem scan on a passive instance of the cluster (when available). Depending on the amount of protected data and the available bandwidth, this verify/sync operation may take a long

time to complete (e.g. a number of hours). During this time, you will not be able to make another server active.

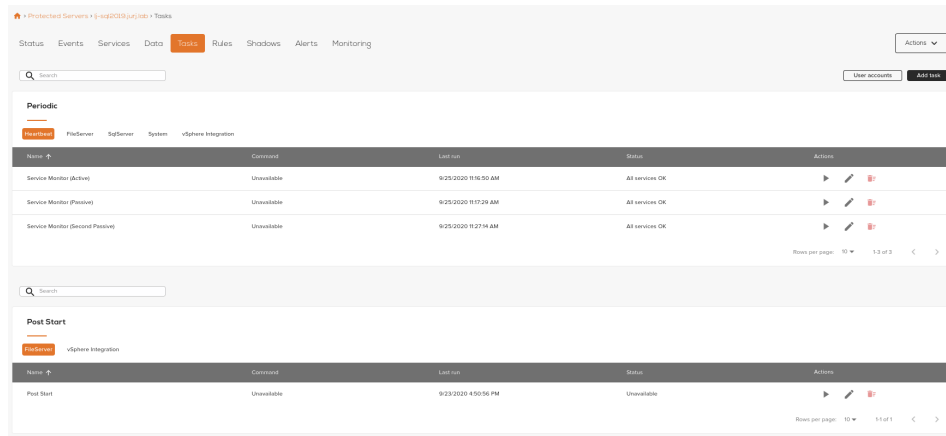
- **Check registry:** opens the Check Registry dialog, allowing you to initiate a registry scan on a passive instance of the cluster (when available).
- **Check orphaned files:** opens the Check Orphaned Files dialog, allowing you to initiate a scan for orphaned files on a passive instance of the cluster (when available).
- **Fast check:** opens the Fast Check dialog, allowing you to enable Fast Check for the filesystem check and initiate it. Fast Check provides an alternative to full data verification by only checking file attributes instead of all the data, and thereby offering greater performance. In addition, delaying application start up until servers connect and replication starts, can significantly improve fast check performance for database servers and Exchange systems.
- **Controlled shutdown:** opens the dialog that allows the cluster to remain in sync, when the server shuts down (or is restarted) during specific hours. Shutdown may be delayed up to a timeout while applications stop and replication data is sent.

You can select to control the instance shutdown on the passive instance(s) and specify the shutdown max time for each instance specifically (for Primary, Secondary or Tertiary). Shutting down active instance will always maintain sync. The schedule table allows you to configure the schedule of the controlled shutdown, for every hour of the days of the week.

- **Orphaned files:** Opens the dialog that allows you to enable the detection of orphaned files and to set the action taken upon detection: delete found files or log the detection to file.

Server Tasks

The **Tasks** view of the Server Details page is where you can define and manage tasks that run on your protected server. Tasks are basically commands scheduled to run in certain scenarios.



The defined tasks are listed in panels, based on the task type (read more about task types below). The types of tasks employed by default by Engine vary depending on the configuration of your protected server (protected applications, etc). For example, the tasks available by default for an SQL Database server can be shown in the following panels:

- Periodic tasks
- Post Start tasks
- Pre Stop tasks
- Rule Action

Each of the available task types (see below) will be represented in its own panel, if at least one task is defined for that type.

Each panel will display a tabbed list of task sponsors (sponsors are either the protected applications for which the task is employed, or User defined). Same tasks with of the same type and having the same sponsor will be grouped under these tabs.

For each available task panel, the task details are presented in the panel's columns:

- The **Name** of the task
- The **Command** which is run by the task

- The **Last run** of the task
- The **Status** of the task
- Task **Actions**:
 - Run task
 - Edit task: allows the user to edit the task name, type, interval (for periodic tasks only) and enabled/disabled state. Note that a sponsored task may have limited configuration options compared to a user defined task.
 - Remove task: allows the user to delete the user defined tasks. The sponsored tasks cannot be deleted since they are installed by Engine plugins.

The task list presented by the panel can be searched using the dedicated **Search** bar at the top of the panel. The search is performed on the task names.

The bottom section of the panel provides list and page navigation controls, useful for longer lists of tasks.

Tasks actions

The top-right side of the Tasks view shows the **User accounts** and **Add task** buttons. Here you can define the user accounts with enough privileges on the target server under which the tasks will be run, and, of course, you can define the tasks themselves.

Clicking the User accounts button opens the **User accounts** dialog, where you can manage the defined user accounts to run your tasks or you can add new ones.

Clicking the Add task button opens the **Add task** dialog, allowing you to define a new task by configuring the following options:

- **Name**
 - **Type**:
 - **Periodic**: runs the command at periodically, based on the defined interval.
 - **Network Configuration**: runs the command whenever the network configuration changes, for example in a switchover.
 - **Pre Start**: runs the command before starting the protected applications.
 - **Post Start**: runs the command after starting the protected applications.
 - **Pre Stop**: runs the command before stopping the protected applications.
-

- **Post Stop**: runs the command after stopping the protected applications.
 - **Pre Shadow**: runs the command before a shadow copy is performed on the instance.
 - **Post Shadow**: runs the command after a shadow copy is performed on the instance.
 - **Rule Action**: runs a command based on a defined rule (see the Rules view of the Server Details page for more details on rules).
- **Command**: the command to be run in the conditions defined by the task type.
 - **Run as**: the user under which the command is run.
 - **Add user account**: opens the User accounts dialog, allowing you to add a new user account directly from this dialog.

Rule Action tasks panel

The **Rule Action** tasks panel is a particular panel type that lists the tasks that are triggered by the rules defined in the **Rules** view of the Server Details page (read the **Rules view** below for more details).

Like other tasks panels, the Rule Action tasks panel presents the same task details in the panel's columns, groups the tasks in sponsors and allows task filtering using the dedicated search box.

Server Rules

The **Rules** view of the Server Details page displays an overview of the Engine rules available for your configuration. The rules are conditional triggers defined by a plugin sponsor, in order to assess protected application health.

The Rules are defined as negative outcome conditions, meaning that a rule action is triggered on the failure of the specified condition. Some rules specify condition(s) and interval of execution (latched rules). Other specify additionally duration (triggered rules): these have cascading failure configurable actions (On first failure, on second failure, on third failure). Tasks defined as Rule Action can be set to run for a rule's failure actions.

The screenshot displays the 'Rules' view for a server configuration. It is divided into two sections: 'FileServer' and 'SqlServer'. Each section contains a table of rules with columns for Rule, Condition, Duration, Result, Status, Triggered, Trigger count, and Actions.

| Rule | Condition | Duration | Result | Status | Triggered | Trigger count | Actions |
|-------------------------|---|-------------|--------------------------|--------|-------------|---------------|---------|
| FileServer | | | | | | | |
| Active Server Sessions | Server sessions > 5 | 30s 100% | Unavailable | OK | Unavailable | 0 | ▶ ✎ |
| Files Open | Files Open > 100 | 30s 100% | Unavailable | OK | Unavailable | 0 | ▶ ✎ |
| Logons / second | Logons / second > 10 | 30s 100% | Unavailable | OK | Unavailable | 0 | ▶ ✎ |
| Mean response time | Mean response time > 1000 milliseconds | Unavailable | Unavailable | OK | Unavailable | 0 | ▶ ✎ |
| SqlServer | | | | | | | |
| Database Availability | Execute sp_help stored procedure on all databases Timeout after 10000 milliseconds. | Unavailable | IS | OK | Unavailable | 0 | ▶ ✎ |
| Database Online Status | Unavailable | Unavailable | ONLINE | OK | Unavailable | 0 | ▶ ✎ |
| DB File Allocated Space | Compare data files allocated space versus used space. Fillup quota 80 %. | Unavailable | Allocated space in check | OK | Unavailable | 0 | ▶ ✎ |

The available rules on the server protection configuration. Engine adds, by default, rules for managing the system and, if configured, the vSphere integration. On top of this, every sponsor (Engine plug-in) may add its particular rules to Engine and make them available in EMS (for example File Server rules or SQL Server rules).

Rules use the following control and decision criteria for evaluation:

- **Name:** the name of the rule.
- **Enabled:** whether the rule is enabled or not.
- **Condition:** the condition being evaluated.
- **Status:** the current status of the rule being evaluation.

-
- **Triggered:** the condition fails to meet configured parameters resulting in initiation of a duration count.
 - **Triggered Count:** a count of the number of times the rule has failed.
 - **Duration:** the length of time the condition exists before triggering the failure action.
 - **Interval:** the length of time between failure actions.
 - **First Failure:** action to take upon first failure. The default is set to Log Warning.
 - **Second Failure:** action to take upon second failure. The default is set to Log Warning.
 - **Third Failure:** action to take upon third failure. The default is set to Log Warning.

Rules panels

The rules are listed in panels corresponding to their sponsors (System, vSphere Integration, File Server, any Engine plugin). Each panel groups the available rules in categories, displayed as tags at the top of the panel. The rules can be also filtered by their name using the search bar above each rule panel.

Each panel lists the rules are table rows, with the following rule details in each column:

- **Rule**
- **Condition**
- **Duration**
- **Interval**
- **Result**
- **Status**
- **Triggered**
- **Trigger count**
- **Actions:** the **Check Now** and **Edit** rule options.

Disabled rules can be manually run by the user. The panel will display those rules as grayed out, with no input action allowed besides editing the rule.

Check rule

To check a rule condition, select the rule in the appropriate panel page and click **Check Now** button (play icon) in the **Actions** column of the panel.

Engine immediately checks the rule conditions of the current configuration against the attributes of the system and application.

Edit rule

Any rule can be edited by using the Edit option from the Actions column. This option provides access to the **Edit rule** dialog, where the user can enable or disable the rule, edit the condition and failure task(s). Only enabled rules can run.

The Rule condition is displayed under the Enable checkbox and is predefined by Engine or one of its plug-ins. The user can edit specific parameters of that condition, like values in a comparison (for example: Server sessions > 5, where 5 is the user editable parameter).

The Rule actions are user selectable and are always triggered on the failure of the condition:

- **On First Failure**
- **On Second Failure**
- **On Third Failure**

The available actions contain the Engine default actions and the Rule Action tasks defined by the user (see the **Tasks view** chapter above):

- **Recover Service:** Restarts the service.
- **Restart Applications:** Restarts the protected application.
- **Log Warning:** Adds an entry to the logs.
- **Switchover:** Initiates a switchover to the currently passive server.
- **Rule Action:** Executes the command or script previously defined as a Rule Action task.

If the installed servers are in a virtual to virtual configuration, the following additional tasks are available as a result of the vSphere Integration Plug-in:

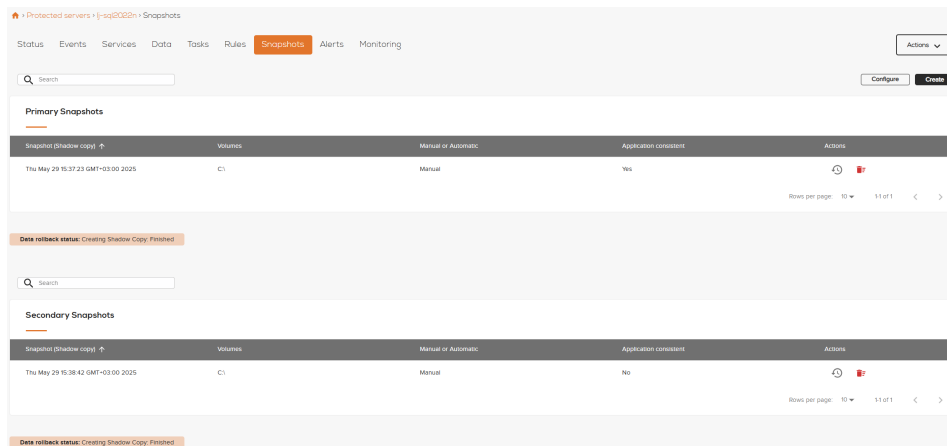
- **vSphere Integration \ RestartVM:** Cleanly shuts down and restarts the Windows OS on the target VM.
-

- **vSphere Integration \ TriggerMigrateVM:** Depending on the parameters specified it can be vMotion, enhanced vMotion or storage vMotion.
- **vSphereIntegration \ TriggerMigrateVMandRestartApplication:** Same as TriggerMigrateVM + application restart.
- **vSphere Integration \ TriggervSphereHaVmReset:** Hard Reset of the VM implemented by integration with VMware HA.

Note: This option requires vSphere HA Application monitoring for the cluster and VM.

Server Shadows

The **Snapshots** or **Shadows** view of the Server Details page provides access to the protected cluster create/rollback shadow functionality. Shadows are data snapshots which can be taken based on a user defined schedule or triggered manually using the UI. Creating shadow copies provides a safe way to roll back data when needed.



The snapshots are listed in panes representing the protected cluster instance on which they were taken: Primary Snapshots, Secondary Snapshots and Tertiary Snapshots.

Each panel presents the available shadow copies as table rows, with the shadow details in each column:

- **Snapshot (Shadow copy):** the name of the shadow copy.
- **Volumes:** the drive volumes on which the snapshot is taken.
- **Manual or Automatic:** the shadow trigger mode. Automatic shadows must be configured using the **Configure** option.
- **Application consistency**
- **Actions:** rollback the current shadow or remove the current shadow.

The shadow panels provide a search functionality for filtering shadows based on their name, as well as pagination options for long shadow lists.

Data rollback status indicates the shadow job status (none, schedule, creating, finished).

Configure automatic shadows

The **Configure** button at the top-right side of the Shadows view provides access to the **Configure Shadows** dialog. Here you can enable automatic shadows creation and maintenance and you can configure the shadow schedule, previous shadow actions and shadow information storage location.

- **Create and maintain shadows automatically:** allows you to enable the automatic shadow creation and maintenance.
- **Create shadow every:** allows you to enable and specify the interval for creating the shadows. Can be 15, 30 or 60 minutes.
- **Create a shadow on the active once per day at:** allows you to enable and specify a once per day schedule for shadow creation on the active instance.
- **Only between the hours:** allows you to enable and specify a time interval in which the shadows can be created.
- **Only on days:** allows you to enable and specify a day interval in which the shadows can be created.
- **For earlier in the current day, keep shadows only at intervals of:** allows you to enable and define the interval between consecutive shadows in the same day that will be kept in storage. The shadows created outside the specified time period will be removed as soon as a new shadow defines the end of the time period.
- **For earlier days in the current week, keep only the shadows nearest:** allows you to enable and the storage of daily shadows, taken during the same week, at (or close to) the specified hour in the day. The daily shadow created at the closest time to the time specified will be kept in storage, the rest will be discarded. This will result in seven stored shadows, all taken daily, as close as possible to the specified hour.
- **For earlier weeks in the current month, keep only the shadows nearest:** similar to the previous option, allows you to enable the storage of weekly shadows taken on the specified day, during a month. The shadows taken in other days of the week will be discarded.
- **Shadow information location:** the storage location of the shadow information. This location must be protected by Neverfail Engine.

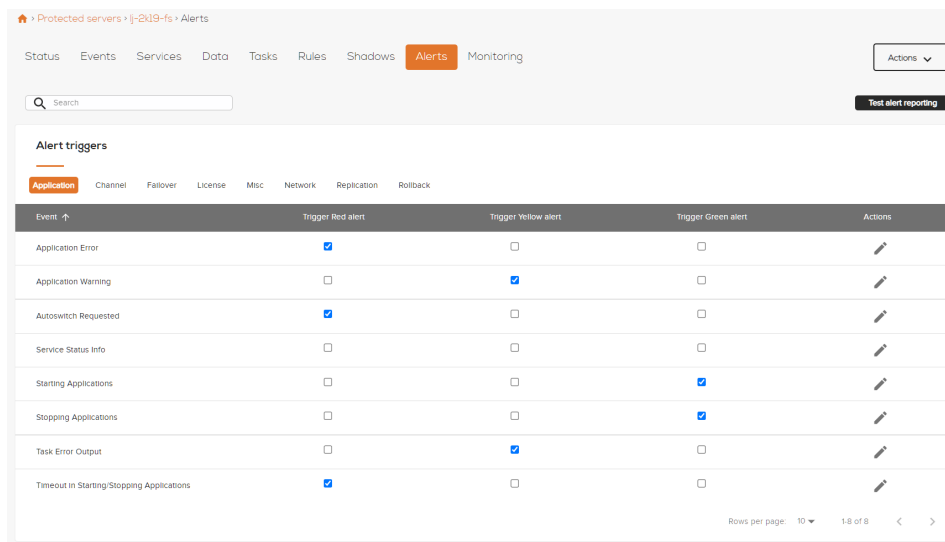
Changing the Shadow information location will cause all old shadows to be deleted.

Create shadows

The Create button, available at the top-right section of the Shadows view, opens the **Create shadow** dialog, allowing you to manually create a shadow for the selected cluster instance.

Server Alerts

The **Alerts** view provides access to the protected server's alert (event notifications) settings. The view's functionalities are split in five panels, allowing you to manage the alert triggers, emailing settings and alerts commands.



The events generated by the alert settings configured here are available in the **Events** view.

Alert triggers

The **Alert triggers** panel displays the available events from each event category (represented in the panel's tab list). For each event, the current alert setup is displayed in the panel's table, along with the edit option, which allows the configuration of the event alerts.

The following event categories are available in the Alert triggers panel as panel tabs:

- Application events
- Channel events
- Failover events
- License events
- Misc events
- Network events

- Replication events
- Rollback events

For each event in each of the above categories, the following alerts can be enabled:

- **Red alert**
- **Yellow alert**
- **Green alert**

These alert types define severity of an event, which is represented by the color and icon displayed in the Events view:

- Blue icon for informative events.
- Green icon for no severity events
- Yellow icon for warning events.
- Red icon for critical events.

The alert email notification recipients, subject and body are also defined by the selected alert type (see Mail Settings). Each alert type can also trigger a specific command, if configured.

The **Actions** button, available in the far right column of each event, allows you to edit the configuration and enable and/or disable the triggering of Red, Yellow and Green alerts for that particular event.

Multiple alerts can be triggered for a single event, as EMS does not limit the number of alert triggers enabled per event. So, theoretically, you could have an event triggering a Red alert, Yellow alert and Green alert in the same time.

The bottom side of the Alert triggers panel provides navigation and display controls for managing long lists of events.

Mail settings

The **Mail settings** panel shows the current email configuration and provides access to the **Configure email settings** dialog.

The following email settings are displayed in the Mail settings panel:

- **Use global configuration**
- **Outgoing mail server for Primary server:** the email server used by the Primary instance to send alert notifications when active.
- **Outgoing mail server for Secondary server:** the email server used by the Secondary instance to send alert notifications when active.
- **Outgoing mail server for Tertiary server:** the email server used by the Tertiary instance to send alert notifications when active.
- **Send mail as:** the email address used as sender in the alert emails.
- **Mail server required authentication:** if enabled, the username and password of the SMTP server must be provided.
 - **Username:** the SMTP server username.
 - **Password:** the SMTP server password.
- **Enable SSL:** if enabled, SSL will be used when connecting to the SMTP server.

Each server instance has its own email server in order to be able to send alert notifications when active, even if each instance is at a different physical site.

The Tertiary instance email server option is available only for trio configurations.

Note: When Engine is protecting an Exchange Server, it is not recommended to configure the alerts to use the protected Exchange server and is advisable if at all possible to use a different Exchange server somewhere else within the organization.

The **Configure** button opens the **Configure email settings** dialog, allowing the user to edit the email configuration detailed above.

Red, Yellow and Green alerts panels

The **Red alerts**, **Yellow alerts** and **Green alerts** panels show the email and command configurations for each of the three alert types. The following details are available:

- **Send mail:** if enabled, the provided email settings will be used to send an alert email whenever an alert from the current category is triggered.
 - **Mail recipients:** the email addresses that will receive the alert email.

- **Mail subject:** the subject of the alert email. Parameters in the form of \$(Parameter-Name) are supported.
- **Mail content:** the body of the alert email. Parameters in the form of \$(Parameter-Name) are supported.
- **Run command:** if enabled, the command specified below will be executed whenever an alert from the current category is triggered.
 - **Command:** any command (or batch file, powershell, executable) configured to be executed on the server when the alert is triggered.

The email subject and body texts can be composed using special parameters provided by EMS. These parameters will be replaced automatically by their corresponding values when the email is sent.

- \$(EventName): the name of the event triggering the alert.
- \$(EventTime): the time of the event occurrence.
- \$(EventDetail): the details of the event triggering the alert.
- \$(EventHostName): the name of the server on which the event occurred.
- \$(EventHostId): the ID of the server on which the event occurred.
- \$(EventHostRole): the role of the server on which the event occurred.

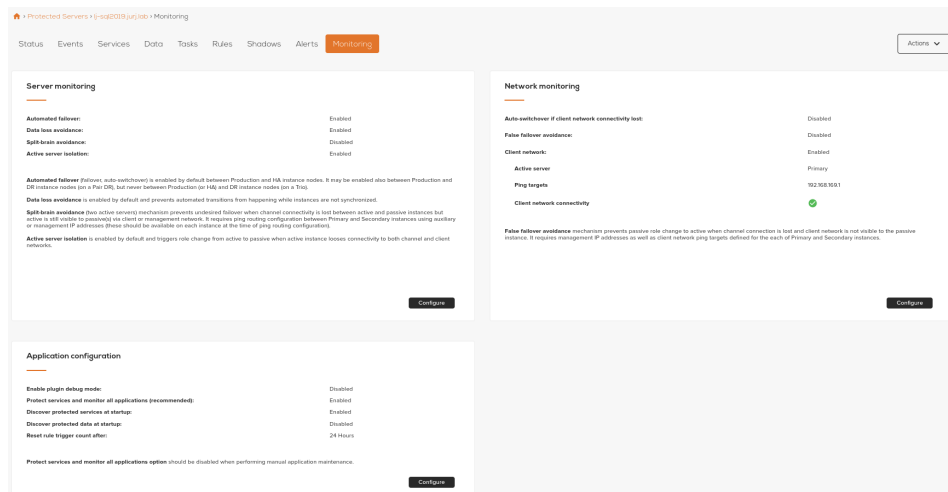
The **Configure** button opens the **Configure [Red, Yellow or Green] alerts** dialog, where the above settings can be configured.

Actions

- Test alert reporting

Server Monitoring

The **Monitoring** view provides access to the server and network monitoring settings as well as the applications monitoring configuration. The three panels available in this view present the current monitoring configuration and provide access to monitoring settings.



Server monitoring

The **Server monitoring** panel shows current state of the monitoring settings applied for the protected server. It also provides a brief description of the available monitoring settings:

- Automated failover
- Data loss avoidance
- Split-brain avoidance
- Active server isolation

The **Configure** button provides access to the **Configure Server Monitoring** dialog, where the options listed above can be configured.

Check out the **Manage Server Monitoring** chapter from the **Manage Server Monitoring** article for more information about the above monitoring options and how to configure them.

Network monitoring

The **Network monitoring** panel shows current state of the available network monitoring settings:

- Auto-switchover if client network connectivity lost
- False failover avoidance
- Client network connectivity

The **Configure** button provides access to the **Configure Network Monitoring** dialog, where the options listed above can be configured.

Check out the **Manage Network Monitoring** chapter from the **Manage Server Monitoring** article for more information about the above monitoring options and how to configure them.

Application configuration

The **Application configuration** panel shows current state of the available application monitoring settings:

- Enable plugin debug mode
- Protect services and monitor all applications (recommended)
- Discover protected services at startup
- Discover protected data at startup
- Reset rule trigger count after

The **Configure** button provides access to the **Edit Application Configuration** dialog, where the options listed above can be configured.

Check out the **Manage Application Monitoring** chapter from the **Manage Server Monitoring** article for more information about the above monitoring options and how to configure them.

Server Actions

The **Actions** button from the top of the Server Details page allows you to quickly initiate actions such:

- **Upgrade:** starts the upgrade Engine procedure on the server.
- **License server:** starts the licensing procedure for the selected server.
- **Make Primary active:** turns the Primary instance to active, if the currently active instance is either the Secondary or Tertiary instance.
- **Make Secondary active:** turns the Secondary instance to active, if the currently active instance is either the Primary or Tertiary instance.
- **Make Tertiary active:** turns the Tertiary instance to active, if the currently active instance is either the Primary or Secondary instance.
- **Start replication:** starts the data replication process between the available instances, if the replication is currently stopped.
- **Stop replication:** stops the data replication process between the available instances, if the replication is currently started.
- **Start applications:** starts the server applications managed by Engine application-specific plugins.
- **Stop applications:** stops the server applications managed by Engine application-specific plugins.
- **Clear application health:** clears the health statistics of the server applications managed by Engine application-specific plugins.
- **Cancel application start/stop:** force stops the starting or stopping procedures of server applications managed by Engine application-specific plugins.
- **Add standby servers:** starts the procedure of defining, reconstructing or reconfiguring a High Availability and / or Disaster Recovery instance.
- **Upgrade applications:** starts the procedure of upgrading the server applications managed by Engine.
- **Reclone Secondary or Tertiary:** starts the procedure of recloning a High Availability and / or Disaster Recovery instance.
- **Create VMware SRM plan step:** Create a script to initiate a switch-over of the current server as part of an SRM recovery plan.

Requires Powershell V2 on the SRM server and permission for powershell scripts to run locally without signing. For servers which are members of Business Application Groups, all members of a group will failover or switchover together. It is recommended to add only the 'First to switch' server of a group to the SRM plan.

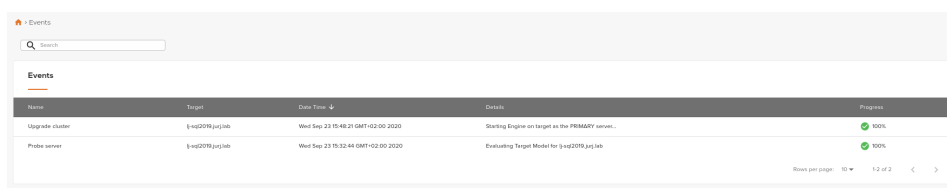
- **Check filesystem:** starts the procedure of filesystem check on the selected server. Depending on the cluster configuration, the check can be done on the available instances (Primary, Secondary, Tertiary).
- **Check Primary registry:** starts the procedure of registry check on the selected server. Depending on the cluster configuration, the check can be done on the available instances (Primary, Secondary, Tertiary).
- **Check Primary for orphaned files:** verifies the selected server for presence of files that no longer exist on the replication source. Depending on the cluster configuration, the verification can be done on the available instances (Primary, Secondary, Tertiary).
- **Startup Engine service:** starts the Engine service on the selected instances of the current server
- **Shutdown Engine service:** stops the Engine service on the selected instances of the current server
- **Uninstall Engine:** removes the Engine service and its corresponding files as well as removes any existing standby (Secondary and/or Tertiary) instances.
- **Remove:** stops the server by being managed by EMS, but does not remove the Engine service or files from the server.

Hovering the mouse cursor over **any unavailable Action** will display the reason for which that specific action is not available.

Events

The **Events** page shows all EMS past and current events, in a table format. The events presented here are scoped across all the servers managed by EMS and only include only the management actions (EMS configuration, protected servers addition, configuration, deletion, etc.).

The activity of the protected clusters and individual instances, as well as their state and their applications states are logged and presented in the Events subsection of each protected server.



The screenshot shows the 'Events' page in the EMS interface. It features a search bar at the top and a table of events below. The table has columns for Name, Target, Date Time, Details, and Progress. Two events are visible: 'Upgrade cluster' and 'Probe server', both showing 100% progress.

| Name | Target | Date Time | Details | Progress |
|-----------------|-----------------|------------------------------------|--|----------|
| Upgrade cluster | §-sp2079.jujlab | Wed Sep 23 15:48:21 GMT+02:00 2020 | Starting Engine on target as the PRIMARY server... | 100% |
| Probe server | §-sp2079.jujlab | Wed Sep 23 15:32:44 GMT+02:00 2020 | Evaluating Target Model for §-sp2079.jujlab | 100% |

The EMS events are displayed in a table format, for each event the following information being presented:

- **Name:** the title of the management action performed on a server.
- **Target:** the server name on which the action has been or is being performed.
- **Date Time:** the date and time when the action has been triggered.
- **Details:** the description of the action, usually describing what actually happens on the protected server.
- **Progress:** the progress and status of the performed action. An action that has started and ended will have a 100% progress, while an action that has not ended will show its current progress (between 0% and 99%, or failed).

Search and navigation

The EMS events can be searched using the top search bar. The search is performed on all string based details presented in the events table.

The bottom side of the events table provides navigation and display controls for managing long lists of EMS events.

Support

The **Support** page of the EMS provides access to:

- **Product feedback:** your feedback is important to us! Use the **Send product feedback** dialog to let us know what we should improve in the Engine Management Service or in Engine. The Neverfail team takes your feedback very seriously and will always work towards perfecting your experience with our products!
- **Help:** everything you need to know so you can manage your Engine infrastructure.
- **Support tickets:** open a new ticket or view your previous tickets and know their statuses and solutions.
- **Knowledge base:** explore the knowledge base for more information about using Engine, how-to articles, best practices or troubleshooting guides.
- **Add-ons:** collection of utility scripts and add-ons designed to extend core features.
- **Available plugins:** explore the complete list of the application plugins available in this version of Engine. Find out quickly if your server application version is supported out-of-the-box by the corresponding plugin.
- **Claim license key:** Claim license key for a deployment not managed via Engine Management Service.

Claim license key
✕

Claim license key
Accept EULA
License key

Thank you for your interest. If you have not already purchased a software subscription, please contact [Neverfail Support](#). If you have purchased a subscription, please provide your supplied credentials along with Scope telemetry XML file, copied from %programdata%\Neverfail-SCOPE\Data\24 Hour Data folder on the active Engine node. Only telemetry data for Engine deployments with valid license subscription is accepted for upload.

Proxy settings can be configured if a direct internet connection is not available to Neverfail CE Management Service. If you have no internet connection, please contact Neverfail Support for offline licensing assistance.

Select file

Customer ID

License activation key

- **Upload telemetry data:** Upload telemetry data for a deployment not managed via Engine Management Service.

Upload telemetry data
✕

Upload file

For successful telemetry data upload, please provide the correct values, in addition to your email address:

- Scope telemetry XML file, copied from %programdata%\Neverfail-SCOPE\Data\24 Hour Data folder on the active Engine node. Only telemetry data for Engine deployments with valid license subscription is accepted for upload. If you have not already purchased a software subscription, please contact [Neverfail Support](#).
- A valid email address where you may be contacted for more details.

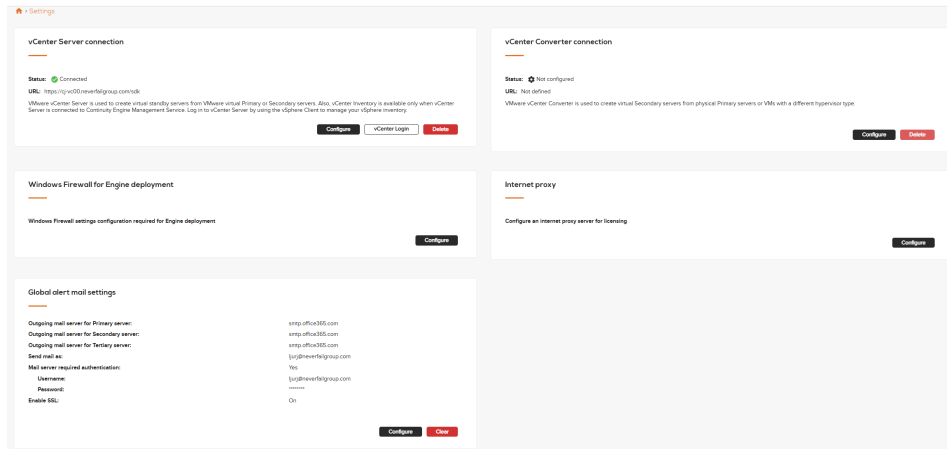
Proxy settings can be configured if a direct internet connection is not available to CE Management Service. If you have no internet connection, please contact Neverfail Support for offline licensing assistance.

Email

Select file

Settings

The **Settings** page of the EMS provides access to the underlying configuration of the Engine Management Service and other software that integrates with it.



vCenter Server connection

The VMware vCenter Server is used to manage the vSphere environments. The connection to vCenter Server enables EMS to connect to your VMware Inventory and to orchestrate the automated server protection process: Engine deployment, virtual standby instances creation (cloning) or recloning. vSphere Client is used to authenticate to vCenter Server.

The connection status and URL are displayed in this subsection, along with the options to configure the connection, log in vCenter or delete the connection.

vCenter Converter connection

The connection to vCenter Converter enables EMS to automatically create (clone-convert) or re-clone the virtual Secondary standby instances from physical Primary servers or from Primary server VMs with a different hypervisor type than ESXi.

The connection status and URL are displayed in this subsection, along with the options to configure the connection or delete the connection.

Windows Firewall for Engine deployment

The Windows Firewall can be configured in order to allow Engine deployment and communication on the Windows servers you want to protect.

Internet Proxy

The proxy server option is used in licensing scenarios where a proxy server is used on your network to access the internet.

Global alert mail settings

Global alert mail configuration will be applied on all the protected servers connected to the management service. If **Force apply** option is checked, this will override any alert mail settings configured locally on a protected server. Any server specific alert mail settings must be (re)configured in the protected server's Alerts section.

Compliance

The **Compliance** page of the EMS provides access to compliance-related artifacts for the Engine product.

SBOM - Software Bill of Materials (CycloneDX)

A Software Bill of Materials (SBOM) is a complete, formally structured list of the components, libraries, and modules that make up a software product. The EMS uses the CycloneDX format, an industry-standard specification for SBOMs.

This subsection lists the SBOMs available for download for the individual product components:

- **Engine SBOM:** the Software Bill of Materials for the Engine component.
- **Engine Management Service SBOM:** the Software Bill of Materials for the Engine Management Service component.
- **Scope SBOM:** the Software Bill of Materials for the Scope component.
- **SBOM Verification Pack:** This package contains the digital signatures and public key required to verify the integrity and authenticity of the Neverfail Engine SBOM JSON files.

Server Protection

The Server Protection chapter provides instructions on how to run the most common flows designed to protect your servers.

- **Protect a New Server**
- **Add a High Availability (HA) Instance**
- **Add a Disaster Recovery (DR) Instance**
- **Add Both HA and DR Instances**
- **Reconstruct or Reconfigure the Protected Cluster**
- **Reclone Secondary or Tertiary Instances**
- **Upgrade Applications**
- **Upgrade Engine on Protected Cluster**
- **Uninstall Engine from Protected Servers**
- **Add Already Protected Server to EMS**
- **Control the Protected Server or Cluster State**
- **Create VMware SRM Plan**
- **Manage Server Monitoring**
- **Manage Server Shadows**
- **License Server**

Protect a New Server

The **Server Protection** is initiated by starting the **Install Engine** flow, using the homonym button available in the actions menu throughout the EMS.

The **Install Engine** flow guides you through the process of configuring a protected, single instance server, taking care of the Engine installation and network configuration. Once the flow is complete, your server is protected by Engine as a single instance, without having any secondary instances like High Availability (HA) or Disaster Recovery (DR). To further enhance the server's protection, you can transform it in a protected cluster by adding standby instances (HA and/or DR). Check out the **Add a High Availability (HA) Instance**, **Add a Disaster Recovery (DR) Instance** and **Add both HA and DR Instances** articles under the Server Protection section of this documentation.

The **Install Engine** action button, an entry point in the server protection flow, is available in either one of these EMS sections:

- Server protection menu, in the top part of the **Dashboard** and **Servers** sections.
- Inventory section, next to every available virtual or physical server in your **vCenter Inventory** panel.

The same server protection flow is started regardless of the EMS section in which the flow was initiated. The only difference is that when initiating the server protection flow from the Inventory section, the Target Server entry is pre-populated with the DNS name of that server.

Server protection flow

Once the server protection is initiated by clicking the **Install Engine** or **Install** buttons, the **Install Engine** dialog is displayed.

Install Engine

Select a target server Validating install Select public IP address Ready to complete

Installing Engine

Enter DNS name or IP address Select from inventory

Target server

dns.name / 192.168.1.1

Username (with full administrator permissions)

administrator

Password

Next

The next steps will guide you through the server protection flow:

1. Start with selecting a **target server** on which the Engine will be installed. This server will become the Primary instance of your protected cluster. There are two methods of selecting the target server; you can choose either one by clicking the corresponding tab:
 - **Enter DNS name or IP address** of the server you want to protect; when you already know the DNS name or IP address of your server, manually enter it in the **Target server** field.
 - **Select from inventory** a server to protect. When a vCenter Server connection is configured, this will show you all the servers available in your vCenter inventory and allow you to search and select the one you want to protect. The **Target server** field is automatically populated with the selected server's DNS name.
2. Provide the administrator user name and password for the server you've selected for protection. Full administrative permissions are required for this account in order to perform the server protection underlying operations.

If User Access Control (UAC) is enabled on the target server, either use the target's built-in local Administrator account or, for Domain member target servers, the a domain user account member in the target's local Administrators group.

If UAC is not enabled, any user account that is member of the target's local Administrators group will work.

3. Click Next to proceed to the **Validating install** step. Here, EMS assesses the target server and validates the provided user account. Any issues with the target server or user account are shown upon the completion of the validation procedure.

The **Validating Result** shows errors, warnings or info about the targeted server. The output messages are displayed as a list sortable by message type. In case of errors, the Server Protection flow can not proceed until the errors are solved.

In case of warnings, the Server Protection flow can proceed only when the warning messages are acknowledged by the user.

The info messages can provide good insights on the targeted system status and may guide you towards optimizing your server.

4. Click Next to continue to the **Select public IP addresses** step. In this step you can select the IP addresses that will be used by clients to access your server. These IP addresses will be replicated on the passive high availability instance when added.

Make sure to unselect the IP addresses which are already assigned as channel addresses or dedicated management addresses.

5. Click Next to proceed to the **Ready to complete** step, in which the **Install Engine** dialog will present summary of the settings previously done. Upon clicking the **Finish** button, Engine will be deployed on the targeted server and will automatically discover applications it can protect.

The applications found available for protection will have their services set to Manual start, in order to allow individual management.

6. EMS will show the Servers page where the Engine deployment triggered by the executed flow is listed in the **Operation in progress** section. Once the deployment is complete, your now protected server will be listed in the Protected Servers panel.

Although your server is now running Engine, which can protect the server's applications, you can continue with enhancing your server's protection by adding a **local high availability instance** (as a cloned virtual machine) and/or a remotely located **disaster recovery instance**.

Add a High Availability (HA) Instance

With a new server completing the Protect Server flow, Engine is able to make use of its plugins to protect the applications running on that server. The newly protected server is now listed in the EMS Servers page, but its protection is not yet foolproof.

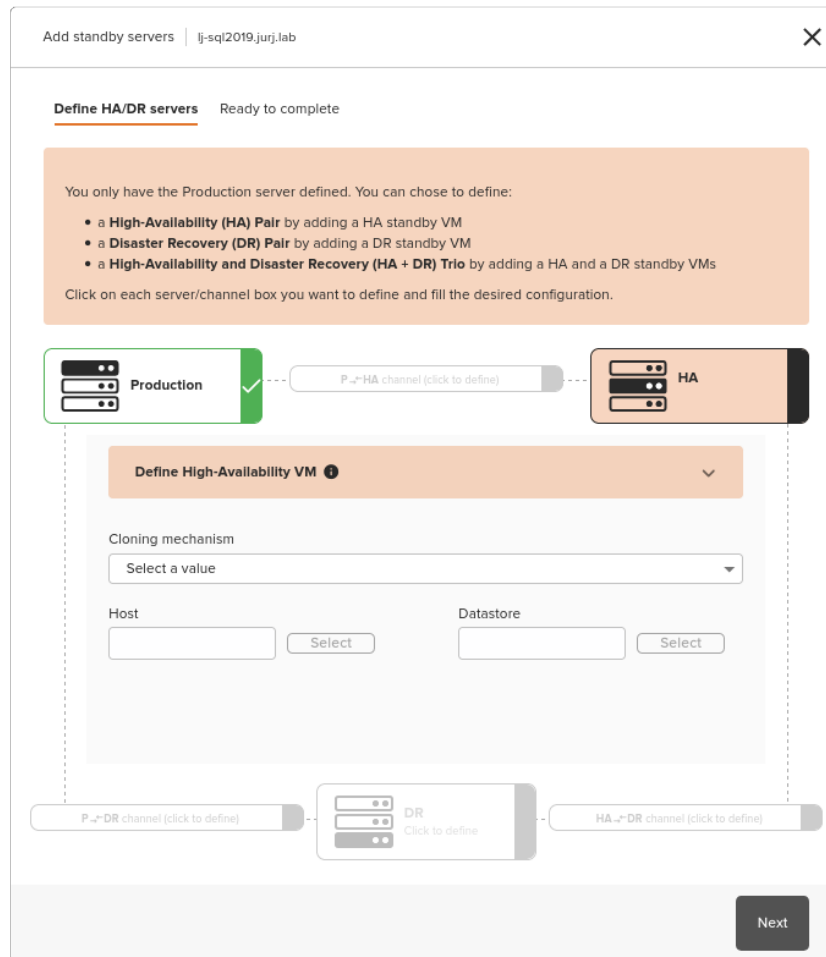
Adding a **High Availability (HA) instance** to your protected Production server enhances this protection. It allows Engine to provide a seamless transition to a perfectly synchronized standby server instance, in case your Production server is affected by any number of issues.

A HA instance is a perfect replica of your Production server, created by Engine as a virtual machine. The HA instance is added as a passive, secondary instance, which turns your production server in an active, primary instance of a protected cluster.

Adding a High Availability instance

Once your server has undergone the Server Protection flow, it is now available in the Servers page of EMS. You can access the **Server Details** page by clicking on the server's name in the Protected Servers list.

The **Status panel** will show the protected server as a single Production instance, offering the possibility to add standby servers (including HA) via the button next to the graphical representation of the server. Clicking the **+ Add standby servers** button will open the **Add standby servers** dialog.



When adding a HA instance, you will need to configure the HA instance itself and the Production-HA Channel. The next steps will guide you through the process of adding a HA instance to your protected server:

1. The **Define HA/DR servers** step of the dialog presents the available options for your current configuration in the form of clickable boxes. Select the **HA box** to configure the High Availability instance for your server. The mid section of the dialog will adjust to provide settings for the selected option (HA in this scenario).
 - Select the desired **Cloning mechanism**. You can choose between Automated or Assisted cloning:
 - **Automated** cloning relies on configured vCenter Server and vCenter Converter connections in order to seamlessly create the standby HA instance VM.
 - **Assisted** cloning is a manual cloning procedure which relies on the user creating the standby HA instance clone using a third party tool. This can be used, for example, when the HA instance is a physical machine or another type of VM.

- For Automated cloning, set the **Host** (the server hosting the clone) and the **Datastore** (clone storage) options. The Host and Datastore should be different than those of the production server, for better protection against server or storage failure.

Once the HA instance has been defined, the **HA box** is marked as defined by a check mark on green background.

2. Select the **P - HA channel box** (between your production instance and the previously defined HA instance) to configure the Production-to-High-Availability Channel. The mid section of the dialog will present the channel options; proceed as follows:
 - Select the network adapter for the channel from the combo-box. The network adapters available on the Production server will be listed as options.
 - In the Production IPv4 address field, enter the IP address to be used on the Production server side of the P - HA channel connection.
 - In the Production subnet mask, either provide the desired subnet mask or leave the field blank for using the default subnet mask.
 - In the HA IPv4 address field, enter the IP address to be used on the HA side of the P - HA channel connection.
 - In the HA subnet mask, either provide the desired subnet mask or leave the field blank for using the default subnet mask.

Once the P - HA channel has been configured, the P - HA channel box is marked as defined by a check mark on green background.

3. Click Next to proceed to the **Ready to complete** step of the procedure. The summary of the configuration is displayed under the High-Availability VM configuration section. You can review your settings and go back to the previous step to edit the configuration if needed.
4. Click Finish to start the operation. The **Operation in progress** section will list the ongoing operation and its real time progress. Once the operation is complete, the HA instance will be listed in the Status view of the Server Details section. Replication between your production instance and the new HA instance will be started automatically.

Once replication is started, the two instances of your protected cluster will reflect the Replicating status. The Primary (Production) instance synchronization should be Active while the Secondary (HA) instance should be synchronized, with a 0 seconds recovery point.

Datacenter: RO Cluj Office

Primary
Production
Active

Channel connections

192.168.60.11 ↔ 192.168.60.21

Public **192.168.169.148**

Status **Replicating**

Synchronization **Active**

Service started **Oct 08, 2020 - 16:35:26**

Network settings

Management

Datacenter: RO Cluj Office

Secondary
High Availability
Passive

Channel connections

192.168.60.21 ↔ 192.168.60.11

Public **192.168.169.148**

Status **Replicating**

Synchronization **Synchronized**
Recovery point: 0.0s

Service started **Oct 08, 2020 - 16:54:27**

Network settings

Management

Make active

Add a Disaster Recovery (DR) Instance

A **Disaster Recovery (DR)** passive server instance is a crucial step in a protected server scenario. It allows a graceful fallback to a perfectly synchronized server instance in situations where the Production server and its data center may not be operational.

Unlike a High Availability instance, a DR instance usually does not reside in the same physical location as the Production or HA instances. It can be added as a single passive instance to a Production active instance (this current article) or as a third passive instance in a protected trio cluster, along with a HA passive instance (see the **Add both HA and DR instances** article).

Like a HA instance, the DR instance is a perfect replica of your Production server, created by Engine as a virtual machine. The DR instance is added as a passive, secondary instance, which turns your production server in an active, primary instance of a protected cluster.

Adding a Disaster Recovery instance

Once your server has undergone the Server Protection flow, it is now available in the Servers page of EMS. You can access the **Server Details** page by clicking on the server's name in the Protected Servers list.

The **Status panel** will show the protected server as a single Production instance, offering the possibility to add standby servers (HA and/or DR) via the button next to the graphical representation of the server. Clicking the **+ Add standby servers** button will open the **Add standby servers** dialog.

Add standby servers | lj-sql2019.jurj.lab

Define HA/DR servers Ready to complete

You only have the Production server defined. You can chose to define:

- a **High-Availability (HA) Pair** by adding a HA standby VM
- a **Disaster Recovery (DR) Pair** by adding a DR standby VM
- a **High-Availability and Disaster Recovery (HA + DR) Trio** by adding a HA and a DR standby VMs

Click on each server/channel box you want to define and fill the desired configuration.

Production P->HA channel (click to define) HA Click to define

Define Disaster Recovery VM

Public IP address
Identical to Production

Cloning mechanism
Select a value

Host Select Datastore Select

P->DR channel (click to define) DR HA->DR channel (click to define)

Next

When adding a DR instance, you will need to configure the DR instance itself and the Production-DR Channel. The next steps will guide you through the process of adding a DR instance to your protected server:

1. The **Define HA/DR servers** step of the dialog presents the available options for your current configuration in the form of clickable boxes. Select the **DR box** to configure the Disaster Recovery instance for your server. The mid section of the dialog will adjust to provide settings for the selected option (DR in this scenario).
 - Select the **Public IP Address** for the DR instance.
 - **Identical to Production:** If your DR instance will be hosted at the same site as your production instance (not recommended), using the same subnet, the same public IP address can be used.
 - **Different than Production:** If your DR instance will run on a different site (recommended), using a different subnet than your production site, the DR instance will require a separate public IP address.

In this case, an account capable of updating the DNS servers must be specified. On switchover or failover, DNS servers will then be updated with the IP address of the active server.

- Select the desired **Cloning mechanism**. You can choose between Automated Powered-On, Automated Powered-Off or Assisted cloning:
 - **Automated Powered-On** cloning relies on configured vCenter Server and vCenter Converter connections but also on a high-bandwidth connection to the remote site where your DR instance will be hosted. The new VMware VM will be created directly on the remote host and started automatically.
 - **Automated Powered-Off** cloning uses the same vCenter Server and Converter connections to create a VMware VM on a temporary host (the same host as the production instance, for example). The resulting DR instance will be powered off after creation and ready to be transferred at its final hosting site using FTP or removable media.
 - **Assisted** cloning is a manual cloning procedure which relies on the user creating the standby DR instance clone using a third party tool. This can be used, for example, when the DR instance is a physical machine or another type of VM.
- For Automated cloning, set the **Host** (the server hosting the clone) and the **Datastore** (clone storage) options. The Host and Datastore should be different than those of the production server, for better protection against server or storage failure.

Once the DR instance has been defined, the **DR box** is marked as defined by a check mark on green background.

2. Select the **P - DR channel box** (between your production instance and the previously defined DR instance) to configure the Production-to-Disaster-Recovery Channel. The mid section of the dialog will present the channel options; proceed as follows:
 - Select the network adapter for the channel from the combo-box. The network adapters available on the Production server will be listed as options.
 - In the Production IPv4 address field, enter the IP address to be used on the Production server side of the P - DR channel connection.
 - In the Production subnet mask, either provide the desired subnet mask or leave the field blank for using the default subnet mask.
 - In the DR IPv4 address field, enter the IP address to be used on the DR side of the P - DR channel connection.

- In the DR subnet mask, either provide the desired subnet mask or leave the field blank for using the default subnet mask.

Once the P - DR channel has been configured, the P - DR channel box is marked as defined by a check mark on green background.

3. Click Next to proceed to the **Ready to complete** step of the procedure. The summary of the configuration is displayed under the Disaster Recovery VM configuration section. You can review your settings and go back to the previous step to edit the configuration if needed.
4. Click Finish to start the operation. The **Operation in progress** section will list the ongoing operation and its real time progress. Once the operation is complete, the DR instance will be listed in the Status view of the Server Details section. Replication between your production instance and the new HA instance will be started automatically.

Once replication is started, the two instances of your protected cluster will reflect the Replicating status. The Primary (Production) instance synchronization should be Active while the Secondary (DR) instance should be synchronized, with a 0 seconds recovery point.

The image displays two side-by-side screenshots of the Neverfail Engine interface, both for a 'Datacenter: RO Cluj Office'.

Left Screenshot (Primary Production):

- Instance:** Primary Production Active (Status: Active, indicated by a green checkmark).
- Channel connections:** 192.168.60.12 ↔ 192.168.60.32.
- Public IP:** 192.168.169.148.
- Status:** Replicating.
- Synchronization:** Active.
- Service started:** Oct 08, 2020 - 17:05:46.
- Network settings:** Disabled (indicated by a grey circle).
- Management:** Disabled (indicated by a grey circle).

Right Screenshot (Secondary Disaster Recovery):

- Instance:** Secondary Disaster Recovery Passive (Status: Passive, indicated by a green checkmark).
- Channel connections:** 192.168.60.32 ↔ 192.168.60.12.
- Public IP:** 192.168.169.148.
- Status:** Replicating.
- Synchronization:** Synchronized (Recovery point: 0.0s).
- Service started:** Oct 08, 2020 - 17:15:32.
- Network settings:** Disabled (indicated by a grey circle).
- Management:** Disabled (indicated by a grey circle).
- Action:** A 'Make active' button is visible at the bottom right.

Add both HA and DR instances

In the previous two articles we've described how to enhance your server's protection by individually adding a High Availability instance or a Disaster Recovery instance. Since configuring a trio protection cluster, with both HA and DR, is a common practice, we'll cover here how to add both HA and DR instances to a protected server in a single step.

Adding both HA and DR instances to a protected server

Since your server has undergone the Server Protection flow, it is available in the Servers page of EMS. You can access the **Server Details** page by clicking on the server's name in the Protected Servers list.

The **Status panel** shows the protected server as a single Production instance. The flow of adding both HA and DR standby instance is similar with the previously demonstrated flows. Clicking the **+ Add standby servers** button will open the **Add standby servers** dialog.

Add standby servers | lj=sq[2019,jur].lab
✕

Define HA/DR servers Ready to complete

You only have the Production server defined. You can chose to define:

- a **High-Availability (HA) Pair** by adding a HA standby VM
- a **Disaster Recovery (DR) Pair** by adding a DR standby VM
- a **High-Availability and Disaster Recovery (HA + DR) Trio** by adding a HA and a DR standby VMs

Click on each server/channel box you want to define and fill the desired configuration.

Production
✓

P-→HA (192.168.60.11 - 192.168.60.21)
✓

HA
Automated cloning
✓

Define HA-DR channel ⓘ

The addresses will be automatically added to each server to allow Engine to communicate and replicate data. A persistent static route should be configured for the channel connection where routing is required.

Select a network adapter for the channel

channel

| | |
|---------------------|------------------------------------|
| HA IPv4 address | HA subnet mask (blank for default) |
| 192 . 168 . 60 . 23 | . . . |
| DR IPv4 address | DR subnet mask (blank for default) |
| 192 . 168 . 60 . 33 | . . . |

P-→DR (192.168.60.12 - 192.168.60.32)
✓

DR
Automated cloning
✓

HA-→DR (192.168.60.23 - 192.168.60.33)
✓

HA and DR configurations complete. You can proceed to the next step. ▶

Next

When adding both HA and DR instances for a trio configuration, you will need to configure both HA and DR instances individually, as well as the channel connections between all three instances.

1. The **Define HA/DR servers** step of the dialog presents the available options for your current configuration in the form of clickable boxes. Select the **HA box** to configure the High Availability instance for your server. The mid section of the dialog will adjust to provide settings for the selected option (HA in this scenario).
 - Set the **Cloning mechanism** to **Automated**. Automated cloning relies on configured vCenter Server and vCenter Converter connections in order to seamlessly create the standby HA instance VM. This particular cloning mechanism is required in order to deploy both HA and DR instances in a single step.

Assisted cloning (manual cloning procedure) implies additional steps (for cloning and/or moving the instances to their host), thus defying the purpose of a unified HA and DR automatic streamlined deployment.

- Configure the **Host** (the server hosting the clone) and the **Datastore** (clone storage) options. The Host and Datastore should be different than those of the production server, for better protection against server or storage failure.

Once the HA instance has been defined, the **HA box** is marked as defined by a check mark on green background.

2. Select the **P - HA channel box** (between your production instance and the previously defined HA instance) to configure the Production-to-High-Availability Channel. The mid section of the dialog will present the channel options; proceed as follows:

- Select the network adapter for the channel from the combo-box. The network adapters available on the Production server will be listed as options.
- In the Production IPv4 address field, enter the IP address to be used on the Production server side of the P - HA channel connection.
- In the Production subnet mask, either provide the desired subnet mask or leave the field blank for using the default subnet mask.
- In the HA IPv4 address field, enter the IP address to be used on the HA side of the P - HA channel connection.
- In the HA subnet mask, either provide the desired subnet mask or leave the field blank for using the default subnet mask.

Once the P - HA channel has been configured, the P - HA channel box is marked as defined by a check mark on green background.

3. Select the **DR box** to configure the Disaster Recovery instance for your server. The mid section of the dialog will adjust to provide settings for the selected option (DR in this scenario).

- Select the **Public IP Address** for the DR instance.
 - **Identical to Production:** If your DR instance will be hosted at the same site as your production instance (not recommended), using the same subnet, the same public IP address can be used.
 - **Different than Production:** If your DR instance will run on a different site (recommended), using a different subnet than your production site, the DR instance will require a separate public IP address.

In this case, an account capable of updating the DNS servers must be specified. On switchover or failover, DNS servers will then be updated with the IP address of the active server.

- Select the desired **Cloning mechanism**. You can choose between Automated Powered-On and Automated Powered-Off.
 - **Automated Powered-On** cloning relies on configured vCenter Server and vCenter Converter connections but also on a high-bandwidth connection to the remote site where your DR instance will be hosted. The new VMware VM will be created directly on the remote host and started automatically.
 - **Automated Powered-Off** cloning uses the same vCenter Server and Converter connections to create a VMware VM on a temporary host (the same host as the production instance, for example). The resulting DR instance will be powered off after creation and ready to be transferred at its final hosting site using FTP or removable media.

Assisted cloning (manual cloning procedure) implies additional steps (for cloning and/or moving the instances to their host), thus defying the purpose of a unified HA and DR automatic streamlined deployment.

- Configure the **Host** (the server hosting the clone) and the **Datastore** (clone storage) options. The Host and Datastore should be different than those of the production or HA instances, for better protection against server or storage failure.

Once the DR instance has been defined, the **DR box** is marked as defined by a check mark on green background.

4. Select the **P - DR channel box** (between your production instance and the previously defined DR instance) to configure the Production-to-Disaster-Recovery Channel. The mid section of the dialog will present the channel options; proceed as follows:
 - Select the network adapter for the channel from the combo-box. The network adapters available on the Production server will be listed as options.
 - In the Production IPv4 address field, enter the IP address to be used on the Production server side of the P - DR channel connection.
 - In the Production subnet mask, either provide the desired subnet mask or leave the field blank for using the default subnet mask.
 - In the DR IPv4 address field, enter the IP address to be used on the DR side of the P - DR channel connection.

- In the DR subnet mask, either provide the desired subnet mask or leave the field blank for using the default subnet mask.

Once the P - DR channel has been configured, the P - DR channel box is marked as defined by a check mark on green background.

5. Select the **HA - DR channel box** (between your HA instance and the DR instance) to configure the High-Availability-to-Disaster-Recovery Channel. The mid section of the dialog will present the channel options; proceed as follows:

- Select the network adapter for the channel from the combo-box. The network adapters available on the Production server will be listed as options.
- In the HA IPv4 address field, enter the IP address to be used on the HA server side of the HA - DR channel connection.
- In the HA subnet mask, either provide the desired subnet mask or leave the field blank for using the default subnet mask.
- In the DR IPv4 address field, enter the IP address to be used on the DR side of the HA - DR channel connection.
- In the DR subnet mask, either provide the desired subnet mask or leave the field blank for using the default subnet mask.

Once the HA - DR channel has been configured, the HA - DR channel box is marked as defined by a check mark on green background.

6. Click Next to proceed to the **Ready to complete** step of the procedure. The summary of the configuration is displayed under the High-Availability VM and Disaster Recovery VM configuration sections. You can review your settings and go back to the previous step to edit the configuration if needed.

7. Click Finish to start the operation. The **Operation in progress** section will list the ongoing operation and its real time progress. Once the operation is complete, the HA and DR instances will be listed in the Status view of the Server Details section. Replication between your production instance and the new standby instances will be started automatically.

Once replication is started, the three instances of your protected cluster will reflect the Replicating status. The Primary (Production) instance synchronization should be Active while the Secondary (HA) and Tertiary (DR) instances should be synchronized, with a 0 seconds recovery point.

The image displays three side-by-side screenshots of the Neverfail Engine interface, each representing a different node in a Datacenter: RO Cluj Office. Each screenshot shows a configuration card with a green header and a green checkmark in the top right corner.

- Primary Node:** Role: Primary, Production: Active. Channel connections: 172.16.155.21 ↔ 172.16.155.22, 172.16.155.26 ↔ 172.16.155.25. Public IP: 192.168.169.30. Status: Replicating. Synchronization: Active. Service started: Aug 21, 2020 - 15:47:23. Network settings and Management are shown with circular icons.
- Secondary Node:** Role: Secondary, High Availability: Passive. Channel connections: 172.16.155.22 ↔ 172.16.155.21, 172.16.155.23 ↔ 172.16.155.24. Public IP: 192.168.169.30. Status: Replicating. Synchronization: Synchronized, Recovery point: 0.0s. Service started: Aug 21, 2020 - 15:57:55. Network settings and Management are shown with circular icons. A "Make active" button is at the bottom.
- Tertiary Node:** Role: Tertiary, Disaster Recovery: Passive. Channel connections: 172.16.155.25 ↔ 172.16.155.26, 172.16.155.24 ↔ 172.16.155.23. Public IP: 192.168.169.30. Status: Replicating. Synchronization: Synchronized, Recovery point: 0.0s. Service started: Aug 11, 2020 - 16:46:40. Network settings and Management are shown with circular icons. A "Make active" button is at the bottom.

Reconstruct or Reconfigure the Protected Cluster

The EMS **Add Standby Servers** dialog, used for creating protected clusters, also allows you to perform maintenance and configuration flows for your protected clusters, like reconstructing a cluster with failed standby instances or reconfiguring existing clusters to different topologies.

Reconstruct a protected cluster

The cluster reconstruction flow covers the scenarios in which the active, production instance is available and running while standby instances are unavailable and need to be recreated. This scenario applies to both duo clusters with an unavailable standby instance or trio clusters, where one or both standby instances are unavailable.



1. In the **Server Details** page of the target cluster, open the **Add Standby Servers** dialog (just like when creating a new instance, as described in the previous articles).
2. Notice the failed instance is not longer highlighted with green in the graphical representation of the cluster topology. The initial deployment settings (cached) are pre-configured in the Add Standby Servers dialog for an easy immediate reconstruction. A failed instance with cached settings is highlighted in yellow.
3. Selecting the yellow highlighted instance will show the configuration screen where all cached settings are already preselected. If any of the settings are not available, you can input them manually at this stage.
4. The channel connection for the failed instance must also be reconfigured, following the same procedure as when adding a new instance. Any cached settings will be highlighted with yellow and ready to be applied.
5. Once the instance and channel configuration is ready, you can proceed with recreating the instance by clicking the **Next** button and finishing the flow.

The cluster reconstruction flow also allows the reconfiguration of the standby instances you need to reconstruct.

Reconfigure a protected cluster

The cluster reconfiguration flow is employed when you need to alter the protected cluster's topology:

- from duo configuration to trio configuration
- from High Availability duo configuration to Disaster Recovery duo configuration

Cluster reconfigurations implies the recreation of the standby instances with different deployment settings. Of course, the reconfiguration of any particular standby instances can be also done by altering the individual standby instance settings.

Duo configuration to trio configuration

When your protected cluster has a P - HA or P - DR topology and you need to extend it to a P - HA - DR trio cluster, the cluster reconfiguration flow implies using the **Add Standby Servers** dialog to extend your duo cluster with another standby instance.

For duo configurations, the **Status Panel** from the Server Details page displays the **Add standby servers** button, which opens the homonym dialog. Here you can follow the procedure described in the previous articles to define the instance you need to add to your cluster:

- for a P - HA duo, you can add a new DR instance, which will be cloned from your existing HA instance.
- for a P - DR duo, you can add a new HA instance, which will be cloned from your Production instance.

Duo with HA to duo with DR

You can employ the reconfiguration flow to change your duo cluster from one type to the other (from P - DR to P - HA or from P - HA to P - DR). Since you cannot reconfigure a healthy cluster,

the applying the reconfiguration flow in this scenario implies making your standby passive instance unavailable (power-off).

1. Once the instance you want to replace is no longer available, click the **Add standby servers** button in the **Status Panel** of the Server Detail page.
2. The instance you've just made unavailable is displayed with a yellow highlight in the **Add standby servers** dialog. Make sure to remove the pre-populated settings that EMS caches.
3. Configure the new instance type to replace the unavailable one, following the same configuration procedure as described in the Adding a HA instance or Adding a DR instance articles from above, depending on your requirements.
4. Once the new instance is configured and applied, the **Status Panel** of the Server Details page in your protected cluster will show the new duo configuration.

Reclone Secondary or Tertiary instances

Recloning the Secondary or Tertiary standby instances flow helps you recreate your passive instances from your active, Production instance of your cluster. This flow can be employed, for example, when you need to update or upgrade applications running on your protected server cluster (read more in the **Upgrade Applications** article).

While recloning a passive instance, you cannot alter the channel connection configuration. Only the storage host of the instance can be configured when recloning.

It can also provide a quick way to reconstruct your standby instances, when you don't need to alter the channel configuration. For example, you can rebuild an unavailable passive instance with the exact same configuration as before, or you could employ this flow to move a passive instance to a different storage location.

The screenshot shows a wizard window titled "Reclone Secondary or Tertiary | lj-w10" with a close button (X) in the top right corner. The wizard has four steps: "Select clone type" (underlined), "Select nodes", "Select location", and "Ready to complete". Under "Automated cloning", there are two radio button options: "Create and power-on the DR cloned VM automatically after cloning" (selected) and "Create a temporary powered-off DR cloned VM locally, so that the VMDK files can be transferred manually". Below this is an option for "Assisted cloning". A large orange box titled "DR VM clone type" contains explanatory text: "You can select to reclone the stand-by VMs immediately, either automated or assisted. Alternatively, you can opt for planning a **Scheduled auto recloning**. If you have a reliable, high-bandwidth connection to the remote site, you can choose to create a stand-by server VM directly on its host. This is recommended only if you have previously cloned VMs to the remote site with success. Alternatively, you can create the VM in a temporary location on a local host. The VM will not be powered-on. You can then transfer the VMDK files to the remote site, e.g. using detachable storage or FTP. Scheduled recloning is not available with this option." At the bottom right of the wizard are "Back" and "Next" buttons.

Note: You can find out more about the use cases in which the Engine's recloning use is recommended here: [When to Use Neverfail Patch Management Options](#)

When triggering a server reclone, certain prerequisites must be met before the procedure starts:

- the Primary node is running (active) and serving applications
- for automated recloning: the Configure Connection to VMware vCenter Server must be set up correctly in the Engine Management Software
- for automated recloning: VMware vCenter Server Converter must be configured if the Primary node is not a VMware virtual machine

When the above prerequisites are met, the cluster is in the Ready State. The Engine cluster may be complete or incomplete: any of the passive servers, Secondary or Tertiary, may be present or not.

Recloning Passive Nodes with configured Static Routes - supported scenarios:

- IPv4 static routes created using the route command.

The route command is used to view and modify the network routing tables of an IP network. For example:

```
route add 192.168.33.63 mask 255.255.255.255 192.168.33.254 IF 12 -p
```

The above command adds a persistent static route for the 192.168.33.63 destination IP address, associated with the NIC interface defined by index 12, using the 192.168.33.254 address as next gateway.

- All the single NIC deployments.
- All virtual-to-virtual (V2V) deployments, where the passive nodes recloning is done via VMware vCenter cloning method.
- All virtual-to-virtual-to-virtual (V2V2V) deployments, where the passive nodes recloning is done via VMware vCenter cloning method.
- Automated, Assisted and Scheduled Automated recloning options, considering the above conditions are met.

To start the recloning flow, open the **Reclone Secondary or Tertiary** dialog from the **Actions** menu.

1. The Select clone type section allows you to choose between **Automated cloning** and **Assisted cloning**:

- The automated cloning option is available only if the standby instance to be reclone was previously cloned in an automated way. While for HA instances, the process is straight forward, there are two automated cloning methods for DR instances:

-
- Powered-on clone, which will create the cloned DR VM and power it on.
 - Powered-off clone, which will create the cloned DR VM without powering it on, allowing the user to transfer the VM manually to a remote host.
 - The assisted cloning option allows the user to perform all the cloning and transfer operation manually.
2. The Select nodes section allows you to choose what to be cloned:
 - **Full cluster reclone** will reclone all standby instances.
 - **Partial cluster reclone** is available only in trio cluster configurations and allows you to select the standby instance to be reclone.
 3. The Select location section allows you to choose the location of the reclone instance:
 - **Current location**, while keeping the original VM.
 - **Current location**, deleting the original VM once reclone.
 - **New location**, in which case you need to specify the **Host** and **Datastore**.
 4. The Ready to complete section will summarize the reclone options. Clicking the **Finish** button will initiate the reclone process and its progress will be visible in the **Operations in progress** section.

Schedule recloning

The schedule reclone feature allows you to trigger the automatic reclone flow at a predefined time and date. This can prove very useful when performing recurring maintenance tasks on the cluster.

To use the Schedule reclone feature, an automated reclone flow must have been previously applied to the standby instances to be scheduled for reclone. Scheduling only works for automated reclone. Assisted reclone cannot be scheduled.

Auto recloning | ij-w10 ✕

Select schedule Ready to complete

A schedule cannot be created for a server that was not fully cloned via vCenter Server.

A schedule cannot be created for a server that was created via vCenter Server, but was a powered-off VM.

A scheduled automated recloning will be executed successfully only when the vCenter Server configured is the same as that of the original clone.

Clear schedule

Disable schedule

Enable schedule

Once every month on the

Twice every month on the

Every second week on

Begin recloning:

Starting at

Some time between the hours of and

Delete original Secondary VM

Delete original Tertiary VM

Next

To configure a scheduled automated reclone, open the **Auto recloning** dialog by clicking the Auto reclone **Configure** button in the Summary Status panel of the Server Details page. The dialog provides the following options:

- **Clear schedule:** deletes the current recloning schedule, if existing.
- **Disable schedule:** stops the existing scheduled reclone, while keeping the schedule configuration for a later enable.
- **Enable schedule:** enables the scheduled reclone operation using the specified settings:
 - Once every month on the specified day of month.
 - Twice every month on the specified days of the month.

Note that the execution days always have a two weeks period between them (unless the first execution is set to the 14th of the month, then the second execution will be triggered in the last day of the month, regardless of how many days the month has).

- Every second week on the specified day.
- Starting time of the recloning operation. It can either be exactly specified or the starting time can be set to an optimal time inside a specified time interval.
- Original VM deletion, which can be enabled individually for Secondary and Tertiary instances. This will remove the clone source VM once the reclone operation is successfully completed.

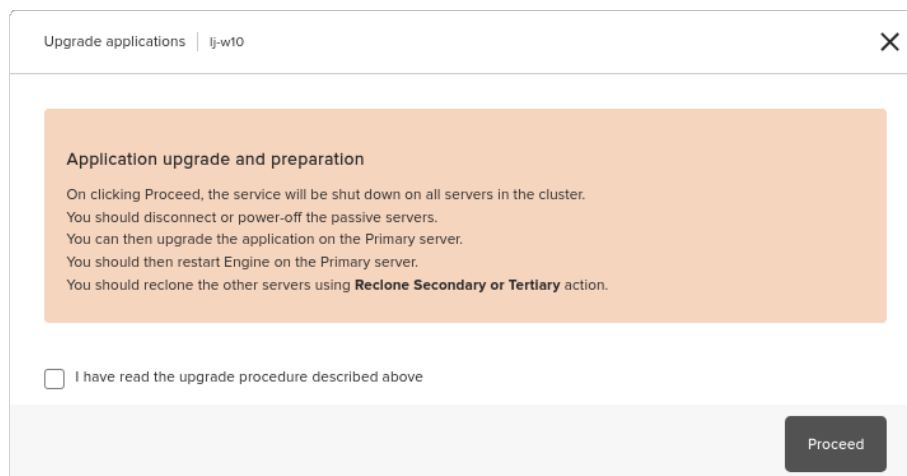
The Ready to complete section of the dialog will resume the recloning schedule. Clicking **Finish** will save and apply the defined schedule.

Upgrade Applications

The upgrade applications flow allows you to install patches and updates (including migrating to new major versions) to your applications running on the protected cluster.

The applications upgrade is done on the Primary instance, while the standby instances are shut down and discarded. Once the applications are upgraded on the Primary instance, the standby instances are recloned from the freshly updated Primary, thus reconstructing the whole cluster with upgraded applications.

To initiate the applications upgrade flow, open the Upgrade applications dialog using the **Upgrade applications** option from the **Actions** menu in the Server Details page.



The **Upgrade applications** dialog describes the steps required safely perform the upgrades and requires you to acknowledge the information in order to proceed. Upon clicking the **Proceed** button, the Neverfail Engine service is stopped on all instances of the cluster, allowing you to continue with the manual steps as follows:

- Power off or disconnect the standby instances manually, using the virtual machine management interface of choice.
At this point, the standby instances are no longer a part of your protected cluster.
- Access the Primary server and perform the required upgrades.
- Using Neverfail EMS, start the Engine service using the **Startup Engine Service** option in the server's Action menu.

- Start the **Reclone Secondary or Tertiary** flow to recreate your standby instances with all the upgrades performed on the Primary.

Once the recloning procedure has successfully completed, your applications running on the cluster will be upgraded across all instances.

Upgrade Engine on Protected Cluster

Upgrading Engine to the latest version ensures that your server's protection benefits from the features and enhancements. Whenever there is a new Engine version available for your protected servers, the **Upgrade** option will be available in the Summary Status panel of the EMS Server Details page.

Clicking the **Upgrade** button will open the **Upgrade Server** dialog, which allows you to upgrade Engine on the whole cluster or on a specific server instance.

Upgrade Server | ij-712

Provide credentials | Validating upgrade | Ready to complete

This server runs an older version of Engine. We recommend upgrading to the latest version.

Username: administrator

Password: Password

I confirm that no users are logged on the Primary, Secondary (or Tertiary) servers

Upgrade all server nodes in cluster (recommended)

Upgrade only a specific server in the cluster

Account and User Access Control (UAC)

If you have UAC enabled on the target server, you must use the built-in local Administrator account.

If UAC is not enabled, you may use any local account with membership in the local Administrators group on the target server.

All server nodes (Primary, Secondary and Tertiary if relevant) will be upgraded, unless a single node upgrade is selected.

Single node upgrades should only be used where upgrade of the whole cluster has failed, for example because of connection loss during upgrade. Single node upgrades require a unique management IP address assigned to the node.

Please see KB 2886 before using single node upgrade

Next

Proceeding with the upgrade flow requires administrative credentials for the server, which should be supplied in the **Provide credentials** step of the dialog. Also at this step, you should confirm that no users are logged in on the server(s) you wish to upgrade.

Next you can select to either upgrade the entire protected cluster or only a specific instance. You should always go with the full cluster upgrade (unless there is a specific reason for upgrading only on one instance) since this will make sure all server instances have the same CE version.

Proceeding with the upgrade flow will open the **Validating upgrade** section of the dialog, where any issues will be highlighted if arising. The upgrade procedure will not continue unless no errors or warnings are found. If any issues are found on the targeted server, you can fix them and use the **Retry** button to re-check the servers for issues.

Upon clicking **Finish**, the upgrade procedure will start and the progress will be visible in the **Operations in progress** section of EMS.

Uninstall Engine from Protected Servers

Uninstalling Neverfail Engine from your protected server will remove all Neverfail software as well as disable the provided protection on any of your server instances.

If vCenter Server connection is configured in EMS, the server VM instances can also be removed upon uninstalling Engine. This helps avoiding IP address and name conflicts.

The **Uninstall Engine** button is available in the EMS Server Details Action menu. The homonym dialog will guide you through the CE removal flow.

Uninstall Engine | lj-712

Uninstalling Engine

Local administrator account: administrator

Password: *****

I confirm that no users are logged onto the Primary, Secondary (or Tertiary) Servers

Secondary

Delete VM (recommended, requires vCenter)

Shutdown VM

Uninstall

To start the removal flow, provide the administrative credentials for the protected server. Also at this step, you should confirm that no users are logged in on the server instances.

Once the credentials are provided, your server's standby instances will be listed, each having the option to be removed from vCenter Server upon removal of Engine (if the vCenter Server connection is defined and working) or to be renamed and their NICs disabled once Engine is removed from all instances. The latter option allows you to manually manage your VM afterwards.

Pressing the **Uninstall** button will begin the removal flow, according to your configuration. The progress will be visible in the **Operations in progress** section of EMS.

Add Already Protected Server to EMS

Neverfail EMS helps you manage your Engine protected servers, but running Engine on a server and managing it through EMS are separate actions. When protecting a new server via EMS, Engine is installed on that server and, once done, the new protected server (running the Engine) is added to EMS for easy management. But since these two operations are basically independent, there might be scenarios when Engine is installed on a server outside of EMS (for example, Engine is manually installed or the server was previously removed from EMS). This means that although you have a Engine protected server, it is not managed through EMS.

These following two functionalities cover the above scenarios, where you have an already protected server and you want to add it to EMS for easy protection management.

Adding a Known Protected Server to EMS

Neverfail EMS helps you in adding an already protected server, with known hostname or IP addresses, to EMS, using the **Add Protected Server** dialog.

This dialog can be accessed either from the **Dashboard** or **Protected Servers** pages of EMS, using the dedicated **Add Protected Server** button in the top right side of the pages.

Add protected server

Add a protected server to be managed by entering the hostname (or public IP address) and port number

Hostname / Public IP address: 192.168.1.1

Port number: 9727

Enter the credentials for connecting to the server

Domain accounts should use the syntax username@domain

Username: administrator

Password: *****

Add

This simple procedure allows you to add an already protected server by providing the following server details:

- **Hostname/Public IP Address:** the host name or public IP address of the server that is protected by Engine but not managed by EMS.

- **Port Number:** the port used by Engine web services on the protected server.
- **Username:** the user name of the account with administrative privileges on the protected server. For domain accounts, the user name should follow the username@domain format.
- **Password:** the password of the user with administrative privileges on the protected server.

Once the above details are provided, clicking the **Add** button will start the procedure of finding and adding the protected server to EMS. Once complete, the protected server will be listed in the **Protected Servers** page and will be available for management.

Discovering a Protected Server and Adding it to EMS

Adding an already protected server is straightforward when you know the IP address of that server. But in many possible scenarios, the IP address is unknown to the user. EMS covers this scenarios with the **Discover Protected Servers** procedure, available in the homonym dialog.

This dialog can be accessed either from the **Dashboard** or **Protected Servers** pages of EMS, using the dedicated **Discover Protected Server** button in the top right side of the pages.

Discover protected servers

Discover protected servers from a range of IP addresses

Start IP address: 192 . 168 . 1 . 1

End IP address: 192 . 168 . 1 . 254

Port number: 9727

Enter the credentials for connecting to the servers

Domain accounts should use the syntax username@domain; depending on the DC configuration the domain may need to be the NETBIOS domain

Username: administrator

Password: *****

Search & add Close

Similar with the Add Protected Server, this procedure requires some details in order to start the discovery. Most important, the IP address range to scan:

- **Start IP Address:** the IP address defining the start of the IP address range to be scanned.
- **End IP Address:** the IP address defining the end of the IP address range to be scanned.
- **Port Number:** the port used by Engine web services on the protected server.
- **Username:** the user name of the account with administrative privileges on the protected server. For domain accounts, the user name should follow the username@domain format.

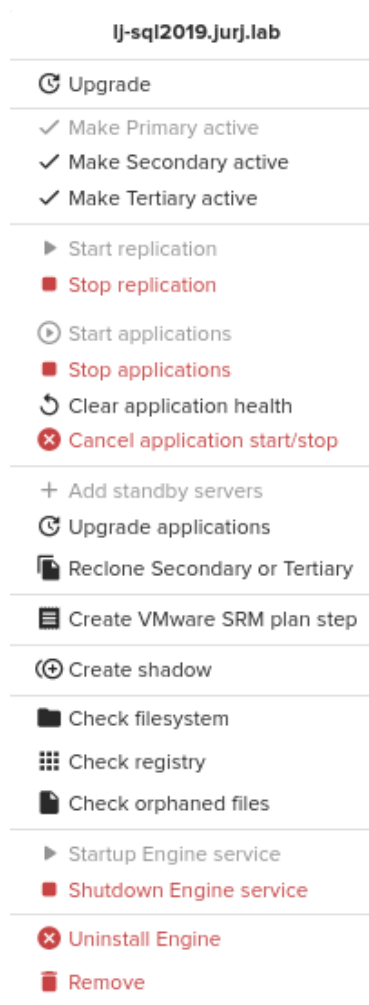
- **Password:** the password of the user with administrative privileges on the protected server.

Once the above details are provided, clicking the **Search & Add** button will begin the discovery process, scanning all the IP addresses in the specified range, and adding the discovered protected servers to EMS. Once the procedure is completed, the discovered protected servers are listed in the **Protected Servers** page and available for management.

Control the Protected Server or Cluster State

Neverfail EMS empowers you to take control of your protected servers and manually manage their states and protection.

The **Actions** menu, available on the **Servers** and **Server Details** pages, provides access to control actions for the selected server (when in the Servers page list) or for the current server (when in a protected server's Server Details page).



Using the options available in the **Actions** menu, you can perform the following control operations:

- Make instance active
- Start/stop replication
- Start/stop applications

-
- Clear applications health
 - Cancel applications start/stop procedure
 - Upgrade applications
 - Reclone Secondary or Tertiary instances
 - Start up Engine service
 - Shut down Engine service
 - Remove server from EMS

Make Instance Active

The **Make Primary active**, **Make Secondary active** and **Make Tertiary active** commands allow you to manually select which server instance in your cluster is the active instance.

The option is only available for the existing passive instances.

Upon selecting this option for a passive server, Engine will initiate a switchover operation between the currently active instance and the passive instance for which the option was triggered.

The initially passive instance will become active thus becoming the source of replication.

Start and Stop Replication

The **Start replication** and **Stop replication** commands are the manual method of controlling the replication state in your protected cluster. You can either stop the replication when it is active or start the replication when it is stopped.

Start and Stop Applications

The **Start applications** and **Stop applications** commands provide the option to control the running state of protected applications on the selected server instance. You can either start all protected applications services on an instance where the services are not running, or stop all protected applications services on an instance where the services are running.

Clear Applications Health

The **Clear applications health** option allows you to dismiss the logged health status for protected applications and reset it to healthy. The protected applications current health status is displayed in the **Applications and Platforms** panel.

Cancel Applications Start or Stop Procedure

The **Cancel applications start/stop** option allows you to gracefully interrupt the starting or stopping procedure of protected application services. EMS checks every 10 seconds for applications start/stop cancellation commands. If an application start/stop cancellation has been requested, EMS behaves as if the start/stop operation has immediately timed out.

Canceling an application start operation will put the application in the **Unmanaged - Unmonitored** state.

Canceling an application stop operation will have additional behavior compared to canceling an application start operation. When an application stop is canceled and times out, then the plan is considered to have failed, and a recovery plan will be generated. The recovery plan will attempt to start applications again.

Reclone Secondary or Tertiary Instances

The **Reclone Secondary or Tertiary** command initiates the reclone procedure, as described in the **Reclone Secondary or Tertiary Instances** article in this same documentation section.

Startup Engine Service

The **Startup Engine service** command allows you to start the Engine service on the selected instances of the current server.

Shut Down Engine Service

The **Shutdown Engine service** command allows you to manually stop the Engine service on the instances selected in the Shutdown dialog. This will stop the Engine protection and its operations (like recloning).

Remove Server from EMS

The **Remove** option allows you to remove the current protected server from EMS, without uninstalling Engine. A removed server will no longer be available in EMS but can be added at any time using the **Add an already protected server** procedure.

Create VMware SRM Plan

Before you begin

The Neverfail Engine Management Service must be installed on vCenter Server in the Recovery and Protected Sites.

Microsoft PowerShell 2.0 must be installed on all SRM servers that will run command files, for example the SRM Servers in the Recovery and Protected sites.

The PowerShell Execution Policy must be set to RemoteSigned on all SRM Servers, use the following PowerShell command:

```
PS C:\> Set-ExecutionPolicy RemoteSigned
```

This feature works to extend the capabilities of VMware's Site Recovery Manager (SRM). While SRM provides the ability to failover virtual servers to a secondary site, this feature integrates Engine physical or virtual servers into the failover process as a natural step in the SRM Site Recovery Plan executed by SRM. It works by allowing the administrator to create an SRM Step that can be added to the SRM Site Recovery Plan thereby allowing servers protected by Engine to participate in failover of servers protected by Site Recovery Manager.

Procedure

1. In the EMS Actions menu of your selected protected server, click the **Create VMware SRM plan step** option.

Important: If the server is a member of a cluster, then select the server from the cluster which is to switchover first. All members of a cluster will switchover when a single member server receives the switchover command.

The **Create VMware SRM Plan Step** dialog will open.

Create VMware SRM plan step | lj-sql2019.jurj.lab
✕

Create a script to initiate a switch-over of lj-sql2019.jurj.lab as part of an SRM recovery plan

Requires Powershell V2 on the SRM server and permission for powershell scripts to run locally without signing. For servers which are members of Business Application Groups, all members of a group will failover or switchover together. It is recommended to add only the 'First to switch' server of a group to the SRM plan.

✔ Authentication token generated for switch-over of lj-sql2019.jurj.lab

1) Choose which server the script will make active. This depends on which server is located on the site for which you are creating a plan. In order to make the server active on either site, you will require two scripts - one for each option.

Make Primary server active
 Make Secondary (or Tertiary) server active

2) If you want the plan to wait for the server to become active, enter the number of seconds. Otherwise, enter 0.

Maximum time to wait

3) Enter alternate IP addresses by which the SRM server can reach the server when passive. Multiples are separated by commas.

Alternate IP address

4) If you want to log script output to a file on the SRM server, enter the path here otherwise leave blank. Recommended for SRM 5.0

Log file for command

5) The script should be saved and copied to the SRM server on the same site as the server being made active. For SRM 5.0, the scripts must have identical names and locations on each SRM server. Use the **save button to save it as a batch file.**

6) Paste this command into the recovery plan in the SRM client, ensuring it matches where you have placed the script on the SRM server.

```
c:\windows\system32\cmd.exe /c c:\nf_make_active_lj-sql2019.jurj.lab.bat
```

2. Select the server to be controlled by the SRM Plan. This depends on which server is located at the site for which you are creating a plan.

To make the server active on either site, you will require two scripts - one for each option. If the SRM Plan Step is being created on the site where the Primary server is located, select Make Primary Server Active.

If the SRM Plan Step is being created on the site where the Secondary server is located, select Make Secondary server active.

3. If you want the SRM plan to wait for the Engine server to switchover and become active before the plan continues with the next step, enter the number of seconds to wait in the **Maximum time to wait** field.

If the Maximum time to wait is set to zero, execution of the SRM Plan will continue without waiting for the Neverfail Engine server to become active.

4. Alternate IP addresses are configured on each server in the cluster so that SRM can switch the servers even when the Protected Site cannot be contacted, for example in times of disaster.

Enter the Alternate IP address that will be used by SRM to contact the Neverfail Engine server in the **Alternate IP address** field, separate multiple IP addresses with a comma. These IP addresses are typically added to the servers as Management IP Addresses.

5. If you want to log the script output to a file on the SRM server, enter a path in the **Log file for command** field (recommended for SRM 5.0), otherwise, leave the field blank.
6. Generate two scripts using the SRMXtender Plug-in.
 - Generate one script with Make Primary Server Active selected.
 - Generate one script with Make Secondary Server Active selected.
7. The scripts should be saved as *.bat files with each being saved to a file share on the SRM server in the same site as the server being made active. Click the **Save** button to save the script as a .bat file.

For SRM 5.0, the scripts must have identical names and locations on each SRM server.

8. Launch the **VMware vSphere Web Client** and connect to the Recovery vCenter Server.
9. Navigate to **Home > Solutions and Applications > Site Recovery Manager** and select the intended Recovery Plan.
10. Select the **Recovery Steps** tab.
11. Add a **Command Step** at the desired point in the Recovery Plan. For example, if the applications running on these servers depend upon the physical server, add a Command Step before the Recover High Priority Machines Step.
12. In the **Add Command Step** dialog enter:
C:\WINDOWS\system32\cmd.exe /c [path_to_saved_file][file_name].bat
Where [path_to_saved_file] is the path where you have saved the [file_name].bat file.
13. Click **OK**.

Repeat the step creation process for each Engine pair that is to participate in the Site Recovery Plan.

Manage Server Monitoring

The **Monitoring** view of the **Server Details** page provides an overview of the server, network and applications monitoring features of EMS. This article covers the configuration options available for monitoring you protected servers.

Manage Server Monitoring

The server monitoring options can be controlled using the **Configure Server Monitoring** dialog. This dialog can be accessed using the **Configure** button in the **Server Monitoring panel** of the Server Details - Monitoring view. It can also be accessed using the Configure buttons available in the **Summary Status** panel of the Server Details - Status view.

It allows the configuration of Automated Failover, Data Loss Avoidance, Split-brain Avoidance, Active Server Isolation and Replication Response Times.

Configure server monitoring
✕

Automated failover ⓘ

Fallover timeout secs

Fallover from Primary server to Secondary server if channel heartbeat is lost for fallover timeout
 Fallover from Secondary server to Primary server if channel heartbeat is lost for fallover timeout

Data loss avoidance ⓘ

Prevent failover or auto switchover while not synchronized (recommended)

Split-brain avoidance ⓘ

Prevent failover if channel heartbeat is lost but Active server is still visible to other servers (recommended)

Ping routes from Primary to Secondary Add

From Secondary to Primary Add

Ping from Ping to 🗑️
 Ping from Ping to 🗑️

Ping interval
 High bandwidth secs
 Low bandwidth secs

Ping echo timeout
 High bandwidth secs
 Low bandwidth secs

Active server isolation ⓘ

Make the server passive if the Channel and Public networks are lost for configured fallover timeout

Replication response times

Time to wait following channel connection before starting replication secs
 Time to wait following channel disconnection before stopping replication secs

Save

Automated Failover

Automated Failover is a feature that enables any of the Primary or Secondary passive instances to become the (new) active instance of the cluster, if the (old) active instance experiences undesired events like power outage, software or hardware failures, protected applications failures, or when the channel or client network connectivity are lost

The Automated Failover is enabled by default for Production and HA instances, in more common P - HA and P - HA - DR configurations, or in case of a P - DR duo configuration, it can also be manually enabled between the Production and the DR instance.

Automated Failover cannot be enabled between a Production or HA instance and a DR instance in a trio configuration.

Automated Failover can be configured using the following options:

- **Failover Timeout:** the time, in seconds, to wait for the Channel heartbeat signal before considering the Primary instance down and initiating the automatic failover.
- **Failover from Primary server to Secondary server if Channel heartbeat is lost for failover timeout:** When enabled, failover will occur from Primary to Secondary when the Channel heartbeat signal is lost for more seconds than the Failover Timeout specifies. Enabled by default for Production and HA instances, can be manually enabled for Production and DR in a P - DR duo configuration.
- **Failover from Secondary server to Primary server if Channel heartbeat is lost for failover timeout:** When enabled, failover will occur from Secondary to Primary when the Channel heartbeat signal is lost for more seconds than the Failover Timeout specifies. Enabled by default for HA and Production instances, can be manually enabled for DR and Production in a P - DR duo configuration.

The protected cluster can also be configured to switchover automatically if there is a failure of the public network, or one of the application monitoring rules is triggered. By editing the application monitoring rules (please see the **Rules** view of the **Server Details** page), the automatic failover/switchover behavior can be overridden so that the user is alerted instead, and has the opportunity to check the system before deciding to make another server active.

Data Loss Avoidance

Data Loss Avoidance is the mechanism that prevents the state transition from passive to active (including Automated Failover) to occur between two instances that are not fully synchronized. Data loss avoidance is controlled by the following option:

- **Prevent failover or auto switchover while not synchronized (recommended):** enabled by default, activates the Data Loss Avoidance mechanism.

It is strongly recommended to leave this option enabled.

Split-brain Avoidance

Split-brain Avoidance is the mechanism that prevents undesired Automated Failovers to occur when the Channel connection is lost between the active and passive instances, but the active in-

stance is still visible to the passive instance through client or management networks. This mechanism ensures that no two instances can be active in the same time in a protected cluster.

Split-brain Avoidance can be configured by the user for Primary or Secondary instances, but not for Tertiary instances. When Tertiary is active, failover is not permitted so there's no risk for split-brain.

The split-brain avoidance can be enabled by activating the following option:

- **Prevent failover if Channel heartbeat is lost but active server is still visible to other servers (recommended).**

Split-brain Avoidance requires that ping routing is configured between the Primary and Secondary instances, using auxiliary or management IP addresses. The Channel network is also used for pinging by default and requires no configuration. Ping routing is required for checking the visibility of the active server instance. If the active instance responds to the ping, the passive instance will not failover even if the Channel connection is lost.

Use the following settings to configure ping routing for split-brain avoidance, once enabled:

- **Ping routes from Primary to Secondary:** allows you to define the ping origin and ping target IP addresses for Primary to Secondary pinging.
 - **Ping from** is the origin IP address from where the ping will be sent.
 - **Ping to** is the target IP address which will receive the ping.
- **From Secondary to Primary:** allows you to define the ping origin and ping target IP addresses for Secondary to Primary pinging.
 - **Ping from** is the origin IP address from where the ping will be sent.
 - **Ping to** is the target IP address which will receive the ping.
- **Ping interval:** allows you to define time intervals for sending the pings. The time intervals can be configured for both **high bandwidth** and **low bandwidth** networking scenarios. The proper value will be used depending on the network bandwidth.
- **Ping echo timeout:** allows you to define the timeout of the ping echo. If this time interval passes without a ping response, the ping will be considered failed. This timeout can also be configured for both high and low bandwidth scenarios, and the proper value will be used depending on the network bandwidth.

Active Server Isolation

Active Server Isolation is the mechanism that triggers the change of roles, from active to passive, when the active instance loses both Channel and client network connectivity. Active Server Isolation, enabled by default, can be manually controlled by activating or deactivating the following option:

- **Make the server passive if the Channel and Public networks are lost for configured failover timeout.**

Replication Response Times

The Replication Response Times allow you to control the time buffer between Channel connection state change and replication state change.

- **Time to wait following channel connection before starting replication:** the amount of time, in seconds, used as a buffer between the Channel connection event and replication start.
- **Time to wait following channel disconnection before stopping replication:** the amount of time, in seconds, used as a buffer between the Channel connection lost event and replication stop.

Manage Network Monitoring

The network monitoring options can be managed using the **Configure Network Monitoring** dialog, which can be accessed using the **Configure** button in the **Network Monitoring panel** of the Server Details - Monitoring view.

It allows the control of Auto-switchover if client network connectivity lost and the configuration of ping targets for checking the client network connectivity and for false failover avoidance.

Configure network monitoring

Client network connectivity and False fallover avoidance ⓘ

Auto-switchover if client network connectivity lost for pings

Ping targets from Primary

Target 1 Target 2 Target 3

Ping targets from Secondary

Target 1 Target 2 Target 3

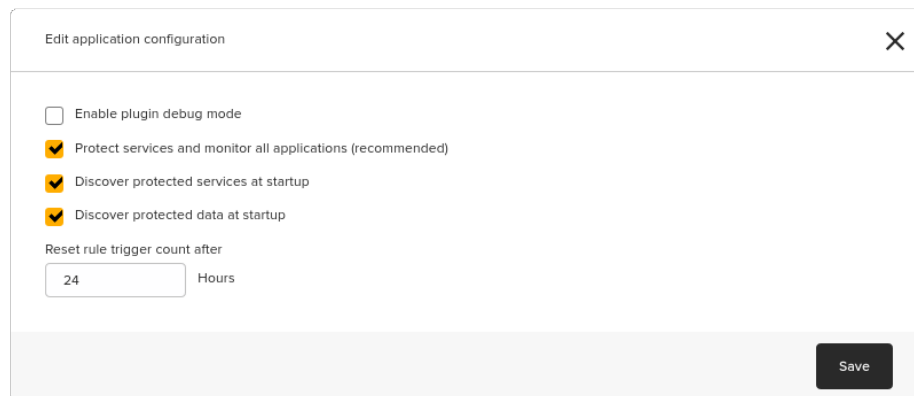
Ping interval secs Ping timeout secs

Save

- **Auto-switchover if client network connectivity lost:** if enabled, the automatic switchover is triggered in case of client network connectivity failure, after the specified number of **lost pings**.
- **Ping targets from Primary:** the network addresses to be pinged by the Primary instance in order to ensure that the connectivity with the client network is operational. By default, the IP addresses of the default gateway, the primary DNS server, and the Global Catalog server are selected.
- **Ping targets from Secondary:** the network addresses to be pinged by the Secondary instance in order to ensure that the connectivity with the client network is operational. By default, the IP addresses of the default gateway, the primary DNS server, and the Global Catalog server are selected.
- **Ping targets from Tertiary:** the network addresses to be pinged by the Tertiary instance in order to ensure that the connectivity with the client network is operational. By default, the IP addresses of the default gateway, the primary DNS server, and the Global Catalog server are selected.
- **Ping interval:** the time interval between two consecutive pings.
- **Ping timeout:** the time to wait for ping response before returning a timeout.

Manage Application Monitoring

Application monitoring options can be managed using the **Edit Application Configuration** dialog, which can be accessed using the **Configure** button in the **Application Configuration panel** of the Server Details - Monitoring view.



Edit application configuration

Enable plugin debug mode

Protect services and monitor all applications (recommended)

Discover protected services at startup

Discover protected data at startup

Reset rule trigger count after

Hours

Save

- **Enable plugin debug mode** allows you to obtain granular logs from your application plugins for debugging purposes.
- **Protect services and monitor all applications (recommended)** is enabled by default and enables Engine to protect running services and monitor applications on your server through Engine application plugins.

This option allows you to disable the protection and monitoring of applications when performing manual application maintenance tasks.

- **Discover protected services at startup** allows Engine to look for protected applications and start monitoring them every time it is started.
- **Discover protected data at startup** enables Engine to look for protected data every time it is started.
- **Reset rule trigger count after** specifies the amount of time after which the triggered (or failed) status of application-related checks and rules is reset to default. This enables application health re-evaluation.

Manage Server Shadows

The Engine Data Rollback Module (DRM) provides a way to rollback data to an earlier point in time. This helps mitigate problems associated with corrupt data such as can result from virus attacks. Before configuring or using any of the DRM features accessed through this page, Neverfail recommends that you read and follow the steps described in the section immediately below, Best Practices for Using Volume Shadow Copy Service & DRM.

Best Practices for Using Volume Shadow Copy Service & DRM

The **Volume Shadow Copy Service (VSS)** component of Windows 2008 and later takes shadow copies and allows you to configure the location and upper limit of shadow copy storage.

Note: Decide which volume to use for storing Shadow Copies before using DRM because you must delete any existing shadow copies before you can change the storage volume. Neverfail recommends that a separate volume be allocated for storing shadow copies. Do not use a volume to store both Engine protected data and unprotected, regularly updated data.

Procedure

1. To configure VSS, right-click on a volume in Windows Explorer, select **Properties**, and then select the **Shadow Copies** tab.

VSS is also used by the Shadow Copies of Shared Folders (SCSF) feature of Windows 2008R2, and consequently, some of the following recommendations are based on Microsoft™ Best Practices for SCSF. For example: do not write backups of data (even temporarily) to a volume that contains Neverfail Engine protected files, as that increases the space required for snapshots.

2. In accordance with the following guidelines from Microsoft: Select a separate volume on another disk as the storage area for shadow copies.

Select a storage area on a volume that is not shadow copied. Using a separate volume on another disk provides two advantages. First, it eliminates the possibility that high I/O load causes deletion of shadow copies. Second, this configuration provides better performance.

3. Be sure to allocate enough space for the retained shadow copies.

This is dependent on the typical load for your application, such as the number and size of emails received per day, or the number and size of transactions per day. The default is only 10% of the shadowed volume size and should be increased. Ideally, you should dedicate an entire volume on a separate disk to shadow storage.

Note: The schedule referred to in the Volume Properties > Shadow Copies > Settings dialog is for Shadow Copies for Shared Folders. This is not used for DRM - the DRM schedule is configured in the Rollback Configuration pane of the Neverfail Advanced Management Client.

4. Configure the schedule to match your clients' working patterns. Considering both the required granularity of data restoration, and the available storage. DRM provides a means of flexibly scheduling the creation of new Shadow Copies, and the deletion of older Shadow Copies. Adjust this to suit the working-patterns of your clients and applications.

For example, do clients tend to work 9am-5pm, Monday-Friday in a single time zone, or throughout the day across multiple time zones? Avoid taking Shadow Copies during an application's maintenance period, such as Exchange defragmentation, or a nightly backup.

In selecting how frequently to create new shadow copies, and how to prune older ones, you must balance the advantages of fine-granularity of restorable points-in-time versus the available disk space and the upper limit of 512 Shadow Copies across all shadowed volumes on the server.

5. Perform a trial-rollback.

After DRM is configured, Neverfail recommends that you perform a trial-rollback, to ensure that you understand how the process works, and that it works correctly. If you do not select the option Restart applications and replication, then you can rollback to Shadow Copies on the passive server without losing the most recent data on the active server.

6. Start the application manually to verify that it can start successfully using the restored data.

Note the following: - The application is stopped on the active during the period of the test. - Following the restoration of data on the passive, it becomes active and visible to clients on the network.

After the test is complete, shut down Engine on both servers. Use the Server Configuration Wizard to swap the active and passive roles, and then restart. This re-synchronizes the ap-

plication data from the active to the passive, and allows you to restart using the application data as it was immediately before the rollback.

7. Monitor Engine to identify any Shadow Copies that are discarded by VSS.

If DRM detects the deletion of any expected Shadow Copies, this is noted in the Engine Event Log. This is an indication that VSS reached its limit of available space or number of Shadow Copies. If many Shadow Copies are automatically discarded, consider adding more storage, or reconfiguring your schedule to create and maintain fewer shadow copies.

Configure Shadow Creation Options

These options set the frequency for shadow creation on the passive and active servers respectively.

Note: No shadows are created when the system status is Out-of-sync or Not Replicating.

Configure Shadows
✕

Create and maintain shadows automatically

Create a shadow every

Create a shadow on the active once per day at

Only between the hours to

Only on days to

For earlier in the current day, keep shadows only at intervals of

For earlier days in the current week, keep only the shadows nearest

For earlier weeks in the current month, keep only the shadows nearest

Shadow information location (must be in a protected location)

ⓘ Changing the Shadow information location will cause all old shadows to be deleted.

Procedure

1. Create a shadow every...

This drop-down list controls how frequently a shadow copy is captured on the passive servers, the default setting is every 30 minutes. The time when the shadow is actually captured is also controlled by **Only between the hours** and **Only on the days**. If either of these are set then shadows are captured at the frequency defined by this drop down list but only within the days/hours defined here.

2. Create a shadow on the Active once per day at...

If the check box is cleared, then no shadows are automatically created on the active. If it is selected, then a Shadow is taken each day at the time selected from the drop down list. The Shadow is taken with "application co-operation", which means that if the application protected by Engine is integrated with VSS, it is informed before the shadow is taken and given the opportunity to perform whatever tidying up it is designed to do when a VSS Shadow is taken.

Note: It is possible to select a time outside of the Only between the hours: range. This prevents creation of the shadow. Whether a shadow is actually taken is also controlled by Only between the hours: and Only on the days:, if either of these are configured, then a shadow is taken only within the days/hours defined by them. The following two options limit the number of shadows taken during periods when the data is not changing.

3. Only between the hours...

If this check box is selected, then the range defined by the two drop down lists are applied to the automatic creation of shadows on either on the passive server(s) (as controlled by Create a shadow every:), or on the active server (as controlled by Create a shadow on the Active once per day at:).

For example, to limit shadow captures to night time hours, you can define a range of 20:00 to 06:00.

4. Only on the days...

When the check box is selected, the range defined by the two drop down lists is applied to the automatic creation of shadows either on the passive server(s) (as controlled by Create a shadow every:) or active server (as controlled by Create a shadow on the Active once per day at:).

For example, to limit shadow captures to weekend days, you can define a range of Saturday to Sunday.

Note: The shadow copy information location is configurable. The default location ensures that the information location includes a copy of the necessary file filters to be used in a rollback. Neverfail recommends that the default setting be used for shadow copy information location.

Configure the Shadow Copy Schedule

DRM can create and delete shadow copies automatically according to a configurable schedule. The aim of the schedule is to provide a balance between providing a fine-granularity of rollback points-in-time on the one hand, and conserving disk space and number of shadow copies on the other. To achieve this balance, the available configuration options reflect the observation that recent events generally are of more interest and value than older ones.

For example, the default schedule maintains one shadow from every day of the last week, and one shadow from every week of the last month.

Engine can be configured to automatically create shadow copies by performing the following steps.

Procedure

1. Navigate to the **Shadows view** in the **Server Details** page and click **Configure**.

The **Configure Shadows** dialog appears.

2. Select the **Create and maintain shadows automatically** check box.

The Create and maintain shadows automatically check box controls the automatic creation and deletion of Shadow copies. When selected, automatic Shadow copies are created and deleted in accordance with other user configuration settings. When cleared, you can still manually create, delete, and rollback shadow copies from the Shadow pane.

Note: Configure the schedule to suit your clients' working patterns; the required granularity of data restoration, and the available storage.

3. Select the frequency and time periods for creating shadows. (See Configure Shadow Creation Options).
4. Select the shadows to keep or remove from earlier time periods. (See Configure Shadow Keep Options).

The Volume Shadow Copy Service (VSS) component of Windows 2008/2012, may automatically delete old shadows because of lack of disk space even when the Create and maintain shadows automatically check box is not selected.

Configure Shadow Keep Options

The purpose of the following three options is to reduce the number of older shadows while preserving a series, which spans the previous 35 days.

Manually created shadows are not deleted automatically, but VSS deletes old shadows (whether manually created or not) whenever it requires additional disk space for the creation of a new shadow. When manually created shadows match the criteria for keeping a shadow from a particular time period, automatic shadows in close proximity are deleted. For example, a manually created shadow is not deleted, but can be used for the "keep algorithm".

Procedure

1. For earlier in the current day, keep shadows only at an interval of...

If the check box is selected, then only the first shadow is kept for each interval as defined by the value (hours) selected from the drop-down list. Earlier in the current day means since Midnight and older than an hour. The intervals are calculated from either at Midnight or if Only between the hours: is selected, then from the start hour. For shadows taken before the start time (as the start time may change), the interval is calculated backwards again starting at the start time.

2. For earlier days in the current week, keep only the shadow nearest...

If the check box is selected, then only the shadow nearest to the time (24 hour clock) selected from the drop-down list is kept for each day. Earlier days in the current week means the previous seven days not including today (as today is covered by the above option). A day is defined as Midnight to Midnight.

If a shadow was taken at 5 minutes to midnight on the previous day it is not considered when calculating the nearest.

3. For earlier weeks in the current month, keep only the shadows nearest...

If the check box is selected, then only the shadow nearest to the selected day is kept for each week. Earlier weeks in the current month means the previous four weeks not including either today or the previous 7 days (as they are covered by the above two options).

To calculate the "nearest", an hour is required. The calculation attempts to use the selected time from For earlier days in the current week, keep only the shadow nearest: if it is selected, otherwise the Only between the hours start time is used if it is selected, finally, when neither of these options are configured, Midnight is used.

All automatic shadows taken more than 35 days ago are deleted. The intervening 35 days are covered by the above three options.

Manually Create Shadow Copies

Shadow Copies can be created manually using the steps below:

- In the **Shadows** view of the **Server Details** page, click **Create** button.
- Select Primary, Secondary or if present, Tertiary node.

A Shadow Copy is created on the selected node.

Delete a Shadow Copy

Should the need arise to delete shadow copies, follow the procedure below:

- To delete a shadow copy, select it in the panel in the Shadows view.
- Click **Delete**.

The selected shadow copy is deleted.

Roll Back Protected Data to a Previous Shadow Copy

Should the need arise to roll data back to a previous point in time, perform the following:

1. Go to the Shadows view of the Server Details page and select an existing Shadow from the Primary, Secondary, or Tertiary server list and click Rollback.
2. A dialog is presented allowing you to create a shadow immediately before the rollback, and select whether to restart applications and replication after the rollback.

Note: Electing to create a shadow before the rollback means that if you change your mind, you can restore to the most recent data. Choosing to restart applications and replication simplifies the restore procedure, but eliminates the chance to examine the data before it is replicated to the other server.

3. Click **OK**. A confirmation dialog is presented.

4. Click **Yes**.

Engine stops the applications and replication, and then restores protected files and the registry from the Shadow Copy. Engine then sets the file and registry filters to those persisted in the Shadow Copy. If the Shadow Copy is on a currently passive server, then this server will become active after the rollback.

If the rollback fails, the reason for the failure is shown in the status display. This may be because a particular file set of files or registry key cannot be accessed. For example, a file may be locked because the application is inadvertently running on the server performing the rollback, or permissions may prevent the SYSTEM account from updating. Rectify the problem and try performing the rollback again.

5. If selected, applications and replication are restarted and the Cluster re-synchronizes with the restored data.

- If you selected not to restart applications and replication automatically, you can now start the application manually. This allows you to check the restored data.
- If you decide to continue using the restored data, click Start on the Engine System Overview pane to re-synchronize using this data.
- If you decide you want to revert to the pre-rollback data, which is still on the other (now passive) server, you can shut down Engine, use the Configure Server Wizard to swap the active and passive roles, and then restart. This re-synchronizes the servers with the pre-rollback data.

As a result of the rollback, the file and registry filters are set to the configuration, which was in use when the shadow copy was taken.

License Server

The **License server** procedure allows you to apply a Neverfail Engine license to your protected server. This procedure is available in the homonym dialog, accessed using the **License server** option from the **Actions** menu.

The screenshot shows a dialog box titled "License server" with a close button (X) in the top right corner. Below the title bar, there are three tabs: "Licensing options" (selected), "Accept EULA", and "Apply license".

An orange informational box contains the following text: "Thank you for your interest. If you have not already purchased a license, please contact [Neverfail Support](#). If you have purchased a license, please provide your supplied credentials to license your server. **Proxy settings** can be configured if a direct internet connection is not available to Neverfail CE Management Service. If you have no internet connection, please contact Neverfail Support for offline licensing assistance."

Under "Licensing options", there are two radio buttons: "Online licensing" (selected) and "Offline licensing".

Below this, there are two input fields: "Customer ID" with the value "f2788ce7-6e75-46ba-a2d9-f062ab1ff718" and "License activation key" with the value "9664bb6e-55cc-4245-b38c-92090079de05". To the right of the license key field is a "Check activation count" button.

There is a checked checkbox for "Enable automated licensing".

Another orange informational box shows: "Activations left 1 (out of total 2) for key 9664bb6e-55cc-4245-b38c-92090079de05" and a note: "Note: If you are refreshing the license key for a renewed License Subscription, just click **Next**, ignoring the number of activations left."

At the bottom right of the dialog is a "Next" button.

Licensing a server implies that you have purchased and own a Engine license.

Online Licensing

The **Online licensing** option allows you to apply your Engine license via internet connection.

1. To apply your license, provide the required information:
 - **Customer ID**: the ID of the license customer.
 - **License authorization ID**: the authorization ID of the purchased license.
2. Once these details are provided, click **Next** to proceed with license application.

3. If EULA corresponding to your purchased license authorization was not yet accepted, the **Accept EULA** page will be presented.
Read and accept the EULA, provide your email address then click **Next** to proceed.
4. On the **Apply License** page, the license application operation progress is displayed.
Once the new license is applied, Engine will restart on your server.

Offline Licensing

The **Offline licensing** option allows you to apply a purchased Engine license when your server does not have an active internet connection.

1. To apply the license offline, provide the following information:
 - **Email address**: the email address used when purchasing the license.
 - **License key**: the license key obtained for offline licensing.
2. Once these details are provided, click **Next** to proceed with the EULA agreement.
Read and accept the EULA, provide your email address then click **Next** to proceed.
3. On the **Apply License** page, the license application operation progress and status are displayed.
Once the new license key is applied successfully, its details will be visible on **Server Summary Status** panel.

Note: If stopped, Engine service will restart on your server after the license is applied successfully.

Best Practices

The Best Practices chapter presents a set of Neverfail recommendations in order to better understand and enhance the protection of your servers.

- **Engine Role Transitions**
- **Configuring Switchover**
- **Configuring Failover**
- **Configuring Auto-Switchover**
- **Public Network Connectivity Loss**
- **Configuring Isolation**
- **Failover and Auto-Switchover**
- **HA Pair, Dedicated NICs**
- **DR Pair, Dedicated NICs**
- **HA Pair, Shared NIC**
- **DR Pair, Shared NIC**
- **HA + DR Trio**

Engine Role Transitions

Switchover

- Always user-initiated action.
- Triggered through **Make Active** action.
- Resulting configuration when executed: connected and replicating cluster with the node selected by the user acting as the new active server.

Failover

- Occurs in case of active server severe failure: site failure, hardware failure, network connectivity failure, etc...
- It is the transition which changes the role of a passive server to active, as a consequence of an unexpected event.
- Can be automated or manual/user-initiated.
- By default it is configured as automated (enabled) for HA deployments and manual/user-initiated (disabled) in DR deployments.
- By design, in Trio deployments, Engine won't failover ever automatically on the Tertiary DR node: user intervention is required for failing-over to Tertiary node.
- Resulting configuration when executed: The previous active server is either crashed or isolated; new active server in a disconnected non-replicating cluster (pair); OR new active server in a partially connected and replicating cluster (for trio).

Auto-switchover

- Could be defined as a special case of failover when some user defined conditions are met, i.e. public network loss, service failure, rule trigger.
- It is more close to a switchover, but configured to occur in an automated manner.
- Resulting configuration when executed: Engine service stopped cleanly on the former active server and moving workloads to a newly active server in an incomplete cluster.

Isolation

- Active server makes itself passive due to some conditions it detects, when it is not safer anymore to act as the active server and serve clients.
- It represents the active-to-passive transition in case of a failover scenario caused by network connectivity loss on the active server.
- Resulting configuration when executed: The active and passive servers are isolated without seeing the outside world or the other node in the cluster.

Configuring Switchover

There's nothing to configure here since this is a user-initiated action.

Note: Switchover is not permitted unless the servers are Connected, in a Replicating state and synchronized.

Configuring Failover

Engine continuously monitors the servers in the pair/trio and the network to ensure availability and uses native logic and a combination of elapsed time, administrator configured rules, current server network status, and configured ping routing to determine if failover or isolation of the active server is warranted should the servers experience missed heartbeats.

To configure failover:

- Connect to the Engine cluster using the EMS.
- Navigate to **Server Details** page > **Monitoring** tab > **Server Monitoring** panel and open the Configure Server Monitoring dialog from the **Configure** button.

- The Failover timeout can be customized by changing the default value (60 seconds) to a custom value. Type a new numeric value (seconds) in the Failover timeout text box or use

the arrow buttons to configure how long Engine waits for a missed heartbeat before it takes a pre-configured action to failover or isolate the active server from the network.

- Select or clear check boxes for the items listed below to select the actions to take if the specified Failover timeout is exceeded. When the configured Failover timeout value has elapsed, Engine will evaluate, in order, the following pre-configured options before taking action:
 1. **Option 1** Failover from Primary server to Secondary server if channel heartbeat is lost for failover timeout
 2. **Option 2** Failover from Secondary server to Primary server if channel heartbeat is lost for failover timeout
 3. **Option 3** Prevent failover or auto switchover while not synchronized (recommended)
 4. **Option 4** Prevent Failover if channel heartbeat is lost but Active server is still visible to other servers (recommended) (Configure **Option 6** "Ping routes from..." for adding routes that can be used to check the visibility of the Active server)
 5. **Option 5** Make the server passive if the Channel and Public networks are lost for the configured failover timeout (Active Server Isolation)

Option 4 requires configuration of a pinging route between passive and active (**Option 6**). Use the **Ping routes...** settings and configure pinging over the Public network using auxiliary/management IP addresses configured on each node.

The Split-brain avoidance Ping settings allows you to change the **Ping Interval** and **Ping Echo Timeout** for both high and low bandwidth instances, but the default values work just fine.

Additional considerations:

- **Option 1** and **Option 2** enable in fact all the automated transitions - failover and auto-switchover - in the indicated directions: from Primary to Secondary and from Secondary to Primary (Automated transition to Tertiary is not permitted by design, that's why this is not considered for failover configuration).
- **Option 3**, **Option 4** and **Option 5** will be evaluated only if **Option 1** and/or **Option 2** are enabled.
- If any of **Option 3**, **Option 4** and **Option 5** are not selected, Engine will skip that option and move to the next one in the list. After all selected options have been evaluated Engine will take action (This assumes at least one of **Option 1** and **Option 2** is/are enabled).

Note: Engine has an additional failover veto conditions compared to versions older than v8.1. It pings from the passive server to the configured public network targets across the Management IP addresses (configured in CSW). If no ping response is received from the public targets, then passive server won't failover. This extra fail safe is enabled when a Management IP is configured on the passive server (if no management IP is defined, the pinging to local passive site GW/DNS/targets is not happening).

Configuring Auto-Switchover

By default Auto-switchover is enabled/permitted for HA installations and disabled/not-permitted for DR installations. However, even if for scenarios where it is permitted, by default it is not configured thus requires user configuration.

Auto-switchover is enabled thus permitted as follows:

- **From Primary to Secondary:** If the failover **Option 1** is enabled - Failover from Primary server to Secondary server if channel heartbeat is lost for failover timeout.
- **From Secondary to Primary:** If the failover **Option 2** is enabled - Failover from Secondary server to Primary server if channel heartbeat is lost for failover timeout.

Note: Auto-switchover to Tertiary server is not permitted by design

Auto-switchover can be configured for the following scenarios:

- **Service failure (Option 9):** By configuring the service failure recovery action to Switchover, as indicated below. By default, all the protected services recover actions in case of failure are set to Recover Service.
- **Rule trigger (Option 10):** By configuring the failure recovery action to Switchover, as indicated below.

Note: as a best practice, the Switchover option should be used as the last failure recovery action; the first two recover actions should retry to recover the service or recheck the rule trigger condition.

Public Network Connectivity Loss

The Public network monitoring feature is enabled by default during the installation of Neverfail Engine.

This feature integrates the polling of the particular targets around the network through the active server's Principal (Public) connection to ensure connectivity with the Public network is operational.

By default, the IP addresses of the default Gateway, the primary DNS server, and the Global Catalog server are all selected as targets. It may however be the case that one or more of the automatically discovered targets are co-located on a physical machine leading to duplication of IP addresses. In such a scenario, the ability to specify additional targets manually becomes an advantage.

Configure network monitoring

Client network connectivity and False failover avoidance

Auto-switchover if client network connectivity lost for pings

Ping targets from Primary

Target 1 Target 2 Target 3

Ping targets from Secondary

Target 1 Target 2 Target 3

Ping interval secs Ping timeout secs

Save

To specify a manual target for the Public network checking, click the Network Monitoring **Configure** button to invoke the **Configure Network Monitoring** dialog. Edit the **Ping targets from... (Option 7)** tab to add to or modify the existing target IP addresses for each server to ping.

To configure the Auto-switchover in this scenario, one should tick (check) the following option: **Auto-switchover if client network connectivity lost for specified number of pings (Option 8)**. Once the failure count of all three targets has exceeded this value, Engine will initiate an auto-switchover.

Configuring Isolation

By design, the active server isolation is enabled for both HA and DR scenarios.

The enablement/disablement is dictated by the failover's configuration **Option 5**; Make the server passive if the Channel and Public networks are lost for the configured failover timeout.

Failover and Auto-Switchover

Failover and Auto-switchover resulting behavior depends on the following variables:

- The above configuration options.
- The Engine deployment topology: HA pair, DR pair, HA+DR Trio.
- Networking configuration: single/shared NIC for channel and public connections; dedicated NICs for channel and public connections.

In the next sections, we'll try to cover the various failover/auto-switchover and isolation scenarios. This will explain how failover works.

Note: All the scenarios presented below will start with the default configuration depending on the installation type. Then we'll indicate how the behavior can be changed depending on the available configuration options.

Best practices

- If possible use 2 separated channel connections. This redundancy will make the channel connectivity more reliable in case that one connection fails.
- Always configure the Configure [Server Monitoring> Ping Configuration] Ping Routing option. You can use a management/auxiliary or any IP address not used by Engine. This IP can sit either in the Public subnet or in any other dedicated subnet different from the channel. This will assure an extra layer of protection against potential false failovers and split-brain syndrome. And this configuration applies to both single and multi NIC topologies.

HA Pair, Dedicated NICs

Default post-install configured options

- **Option 1** [Configure Server Monitoring] Failover from Primary server to Secondary server if channel heartbeat is lost for failover timeout: **Enabled**
- **Option 2** [Configure Server Monitoring] Failover from Secondary server to Primary server if channel heartbeat is lost for failover timeout: **Enabled**
- **Option 3** [Configure Server Monitoring] Prevent failover or auto switchover while not synchronized (recommended): **Enabled**
- **Option 4** [Configure Server Monitoring] Prevent Failover if channel heartbeat is lost but Active server is still visible to other servers (recommended): **Enabled**

Please use Ping configuration to set a route that can be used to check the visibility of the Active server.

- **Option 5** [Configure Server Monitoring] Make the server passive if the Channel and Public networks are lost for the configured failover timeout: **Enabled**
- **Option 6** [Configure Server Monitoring] Ping Routes: **Not Configured**
- **Option 7** [Configure Network Monitoring] Ping targets: **Configured**
- **Option 8** [Configure Network Monitoring] Auto-switchover if client network connectivity lost for specified number of pings: **Disabled**
- **Option 9** [Service failure recovery action] Switchover: **Not configured**
- **Option 10** [Rule trigger action] Switchover: **Not configured**

Scenario 1 - Failover when the active server has failed and is no longer available or visible on any network

Behavior explanation

Upon detection of missed heartbeats, Engine on the passive server performs the following steps:

1. As soon as the passive server detects that the Neverfail Channel is experiencing missed heartbeats, it will determine if itself is a valid failover target to the currently active server.

2. As soon as the passive server detects that the Neverfail Channel is experiencing missed heartbeats, it will attempt to ping the active server's Management IP address via the Public network using the passive server's NIC configured with the Management IP address. If the ping is successful, the passive server will veto the failover. If the ping is unsuccessful, it will continue to the next step.

NOTE: Since the passive server assumes that active server has failed, the passive server will not attempt to verify synchronization with the active server.

3. At this point, the passive server checks the configured value of the Failover timeout and starts a "Heartbeat lost" countdown. The passive server continues with the next step.
4. At this point, failover to the passive server is postponed until the value of the Failover timeout has elapsed.
5. The passive server changes its role to active, removes the packet filter, and starts all services.
6. As the new active server, it will begin accepting traffic from clients.

Scenario 2 - Active Server Isolation when the active server has lost connection with the passive server and public network

Behavior explanation

Network Isolation Workflow Diagram:



Upon detection of missed heartbeats Engine performs the following steps:

1. As soon as the active server detects that the Neverfail Channel is experiencing missed heartbeats, it will determine if a valid failover target (the passive server) is present. Simultaneously, once the passive server detects missed heartbeats, it will determine if it is a valid failover target.
2. Next, the active server will determine if it is synchronized with the failover target (the passive server). If synchronized, it will continue to the next step. If it is not synchronized, it will veto a failover. Simultaneously, the passive server checks to see if it is synchronized with the active server. If synchronized, it will continue to the next step. If it is not synchronized, it will veto a failover.
3. At this point, both the active and passive servers check the configured value of the Failover timeout and start a "Heartbeat lost" countdown. Both servers should start the countdown at approximately the same time.

4. Failover or isolation of the active server is postponed until the configured Failover timeout value (in seconds) has elapsed and it is during this period that both servers accomplish steps 1 & 2.
5. Once the configured Failover timeout period has elapsed, the active server assumes the Neverfail Channel is lost and will attempt to ping the failover target (passive server) via the Public network. If the ping is successful, active server isolation is vetoed. If the attempt to ping the failover target is unsuccessful, the active server will proceed to the next step. Simultaneously, the passive server assumes the Neverfail Channel is lost and attempts to ping the active server via the Public network. If the ping is successful, failover is vetoed. If the ping attempt is unsuccessful, the passive server proceeds to the next step.

NOTE: If the servers have reached this point, then neither server can see the other server.

6. The active server checks only its own network connectivity to the Public network. If the active server has lost connectivity to the Public network, it will isolate itself by making itself passive (potential active).
7. Both the active and passive servers will check their connectivity to the Public network. If the active server has lost connectivity to the Public network, it will isolate itself by making itself passive (potential active). Should the active server reconnect with the passive, it will become active again. Otherwise, it will remain passive. If the passive server has lost connectivity to the Public network, it will veto a failover.

Scenario 3 - Engine service fails unexpectedly on the active server

Result

Failover prevented.

Behavior explanation

- Channel disconnection event triggered due to heartbeat missed.
- Countdown is started.
- Failover is prevented because active server is still visible to passive - passive pings active across the channel using the channel IP addresses (default, built in behavior)

Scenario 4 - Channel connection is lost between active and passive

Result

Failover is not prevented, but isolation doesn't happen, leading to undesired split-brain syndrome risk.

Behavior explanation

- Channel disconnection event triggered due to heartbeat missed.
- Countdown is started.
- Failover happens because active server is not visible to passive.
- In the same time isolation of former active is prevented as it still sees the public network.
- Undesired result: split-brain, two active servers.

NOTE: Behavior change since version v8.5 (see above).

How to avoid this scenario: **configure Server Monitoring > Ping Routing - Option 6** from the table above. This means adding auxiliary/management IP addresses on the nodes and configure the ping routing between them. Therefore that **Option 6** rule will be effectively applied and prevent failover.

Scenario 5 - Active server loses connectivity with public network (nodes remain connected though channel). Auto-switchover if client network connectivity lost for Option 8 is not enabled

Result

Failover is not happening.

Behavior explanation

- Failover is not happening because channel is still connected. This is not quite desired because even if cluster is connected the active server cannot serve clients

How to avoid this scenario: **enable Option 8 option** from above (treated in next scenario).

Scenario 6 - Active server loses connectivity with public network (nodes remain connected though channel). Auto-switchover if client network connectivity lost for Option 8 is enabled

Result

Auto-switchover happens.

Behavior explanation

- Public network missed pings starts the countdown.
- Auto-switchover occurs because auto-switchover condition **Option 8** is enabled.

Scenario 7 - Active suffers severe failure or total network connectivity failure (channel and public) WHILE not in sync with the passive

Result

Failover / auto-switchover is prevented as effect of Option 3.

Behavior explanation

- **Option 3** - servers not in sync prevent failover.

Scenario 8 - Auto-switchover when service fails

Behavior explanation

- Auto-switchover will happen if **Option 9** is configured and the configured server fails. It is recommended to configure the recovery actions Recover/Recover/Switchover for the critical services intended to trigger failover if they cannot be recovered.

Scenario 9 - Auto-switchover when rule is triggered

Behavior explanation

- Auto-switchover will happen if **Option 10** is configured and the rule triggers.
-

DR Pair, Dedicated NICs

Default post-install configured options

- **Option 1** [Configure Server Monitoring] Failover from Primary server to Secondary server if channel heartbeat is lost for failover timeout: **Disabled**
- **Option 2** [Configure Server Monitoring] Failover from Secondary server to Primary server if channel heartbeat is lost for failover timeout: **Disabled**
- **Option 3** [Configure Server Monitoring] Prevent failover or auto switchover while not synchronized (recommended): **Enabled**
- **Option 4** [Configure Server Monitoring] Prevent Failover if channel heartbeat is lost but Active server is still visible to other servers (recommended): **Enabled**.

Please use " Configure Pings..." to configure a route that can be used to check the visibility of the Active server.

- **Option 5** [Configure Server Monitoring] Make the server passive if the Channel and Public networks are lost for the configured failover timeout: **Enabled**
- **Option 6** [Configure Server Monitoring] Ping Routing: **Not Configured**
- **Option 7** [Configure Network Monitoring] Ping targets: **Configured**
- **Option 8** [Configure Network Monitoring] Auto-switchover if client network connectivity lost for specified number of pings: **Disabled**
- **Option 9** [Service failure recovery action] Switchover: **Not configured**
- **Option 10** [Rule trigger action] Switchover: **Not configured**

By default, any automated transitions (failover, auto-switchover, isolation) are disabled in this DR scenario, i.e. **Option 1** and **Option 2** conditions. If they're enabled post-install, the cluster will behave exactly like a Pair HA, and all of the above scenarios will apply.

The next scenarios will explain the behavior when automated failover/auto-switchover are not enabled, and the Engine transitions require user intervention

Scenario 10 - Manual failover in case of active server failure in a DR pair

Behavior explanation

- Passive server stays passive.
- User can initiate the manual failover to passive server after the disconnection period elapses. This can be done from the Advanced Management Client, Make Active... action.

Scenario 11 - Auto-switchover not permitted in case of public network connectivity loss in a DR pair

Behavior explanation

- Even if condition is met (public targets are not reachable anymore), autoswitchover is not permitted when failover is not enabled; an user warning is raised in this sense.

Scenario 12 - Auto-switchover not permitted in case of service failure in a DR pair

Behavior explanation

- Even if condition is met, auto-switchover is not permitted when failover is not enabled; an user warning is raised in this sense.

Scenario 13 - Auto-switchover not permitted in case of rule triggered in a DR pair

Behavior explanation

- User is warned when a rule is configured to switchover.
- Even if condition is met, auto-switchover is not permitted when failover is not enabled; an user warning is raised in this sense.

HA Pair, Shared NIC

The same failover options are configured as for Pair HA multi NIC deployment.

- **Option 1** [Configure Server Monitoring] Failover from Primary server to Secondary server if channel heartbeat is lost for failover timeout: **Enabled**
- **Option 2** [Configure Server Monitoring] Failover from Secondary server to Primary server if channel heartbeat is lost for failover timeout: **Enabled**
- **Option 3** [Configure Server Monitoring] Prevent failover or auto switchover while not synchronized (recommended): **Enabled**
- **Option 4** [Configure Server Monitoring] Prevent Failover if channel heartbeat is lost but Active server is still visible to other servers (recommended): **Enabled**

Please use "Configure Pings..." to configure a route that can be used to check the visibility of the Active server.

- **Option 5** [Configure Server Monitoring] Make the server passive if the Channel and Public networks are lost for the configured failover timeout: **Enabled**
- **Option 6** [Configure Server Monitoring] Ping Routing: **Not Configured**
- **Option 7** [Configure Network Monitoring] Ping targets: **Configured**
- **Option 8** [Configure Network Monitoring] Auto-switchover if client network connectivity lost for specified number of pings: **Disabled**
- **Option 9** [Service failure recovery action] Switchover: **Not configured**
- **Option 10** [Rule trigger action] Switchover: **Not configured**

This can be considered a simplified use-case of a Pair HA multi NIC deployment, having both dedicated/separated channel and public NICs. Thus all the Pair HA multi NIC deployment scenarios (1 to 9) apply here too, except scenarios 4, 5, 6 which become now:

Scenario 14 - Failover in case of active losing channel and public network connectivity for Pair HA single NIC deployment

Behavior explanation

- The active will isolate as per **Option 5**.

- The passive will become active as it cannot see anymore previous active (also, in v8.5 or newer because it can see the public network).

Scenario 15 - Passive losing channel and public network connectivity in Pair HA single NIC deployment. Active still connected to public network

Behavior explanation

- The active will stay active because it can see the public network.
- In v8.1 passive will become active as it cannot see anymore previous active. Here's a potential risk of split brain if network connectivity is restored at passive end after this become active.
- In v8.5 or newer passive won't become active because it can't see the public network.

DR Pair, Shared NIC

The same failover options are configured as for Pair DR multi NIC deployment.

- **Option 1** [Configure Server Monitoring] Failover from Primary server to Secondary server if channel heartbeat is lost for failover timeout: **Disabled**
- **Option 2** [Configure Server Monitoring] Failover from Secondary server to Primary server if channel heartbeat is lost for failover timeout: **Disabled**
- **Option 3** [Configure Server Monitoring] Prevent failover or auto switchover while not synchronized (recommended): **Enabled**
- **Option 4** [Configure Server Monitoring] Prevent Failover if channel heartbeat is lost but Active server is still visible to other servers (recommended): **Enabled**

Please use Ping configuration to set up a route that can be used to check the visibility of the Active server.

- **Option 5** [Configure Server Monitoring] Make the server passive if the Channel and Public networks are lost for the configured failover timeout: **Enabled**
- **Option 6** [Configure Server Monitoring] Ping Routing: **Not Configured**
- **Option 7** [Configure Network Monitoring] Ping targets: **Configured**
- **Option 8** [Configure Network Monitoring] Auto-switchover if client network connectivity lost for specified number of pings: **Disabled**
- **Option 9** [Service failure recovery action] Switchover: **Not configured**
- **Option 10** [Rule trigger action] Switchover: **Not configured**

By default, any automated transitions (failover, auto-switchover, isolation) are disabled in this DR scenario, i.e. **Option 1** and **Option 2** conditions. If they're enabled post-install, the cluster will behave exactly like a Pair HA single NIC, and all of the above scenarios will apply, i.e. 1, 2, 3, 7, 8, 9, 14, 15.

With regards to the manual failover all of the above Pair DR multi NIC scenarios apply here too, i.e. 10, 11, 12, 13.

HA + DR Trio

The same failover options are configured as for Pair HA:

- **Option 1** [Configure Failover] Failover from Primary server to Secondary server if channel heartbeat is lost for failover timeout: **Enabled**
- **Option 2** [Configure Failover] Failover from Secondary server to Primary server if channel heartbeat is lost for failover timeout: **Enabled**
- **Option 3** [Configure Failover] Prevent failover or auto switchover while not synchronized (recommended): **Enabled**
- **Option 4** [Configure Failover] Prevent Failover if channel heartbeat is lost but Active server is still visible to other servers (recommended): **Enabled**

Please use " Configure Pings..." to configure a route that can be used to check the visibility of the Active server.

- **Option 5** [Configure Failover] Make the server passive if the Channel and Public networks are lost for the configured failover timeout: **Enabled**
- **Option 6** [Server Monitoring> Ping Configuration] Ping Routing: **Not Configured**
- **Option 7** [Network Monitoring> Ping Configuration] Ping Routing > Ping targets: **Configured**
- **Option 8** [Network Monitoring] Auto-switchover if client network connectivity lost for specified number of pings: **Disabled**
- **Option 9** [Service failure recovery action] Switchover: **Not configured**
- **Option 10** [Rule trigger action] Switchover: **Not configured**

With regards to the failover behavior, this scenario can be considered as a mix of a Pair HA between Primary and Secondary with a Pair DR between any of the HA nodes (Primary and Secondary) and the Tertiary/DR node:

- Automated failover, auto-switchover, isolation are enabled by default and permitted between the HA nodes (Primary and Secondary).
- Automated failover or switchover is never permitted on the Tertiary node by design. This means Failover to Tertiary is never done automatically and same for Failover from Tertiary

to any of the P and S. Tertiary won't ever isolate. In other words, once Tertiary is made manually active (manual failover) it stays so till its state is changed by user intervention.

- No matter Primary or Secondary is the active server, all the time Tertiary is placed at the end of the replication chain, i.e, $P > S > T$ or $S > P > T$.

Regarding the failover behavior, automated failover/switchover and isolation scenarios described above apply entirely to the transitions between Primary and Secondary nodes (we can call this the HA pair part of a trio). The automated transitions behavior is the same no matter if Tertiary is present in the cluster or absent (incomplete cluster because trio is done, powered off, etc). Thus:

- For a separated/dedicated NICs for P and S Channel and Public IP addresses deployment all of the above Pair HA multi NIC scenarios apply.
- For a shared/single NIC for P and S Channel and Public IP addresses deployment all of the above Pair HA single NIC scenarios apply.

The DR pair part of a trio can be formed by any of the HA nodes and the DR node, i.e. (P, T) or (S, T). For these use cases when one of the HA nodes is not available, the same manual transitions behavior applies as for a Pair DR:

- For a separated/dedicated NICs for P/S and T Channel and Public IP addresses deployment all of the above Pair DR multi NIC scenarios apply.
- For a shared/single NIC for P/S and T Channel and Public IP addresses deployment all of the above Pair DR single NIC scenarios apply.