




Fault Tolerance of WVMS Enterprise



Ensuring a robust and reliable video management system is crucial for effective surveillance. Wavesys WVMS Enterprise recognizes the significance of fault tolerance mechanisms and offers a range of redundancy features to minimize disruptions and maintain data integrity. These features include database backup, management server mirroring, failover, fallback storage, replication, and edge recording. By combining these capabilities, Wavesys WVMS Enterprise provides a high availability service that meets stringent redundancy requirements, such as those outlined by the Security Monitoring Standards from the Monitoring & Control Centre (MCC) of UAE.

In this article, we will delve into the failover mechanism for recording servers within the Wavesys WVMS Enterprise system. We will explore how this mechanism can be finely tuned to meet the highest standards of redundancy, ensuring uninterrupted video recordings and fulfilling the stringent demands of security monitoring.

What is failover?


The failover feature enables quick automated recovery of the recording system in case of a sudden failure. WVMS Enterprise server carefully tracks the operation of all the servers in the system by exchanging heartbeat messages with each one; if there is no response for too long, the central management server assumes that the target recording server is down and immediately replaces that server with a spare one.

What does actually happen when a faulty server is replaced? The central management server remembers the configuration of all the servers in the system; so, a secondary (spare) server is assigned a configuration that is identical to the one from the primary recording server. Thus, the fault has no consequences on everyone else in the system as all the live and recorded video streams are continuously available. When the malfunctioned machine comes back online, it is returned to operation either automatically or manually by the administrator, depending on the desired system configuration.

In order to achieve this, recording servers are split into groups — clusters — based on their location, purpose, importance, configuration and/or other attributes or their combinations. Inside the cluster, there may be any proportion of actual (primary) recording servers and 'spare parts' — failover nodes: thus, any required redundancy level can be reached. There are no limitations on the number of clusters or on the number of servers in a cluster.

To ensure seamless failover in Wavesys WVMS Enterprise, certain requirements must be met for the failover servers to effectively take over any server within the same cluster:

1. Failover nodes should be located within the same environment to ensure access to all devices.
2. The hardware specifications of failover nodes should be equal to or higher than the top-spec recording server(s) within the cluster. If there is a discrepancy, it may be necessary to review and reorganize the clusters accordingly.
3. Failover servers should not have any permanent device configurations assigned to them. Their configurations should remain empty while the failover node is idle.



In Wavesys WVMS Enterprise, both primary recorders and failover servers utilize the same installation package, namely the WVMS Global Recording Server. The server roles are assigned in their respective settings. It's important to note that the Global server itself does not participate in failover clustering and therefore is not covered by failover mechanisms. While it is allowed to use the Global server for recording purposes, it is recommended to use it primarily for management functions. Achieving high availability for the central management server is accomplished through a separate redundancy feature called central server mirroring.

One time setup

Failover is configured via Wavesys WVMS Console application. Generally, the plan is:

1. Create a failover cluster
2. Add all necessary recording servers to the WVMS Global configuration — manually or using server auto discovery
3. Assign server roles (primary recorder or failover node) and define their settings
4. Put the servers into the cluster(s)

The order of these steps is not crucial, e.g., you can first add the servers and create the failover cluster later.

We shall consider a system with one cluster that contains two primary recording servers and two spare servers. All these servers (four of them) are located in the same local network and, as a result, have access to the same cameras. This network can be different from the Global server network.

Recording Servers Setup

To configure failover settings for a recording server in Wavesys WVMS Enterprise, follow these steps:

1. Double-click on a recording server to access its properties editing dialog box.
2. Switch to the Failover tab (remember to fill in the rest of the tabs afterward).
3. Click on the "Change" button to select a pre-created failover cluster from the list. If you haven't created a cluster yet, you can create one by clicking the corresponding button at the bottom of the cluster list.
4. In the "Current failover server" field, it should currently display "none." This field will indicate the currently active failover node once the target recording server fails and is replaced by a failover server.

Here are the explanations for the remaining failover settings:

- **Failover timeout:** Specifies the duration for which the Enterprise server will wait after the last heartbeat message before marking the target server as faulty and replacing it with a failover node. Set this to the minimum value, which is typically ten seconds.





- **Central server connection timeout:** Determines the duration for which the remote server should attempt to receive configuration from the Enterprise server before giving up and using the last known good configuration. Set this to the minimum value, usually one minute.
- **Auto recovery:** If enabled, the remote recording server will automatically start operating once it is back online. If the central server connection is available at that point, the failover server will be stopped by the Enterprise server. Enabling this ensures autonomous system operation, as the failover will continue operating until manually replaced by the primary server if this option is not enabled.
- **Recovery timeout:** Specifies the delay before the Enterprise server activates the target server once it is back online. Set this to zero to allow the recovered server to resume operation without delay.

Repeat these steps for every primary recording server in the list to configure failover settings for each one.

Setup of Failover Nodes

To configure a failover server in Wavesys WVMS Enterprise, follow these steps:

1. Double-click on the server that will serve as the failover node and stay on the Details tab.
2. Enable the Failover node role in the server settings.
3. Switch to the Failover tab. Note that there are fewer settings here compared to a regular recording server since auto recovery settings are not required for failover nodes.
4. Set the Failover timeout parameter to ten seconds. This determines how long the system will wait before replacing a faulty failover server with another failover node.
5. Similarly, set the Central server connection timeout parameter to one minute. This defines the duration for which the failover server should attempt to establish a connection with the Global server before using the last known good configuration.
6. Assign the failover node to the same failover cluster as the primary servers.
7. Save the server configuration by clicking OK.


By following these steps, you have successfully enabled the failover functionality within your Wavesys WVMS Enterprise system for the specified servers.

Realtime Status

In the Monitoring section of WVMS Console, under Servers, you can view the current status of your recording servers as well as the hardware load for connected machines.

If a recording server goes offline and its responsibilities are taken over by a failover node, the failover server's status will be shown as "Substituted" and its failover configuration will display the name of the primary recording server whose configuration is currently being used.





Simultaneously, the faulty recording server will be marked as "Unknown" and displayed in red, indicating its unavailability.

Furthermore, in the server settings under the Failover tab, you can see the current failover substitute for each primary recording server.

By checking the Monitoring section, you can easily monitor the status of your servers and identify any issues or failover situations in your Wavesys WVMS Enterprise system.

What to expect

Once you have successfully clustered and configured the servers as described above, Wavesys WVMS Enterprise is prepared to handle any misbehaviour of the recording servers. If any of the primary recording servers fails, it will be immediately replaced by a failover server. The failover server not only takes over the camera configuration but also maintains the state of the server event and action configuration.

From the perspective of connected clients, such as WVMS Monitor, mobile apps, and others, all failover operations are transparent, ensuring a seamless user experience. Clients continue to receive the requested live streams and recordings without any interruption or indication of server failures. Users may not even be aware that a server switch has occurred.

Recordings made on failover servers are retained until they are erased based on the defined quotas. The individual archive duration quotas set for specific channels apply to all servers, including failover servers. This means that outdated recordings will be automatically erased from the failover servers as well, provided that they are online and connected to the Global server.

With Wavesys WVMS Enterprise failover mechanism in place, your system can maintain continuous and uninterrupted video surveillance, ensuring data integrity and availability even in the event of server failures.



Follow us on

