



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

GVP Deployment Guide

Resource Manager High Availability Solutions

5/13/2025

Resource Manager High Availability Solutions

SIP Server with Resource Manager (RM) and GVP with RM offer several options for HA:

- External Load Balancer
- Virtual IP Takeover Solution Windows
- Virtual IP Takeover Solution Linux
- Microsoft NLB Windows

Each option has conditions and limitations.

External Load Balancer

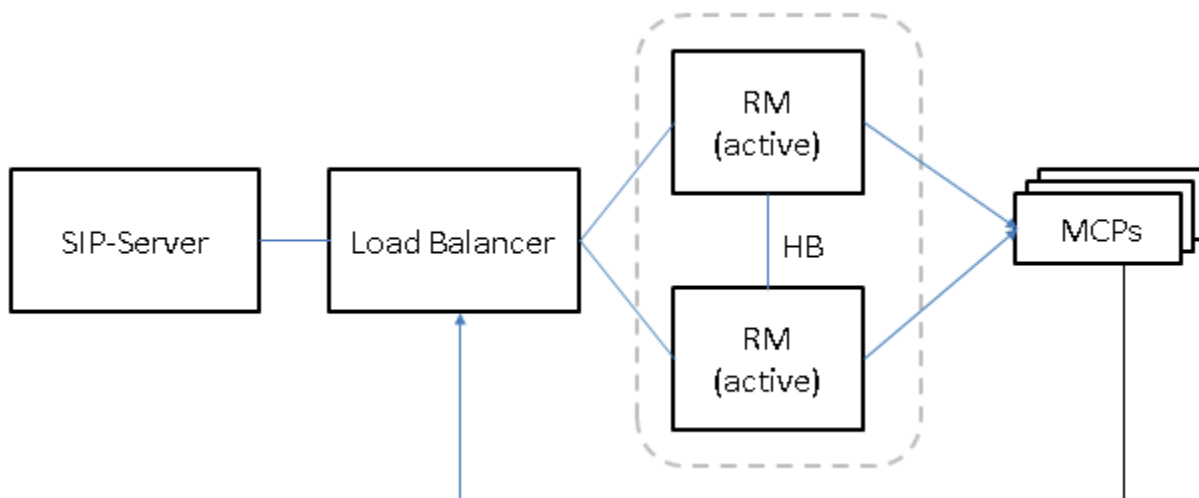


Figure: External Load Balancer (F5/Radware) for RM active-cluster

- This configuration is applicable for RM active-cluster.
- GVP deployment can be Windows or Linux.
- There are separate hosts for SIP-Server, RM and MCPs (preferred).
- SIP-Server goes through the load balancer to RMs.
- RM inserts the Load-Balancer IP address in the Record-Route header so that messages sent within the dialog traverse through the Load-Balancer.
- It is possible, but not preferred, to have RM and MCP on the same host.
- The load balancer must reside on its own host.

Virtual IP Takeover Solution Windows

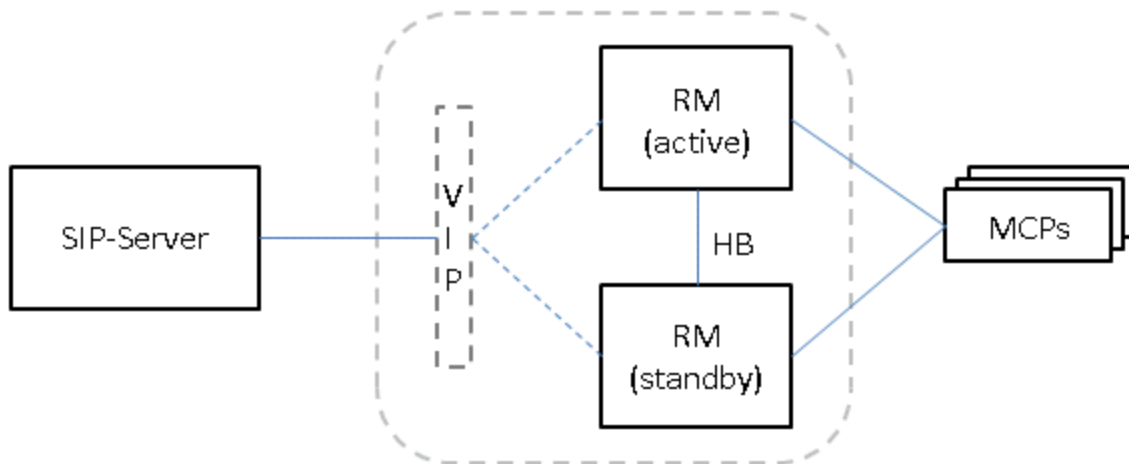


Figure: Virtual IP Takeover Solution (Windows) RM active-standby

- This configuration is applicable for RM active-standby.
- GVP deployment is for Windows.
- Separate hosts are required for SIP-Server and RMs.
- MCPs can be in the same host or in a separate host (preferred) from Resource Manager.
- Virtual IP Takeover is tied to just one NIC.
- Genesys recommends you configure alarm conditions and reaction scripts for handling the failover/ switchover condition in this case.
- Due to ARP cache update issues in Windows 2008, this solution uses the third-party utility arping.

Virtual IP Takeover Solution Linux

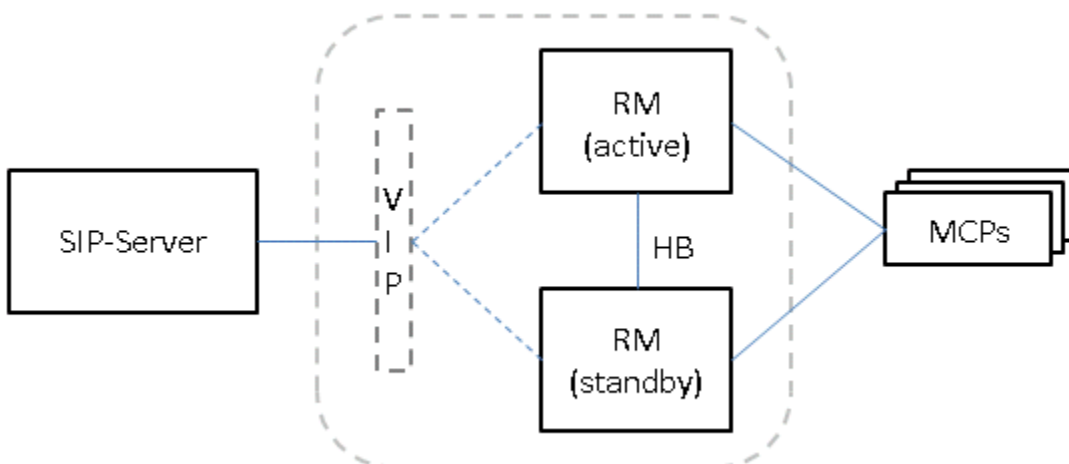


Figure: Virtual IP Takeover Solution (Linux) RM active-standby

- This configuration is applicable for RM active-standby.
- GVP deployment is for Linux.
- Simple IP Takeover for single NIC is supported; if multiple NICs are present in the system, then a Linux bonding driver can be used.
- Separate hosts are required for SIP-Server, RMs and MCPs.
- Genesys recommends that you configure alarm conditions and reaction scripts for handling the failover/ switchover condition in this case.

Microsoft NLB Windows

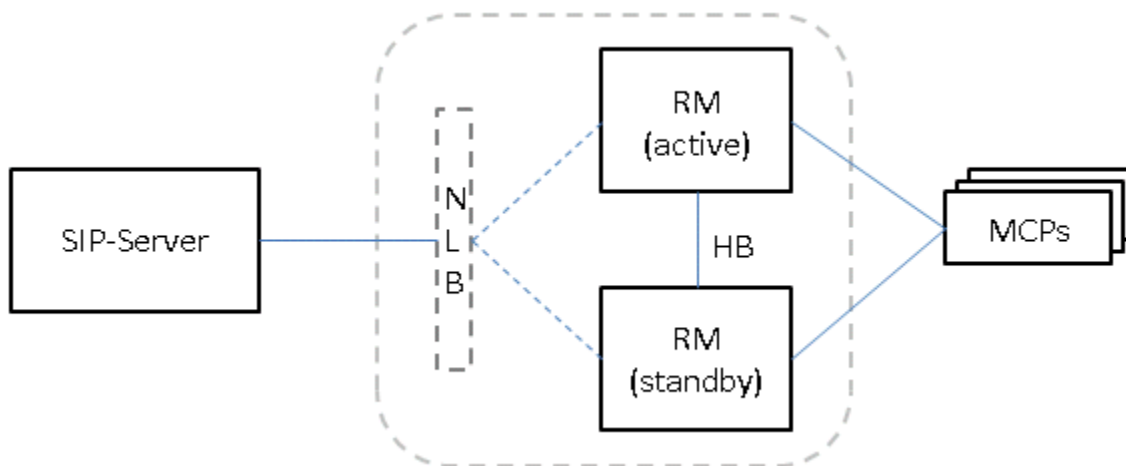


Figure: Microsoft NLB (Windows) RM active-standby

- This configuration is applicable for RM active-standby. GVP deployment is for Windows.
- NLB is used in unicast mode. NLB ensures that other elements can communicate with RMs from outside local network, or within the same subnet.
- NLB configuration requires that RMs be in separate hosts from SIPs or MCPs.
- Each RM host must have multiple NICs; one dedicated to this NLB cluster communications, and the other NIC for non-NLB communications.

HA Using Virtual IP

With this HA solution, multiple hosts share the same virtual IP address, with only one instance actively receiving network traffic. A switchover from active to backup instances in the HA pair can occur with no apparent change in IP address, as far as the SIP dialog is concerned.

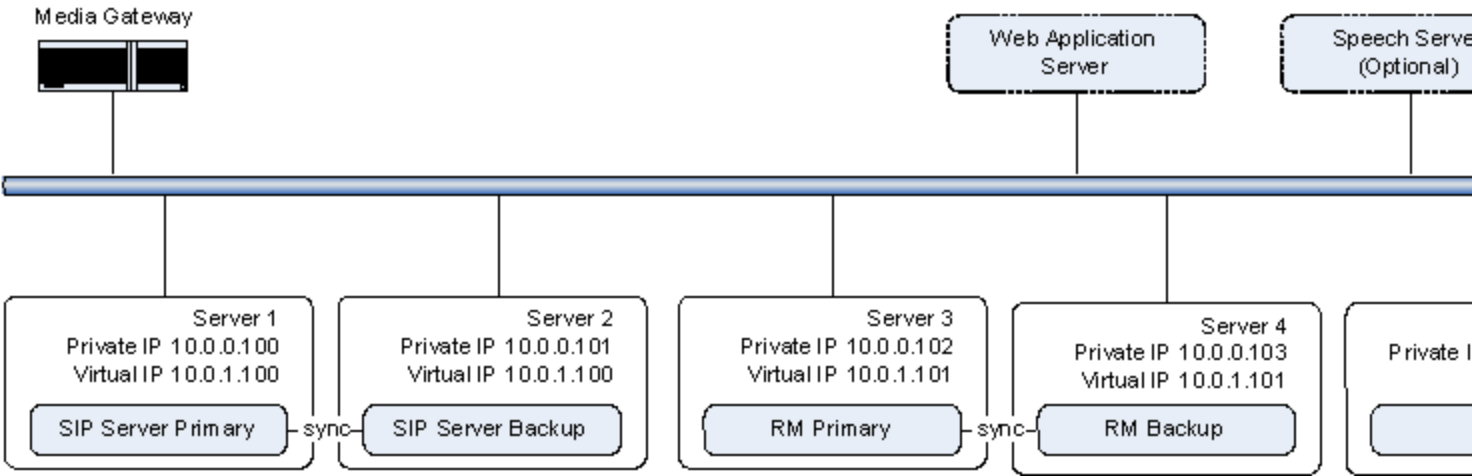


Figure: High Availability Using a Virtual IP Configuration

Feature Limitation

When using Windows Network Load Balancing (NLB) for virtual IP-based HA, only processes running outside the Windows NLB cluster can address that cluster. If SIP Server uses Windows NLB and Resource Manager/Media Control Platform are running on the same machine as SIP Server, then RM/MCP cannot address SIP Server using the cluster address.

With Windows NLB, a local process will always resolve a virtual IP address to the local host. This means that if an MCP process on a particular server tries to contact a failed Resource Manager, Windows NLB will resolve the virtual IP address in the configuration to the local host, and the same local Resource Manager will be contacted, instead of the backup Resource Manager on the backup host.

Sample Configuration

The following task table outlines the basic steps required for an HA deployment on Windows, with the following assumptions:

- This is a two-machine deployment, with one SIP Server and Resource Manager instance co-deployed on each machine.

Task Summary: Configuring HA through Virtual IP for Windows

Objective	Key Procedures and Actions
1. Configure SIP Server instances using Windows NLB.	See the SIP Server Deployment Guide for more information.
2. Configure RM applications	Go to: Provisioning > Environment > Applications

Objective	Key Procedures and Actions
	<ul style="list-style-type: none"> Set the Resource Manager option <code>cluster.ha-mode</code> to <code>active-active</code>.
3. Configure GVP components.	See the GVP 8.5 Deployment Guide for more information.
4. Configure GVP DNs.	<p>Go to: Provisioning > Switching > Switches</p> <ul style="list-style-type: none"> In the GVP Trunk DN, set the contact and contact-backup options to the two Resource Manager IP addresses. For the Voice over service IP Service DN, create two separate DNs with each contact option set to one of the Resource Manager IP addresses.

HA Using an External Load Balancer

This HA method uses an external hardware load balancer to manage clusters of active nodes. The load balancer owns a virtual IP address that is used to forward requests to the active cluster. The load balancer can apply its own load-balancing rules when forwarding the requests.

Use of Active SIP Server Pairs

In some deployments, the customer network does not allow the use of virtual IP addresses. In this case, both SIP Server and Resource Manager pairs can be deployed as active instances. The SIP Server instances are deployed as separate active instances without synchronization. The load balancer will load-balance across the set of active SIP Server hosts. Without the active-backup relationship between the SIP Server instances, SIP Server will lose the state of mid-dialogs if it fails, even though the call will not be immediately dropped.

Important

Genesys does not recommend this HA solution for deployments that require the use of URS routing strategies. If routing is required, Genesys recommends using the load balancer with active-backup detection, so that SIP Server can be deployed as an active-backup HA pair.

Use of Active-Active Clusters for Resource Manager

Resource Manager can be deployed as an active-active cluster, where both instances run together as the active instance, each with a unique IP address. The active pair then synchronizes active session information, so that both instances can correctly route incoming requests.

The next figure shows a sample deployment with an external load balancer, using the following assumptions:

- SIP Server instances are configured as two separate active instances with no synchronization.
- Resource Manager instances are configured as an active-active cluster with synchronization.

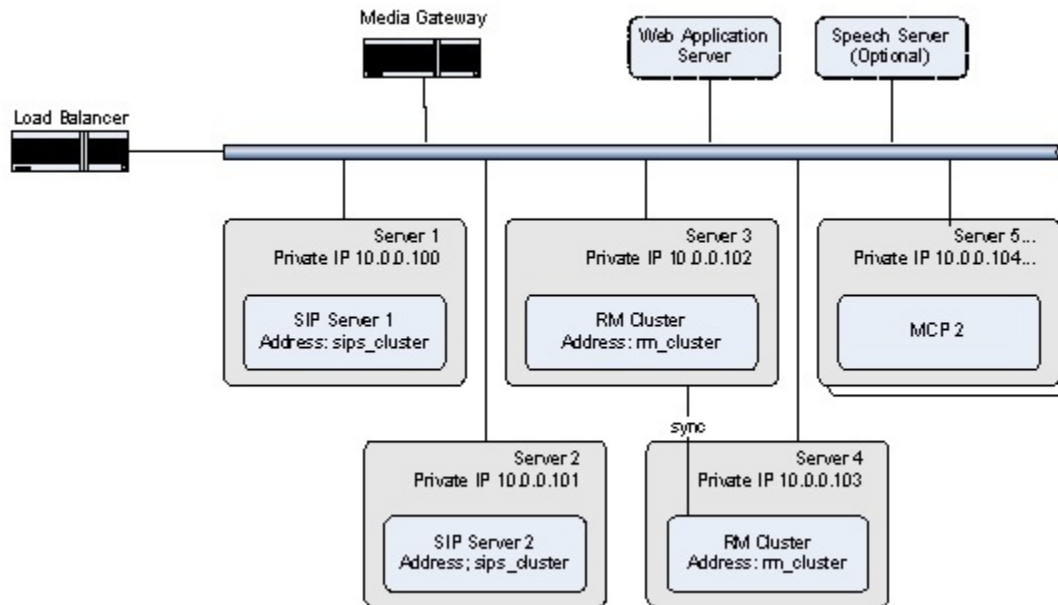


Figure: High Availability Using An External Load Balancer

Resource Manager

GVP 8.5.x supports HA for Resource Manager on both Windows and Linux operating systems.

HA (Windows)

Windows Network Load Balancing (NLB) provides HA for the Resource Manager. You can configure two Resource Managers to run as hot standby, or warm active standby pairs that have a common virtual IP.

Incoming IP traffic is load-balanced by using NLB, in which two Resource Manager servers use a virtual IP number to switch the load to the appropriate server during failover. The network interface cards (NICs) in each Resource Manager host in a NLB cluster are monitored to determine when network errors occur. If any of the NICs encounter an error, the Resource Manager considers the network down, and the load balancing of the incoming IP traffic is adjusted accordingly.

To determine the current status of the Resource Manager at any time, check the traps in the SNMP Manager Trap Console to which the traps are being sent. In the Console, check the most recent trap from each Resource Manager in the HA-pair. If the specific trap ID is 1121, the Resource Manager is active. If the specific trap ID is 1122 the Resource Manager is in standby mode.

In the following example, 5898: Specific trap #1121 trap(v1) received from: 170.56.129.31 at 12/8/2008 6:34:31 PM, the IP address 170.56.129.31 represents the Resource Manager that is active,

therefore, the other Resource Manager in the HA-pair is in standby mode.

Scalability

NLB also provides scalability, because adding Resource Manager hosts to a cluster increases the management capabilities and computing power of the Resource Manager function in the GVP deployment.

Multiple clusters of Resource Manager instances that operate largely independently of one another can be deployed to support large-scale deployments, such as those that involve multiple sites. Each Resource Manager cluster manages its own pool of resources.

Tip

GVP 8.5.x does not support more than two Resource Managers in a cluster.

HA (Linux)

There are two ways to achieve HA for Resource Manager on Linux: by using Simple Virtual IP failover or Bonding Driver failover.

In each of these options, each host in the cluster maintains a static IP address, but all of the hosts share a virtual public IP address that external SIP endpoints use to interact with the Resource Manager hosts in the cluster. If an instance of the Resource Manager fails on any host, the virtual IP address remains valid and provides failover.

When the Bonding Driver failover option is used, two or more network cards are required for the same server and the bonding driver controls the active standby capabilities for the network interfaces.

Resource Groups in HA Environments

When the call-processing components are provisioned in Resource groups, the Resource Manager provides HA for GVP resources in the same way that it normally manages, monitors, and load-balances the resource groups. For example, provided that more than one instance of the Media Control Platform has been provisioned in a VoiceXML resource group, the Media Control Platform service is still available to other VPS components, even if one of the HA provisioned instances is not available.

You can set up your HA Resource Manager environment in one of two ways, depending on the Windows or Linux OS version:

- The **active standby** configuration, in which the active Resource Manager instance only processes SIP requests. Windows 2003, Windows 2008, or RHE Linux 4 or 5 is required.
- The **active active** configuration, in which an external load balancer is used and either of the active nodes can process SIP requests. Windows 2003 or Windows 2008 is required.

For information about configuring the Resource Manager for HA, see [Resource Manager High](#)

Availability.

MRCP Proxy

GVP 8.5 supports the MRCP Proxy in HA mode to provide highly available MRCPv1 services to the Media Control Platform through a warm active standby HA configuration.

To support the MRCP Proxy in HA mode, the latest versions of Management Framework and LCA must be installed and the Solution Control Server (SCS) Application configured to support HA licenses. For more information about HA licenses for the SCS, see the [Framework 8.5 Deployment Guide](#) and the [Framework 8.5 Management Layer User's Guide](#).

Policy Server

GVP 8.5 supports the Policy Server in warm active standby HA mode. The active standby status is determined by the Solution Control Server (SCS), which must be configured to support HA licenses. (See MRCP Proxy, above.) Also, the Policy Server is stateless, therefore, data does not require synchronization.

Reporting Server

GVP 8.5 supports HA for Reporting Server by using a primary/backup paradigm and an Active MQ message store in one of two solutions Segregated Storage or Shared Storage.

Active MQ

The Reporting Server has JMS queues to which data from reporting clients is submitted. The JMS queues implementation used in GVP 8.5 is Active MQ. If the Oracle or Microsoft SQL database is unavailable, but the Reporting Server is still operational, the Active MQ must persist and store the submitted data to the hard disk drive (HDD), thereby ensuring that the data that was submitted to the Reporting Server is not lost if the Reporting Server fails before the database server is restored.

Segregated Solution

In the Segregated Solution, each Reporting Server instance in the cluster uses its own independent Active MQ message store. However, only the server that is designated primary activates its message store. The other independent message store is activated only if the primary Reporting Server fails.

When the Segregated Solution is used, the backup and primary Reporting Servers are configured in Genesys Administrator.

Shared Storage Solution

In the Shared Storage Solution, the Reporting Servers in a cluster share a connection to one Active

MQ message store that receives, queues, and dequeues data from Reporting Clients. Only one Reporting Server instance obtains and holds the Active MQ message-store lock.

When the Microsoft Cluster Service (MSCS) is used, two distinct Reporting Servers access a single shared drive. Switch-over to the backup server occurs when the primary server goes down. After the MSCS is configured, the administrative user interface can be used to manage the primary and backup Reporting Servers.

For more information about configuring the Reporting Server for HA, see [Reporting Server High Availability](#).