

HVR on Google Cloud

The Google Cloud is a commonly selected cloud provider for small and large organizations because of its rich functionality and ease of use. As an all-in-one solution, HVR enables organizations to efficiently replicate data in real-time from a variety of on-premise and cloud platforms to Google Cloud. This datasheet describes HVR's continuous data integration support for Google Cloud technology.



Why HVR for Google Cloud

HVR enables organizations to perform real-time high-volume data replication from on-premises to the cloud, from the cloud to on-premises, and in cloud-to-cloud environments. HVR's distributed and modular architecture is built for the cloud and is designed to give you control over your environment. One set-up of HVR enables you to deploy often and feed streams of data to multiple destinations, continuously and at once.

Installing HVR in Your Environment

An HVR installation is low impact and its modular architecture is designed to perform on relatively high latency, low bandwidth networks. Each installation of HVR includes a hub. The hub is referred to as "the coordinator" because it always initiates communication and determines which firewall must be opened. At any point in time, you can choose where to run the hub (e.g. many on-premises to Google Cloud setups run the hub on-premises and only open the firewall to Google Cloud).

An agent, which is an additional installation of software on or near the source or target, facilitates the movement of changes that occur between systems. Agents don't have to be installed on database servers, but agents

Using tables in an existing database or application, HVR enables you to:



Create tables automatically by mapping your source system data types to compatible loss-less data types in the Google Cloud destination.



Perform an initial load integrated with continuous log-based Change Data Capture(CDC).



Efficiently and continuously move data in real-time to your Google Cloud technologies using log-based CDC.



Check your data with data validation and repair. This is a unique function ensures the data you have moved is correct. Monitor your data statistics using the management console. In the console, you can also set up the data flows and automated alerts.

should be used when communicating across a Wide Area Network (WAN), either between on-premises and the cloud or when moving data between availability zones or different providers in the cloud. When moving data using HVR, the following features ensure performance and security:

- Proprietary compression combined with large data block transfer ensures fast performance maximizing bandwidth utilization irrespective of latency.
- Network communication is secured using SSL/TLS encryption with the use of explicit certificates.
- Certificates are also used for secure two-factor authentication on top of username/password validation.
- Data is securely routed through a proxy so individual systems or databases aren't exposed through the firewall, for example in the demilitarized zone (DMZ)

Supported Services

The majority of data storage on Google Cloud is available on the Google Compute Engine. Some technology vendors such as Snowflake use the Google Cloud infrastructure to run their branded cloud services. We strongly recommend that you use an HVR agent in the availability zone of the source and/or destination endpoint in order to achieve optimal efficiency and performance.

Agent sizing is dependent on the endpoint technology but can be very modest. For automatic scaling, and to provide high availability, HVR can take advantage of Google Cloud's infrastructure service Cloud Load Balancing (CLB).

COMPUTE ENGINE SERVERS

HVR supports all of the technologies installed on the Google Cloud Compute Engine servers as if they were on-premises installations. For example, HVR customers use SQL Server Databases directly running on Google Cloud servers, Greenplum running on Google Cloud, and various other databases and technologies. The list of supported technologies continues to expand, please refer to the [HVR platform support page](#) on the HVR website for a current overview of supported technologies.

CLOUD SQL

Google Cloud supports three different databases; MySQL, PostgreSQL, and SQL Server. HVR supports all three of these databases as a target for data replication. MySQL is also supported as a source for log-based data replication, and changes into SQL Server can be retrieved in near real-time through trigger-based replication.

BIG DATA TARGETS

HVR supports both Google Cloud Storage (GCS) and Google Cloud Big Query as a target for replication.

THIRD-PARTY SERVICES

HVR supports Snowflake running on Google Cloud, with GCS as the staging area for optimum delivery of the initial data set and incremental changes.

[Test drive a cloud-based instance of HVR](#)

[Try now](#)