# SAP HANA ™ – High Availability FAQ


## About this Document

This document answers some frequently asked Questions regarding SAP HANA's High Availability support for Fault and Disaster Recovery. For more information, see the references at the end of this document.

***Updated for SPS7.***

## Legal Disclaimer

## What is High Availability?

***High Availability*** is a set of techniques, engineering practices and design principles for Business Continuity. This is achieved by eliminating single points of failure, and providing the ability to rapidly resume operations after a system outage with minimal business loss. ***Fault Recovery*** is the process of recovering and resuming operations after an outage due to a fault. ***Disaster Recovery*** deals with recovering operations after an outage due to a prolonged datacenter or site failure.

SAP HANA provides a range of fault and disaster recovery solutions to achieve High Availability.

## What is "Enterprise Ready" High Availability?

Enterprises typically have a high level of expectation, regarding business continuity, expressed in the form of key measures around the recovery from faults. One such key measure is RPO: the maximal permissible period of time during which operational data may be lost. Enterprises often design their systems for an RPO of zero; no data should ever be lost. A second common measure is RTO, the maximal permissible time it takes to recover the system at full performance. Enterprise systems are typically designed for very low RTO values, ranging from near-zero to at most a few hours for the most critical business processes.

## What do I need to do to ensure business continuity?

High Availability consists, broadly speaking, of four stages:

- The planning phase
- Prevention and Preparation
- Failure Detection and Recovery
- Fix failure and Failback

During the planning phase, business needs are evaluated against the cost and complexity of the different solutions. Prevention and preparation includes the selection, installation and configuration of redundant components, standby systems and replication. Failure detection involves the integration within existing network/system monitoring and management systems that provide fault indicators via alarms and other messages. Based on the selected high-availability solution, recovery from failures may be automatic or manual. Failback is the process of restoring the redundancy/replication after recovery, to be prepared for a next failure event.

## What downtime scenarios have been considered?

The SAP HANA High Availability solutions address all major faults, such as OS crash, software errors, Operator error, Data Corruption, Disk crash, Component failures (e.g. fan breakdown or PSU burnout), Host crash (general hardware fault), Intermittent memory errors, power outage, cooling failure, network faults, severed cables. In addition, SAP HANA provides recovery solutions in case of entire data center failures (aka disasters). These solutions are also used to manage Planned Downtimes and upgrades.

## What is k-safety?

K-safety is a term used by vendors of databases without reliable two-phase commits, to describe the number of nodes that can fail before loss of data occurs. SAP HANA provides reliability via redo-log saves (thus, the problem of data-loss due to local failure never exists), and additionally supports full system replication to a secondary system.

## Does SAP HANA support clustering?

Other vendors sometimes support a concept they call clustering. Clusters are groups of symmetric hosts that are synchronized on the basis of a single shared disk image of the data. SAP HANA avoids this I/O performance bottleneck by instead employing in-memory data distribution over multiple collaborating hosts. This avoids unwanted data sharing and the associated locks, resulting in enhanced scaling and efficiency.

More details can be found in [6].

## What solutions does SAP offer to ensure business continuity of SAP HANA?

The following table summarizes the main SAP HANA High Availability solutions. See [1] for more details.

| Solution | Supported Downtime Scenarios |
|---|---|
| System Replication | Business continuity for planned downtime, software fault, host crash or disaster |
| Storage Replication | Business resumption after disaster; recovery from local storage corruption |
| Persistence | Restart from software fault, power outage or host crash |
| Backups | Recovery from storage or data corruption, host crash, or disaster |
| Host Auto Failover | Automatic recovery from software fault or host crash |
| Service Auto Restart | Automatic recovery from software fault |

## HANA is an in-memory database. What happens when the power goes down?

Indeed, SAP HANA uses in-memory technology, but of course it fully persists any transaction that changes the data, such as row insertions, deletions and updates. Thus, the database can fully resume from an unexpected power-outage *without any loss of data*, by rolling back incomplete transactions, and recovering fully completed transactions.

## How does HANA recover from a server (host) hardware failure?

SAP offers several complementary solutions, which involve the use of redundant or standby equipment. The main solutions are:

- Host Auto-failover: this solution involves adding one or more passive standby hosts to a SAP HANA system. When one of the active primary hosts fails, due to a hardware failure for instance, the standby host automatically takes over the role of the failed host.
- System Replication: this solution involves the installation of a fully redundant standby system and continuous replication of data changes to the standby system. When the primary system fails, e.g. due to a host-crash, the secondary standby system takes over.

These solutions can be combined, and further combined with traditional backups, which provide additional support for point-in-time recovery, in case of data corruption.

## Does HANA support Disaster Recovery?

Yes. In addition to recovery from traditional backups, SAP HANA provides two solutions that can be used for disaster recovery:

- Storage Replication: partner-delivered solution is a *storage-level replication* solution, which delivers a backup of the volumes or file-system to a remote, networked storage system. Upon failure of the primary system, a secondary can be brought up remotely, using the secondary storage.
- System Replication: this solution involves the installation of a fully redundant standby system and continuous replication of data changes to the standby system. When the primary system fails, e.g. due to a full data-center outage, the secondary standby system takes over immediately.

## What is the difference between a DB restart and a replicated system takeover?

A restart requires most tables to be reloaded to memory, before the database can effectively respond to queries. During a takeover, since column store tables are already loaded into the memory of the secondary system, only a small delta must be loaded, resulting in a very short RTO.

## How does SAP HANA support point-in-time recovery?

SAP HANA provides complete data backups and data snapshots as well as incremental log backups. Combining these backups together with the redo log entries, the database can be recovered to a precise point in time. SAP HANA does not support "incremental backups".

## Are SAP HANA backups compressed and encrypted?

Yes, SAP HANA backups represent the in-memory data image, and therefore compressed. With "Backint for SAP HANA" we offer a 3rd party backup tool API, for instance to stream directly to external backup devices, or to support backup encryption, for example with NetBackup from Symantec. Database administrators also often use storage-level encryption.

## Can I use standby hosts for development/testing?

Yes, if you fully understand the implications of dual usage, then you may use standby hosts for development or testing, on condition that these operations do not interfere with the regular replication or failover functionality.

In the case of storage-replication, the secondary standby hosts can be used for any other purposes, until they are required to assume their active SAP HANA roles. Standby hosts must have enough memory for the intended dual usage. This is particularly true in the case of full system replication, where the memory of the active standby hosts reflects the corresponding primary hosts. Of course, parallel activities must be shut down before a potential failover to the standby system.

See also "Using Secondary Servers for Non-Productive systems" in [2].

## What is the advantage of System Replication without Data Preload?

System Replication without data preload provides an economic option for organizations that require failure or disaster recover without data-loss, but prefer not to dedicate a completely identical secondary system. Without data preload, the secondary system uses a very small amount of memory, and can therefore be used for additional purposes, such as development, QA or staging, while standing by. The tradeoff, of course, is that upon take-over, there will be a period of downtime, during which the secondary system is loaded.

See also "Using Secondary Servers for Non-Productive systems" in [2].

## How about Planned Downtime (near zero-downtime upgrades)?

Absolutely. The System Replication solution can be used to manage planned downtime scenarios with minimal interruption to business processes. Simply perform the required maintenance on the secondary first, fail-over to the secondary. Then, perform the maintenance on the primary and - if required - fail back.

See also "Using System Replication for Near Zero Downtime Upgrades" in [2].

## What are typical recovery timelines?

This depends on the solution. Fundamentally, database recovery consists of the following aspects: fault detection, getting the recovery hardware up and running, restarting the database to accept queries, and, to reach full performance, loading most of the tables back into memory.

SAP HANA's Service Auto-Restart and Host Auto-Failover solutions automatically detect faults in less than a minute. For other situations, e.g. full power outage, fault detection will be part of the data center operating procedures.

Using on-line standby hosts avoids the extra effort and time involved in acquiring, connecting, and rebooting off-line systems, which may take considerable time. Once the hardware is up and running, loading the HANA database is a matter of minutes in the absence of large pre-loading tables, such as row-store tables. Pre-loading tables may take several minutes or more, depending on the size of these tables.

Once preloading is completed, the HANA database is ready to accept queries, and non-preloading tables (by default, column-store tables) will load in parallel.

System Replication with data pre-load keeps all columnar tables in memory in the secondary system. Thus, in addition to a very short recovery time (typically under two minutes for columnar tables), this solution also provides immediate full performance upon recovery. Note: replicated system recovery will take more time if it contains large row tables, which still need to be loaded upon failover; this is currently a SAP HANA limitation, which can be mitigated by converting tables to columnar structure.

Restoring from a *backup* proceeds typically at a rate of about 400-500 GB per hour, however this depends greatly on the storage attributes.

## So, what do you recommend?

As a first level of defense, employ hardware redundancy where possible. In particular, this means following the hardware vendors' guidelines that include a well-mirrored (RAID-enabled) storage system with good auto-detection and recovery from storage failures.

On top of this, SAP HANA System Replication provides the best protection against a variety of failures and downtime cases, and by far the best recovery solution, when considering recovery time, and disaster scenarios. It is therefore the recommended solution for highly critical business systems.

For very large systems, this solution can be further augmented with Host-Auto Failover, using additional standby hosts at both the primary and secondary site, to ensure that both sites remain operational, even in the presence of a local host crash. Furthermore, regular backups should be enabled, to further protect against data-corruptions or operator errors.

However, if a fully redundant, N+N system standby solution with disaster recovery abilities cannot be justified due to implementation cost or complexity, a simpler Host Auto-Failover only solution could be implemented, typically in the form of an N+1 configuration on a single site, augmented by a reliable backup-shipping method to support disaster recovery.

Realistically, the High Availability solution will be further influenced by considerations like timelines, budgets and customer paradigm-preferences which should be discussed in more detail with professional consultants.

## Host Auto-Failover : When would I use more than 1 standby host?

Adding more than one standby host in an Auto-Failover configuration increases the ability to handle multiple faults. For instance, if the chance of a host failing is 1%, the chance of two active hosts *independently* going out of order at the same time would be approximately 1% x 1% = 0.01%. Thus, a decision to deploy more than one standby host is fundamentally driven by *what-if* availability requirements.

## Host Auto-failover: How does HANA avoid "split brain" problems?

One classical problem in auto-failover is a situation called "split brain": rather than a real failure, two hosts are separated due to a network disconnection. If each host thinks it is in charge, this may lead to data corruption.

The HANA Host auto-failover solution avoids a "split-brain" condition, through the use of exclusive storage locking. When a host is determined to be lost (either due to a crash, or loss of communication), one of the standby hosts will attempt to execute a failover, by attempting to take over the storage lock. If this proves unsuccessful (due to the existing host being alive), the failover will abort.

## System Replication: How does distance affect the latency of transactions?

In case of synchronous replication, transactions are delayed by the time of a round-trip network message from the primary to backup site. The longer the distance, the longer the delay. Up to 100 km, delays can be limited to about 5 µseconds on optical fiber. Over longer distances it is recommended to use asynchronous replication, which do not add latency, at a slight data-loss risk (data lost in transit) upon failure.

## Can I use a Standby System for read-only queries?

Not yet, but this capability is on the SAP HANA roadmap.

## What HA/DR features are offered for Business One on HANA?

It is possible to deploy B1 on HANA with SAP HANA System Replication, providing full HA and DR support. Please contact your VAR for more information.

## What is multi-tier System Replication?

From SPS07 you can use System Replication to set up a local secondary system, that acts as the primary for a third (remote) site. One common use is to synchronously replicate to the local secondary system, which asynchronously replicates to the remote system.

## How can clients be configured to be failover-safe?

Clients are typically configured to connect to the SAP HANA system via its domain name. A failover to another server poses a problem – how do clients re-connect upon a failover?

There are several possible approaches to configure HA solutions to allow clients to recover after a failover.

One approach involves client configuration for multiple-connections (supported by the HANA JDBC/ODBC, and (for instance) reverse-proxy for HTTP-based clients. But addressing the problem at this level is undesirable, as it is highly client-dependent, and the solution is not uniform. Additionally, it can lead to unintended results due to "split-brain" scenarios.

Thus, solutions at the network layer are preferred. One approach is to use the IP to domain-name binding, or DNS. Upon failover, the DNS server is configured to return the IP address of the standby system. This solution has the advantage that there are no client-specific configurations. Further, it supports DR confgurations, where the primary and standby systems may be in two different networks (separated by routers). On the other hand, modifying DNS mappings requires a vendor-specific solution, and secondly, due to DNS caching in network elements (both clients and intermediate routers), it can take a long time for DNS changes to propagate.

The alternate approach is at the IP layer, or the binding between IP addresses and MAC addresses, using IP redirection or Virtual IP (VIP) redirection. Provided that the primary and standby system share the same L2 subnet, this is the most comprehensive and recommended solution.

## What happens to "queries in flight" during a failure?

During a failure, the connection to clients is lost. Incomplete in-progress transactions will be rolled back. This is as it should be. Once the failover is complete, clients reconnect to the newly active system, which resumes the handling of queries and transactions.

## SAP HANA system replication: more details

**How does system replication start up?**

After setting up system replication, the primary sends its topology information and immediately starts log shipping to the secondary. It also prepares a full data set, based on snapshots, and transmits it.

From time to time (default every 10 minutes) the *data* that has changed since the last data transport is transmitted to the secondary. (with SPS8, continuous log shipping will make data shipping unneccesary).

**What happens, if the connection between primary and secondary is lost?**

The primary stops shipping of the redo logs, if the secondary does not respond within a configurable timeout (default 30 sec). The secondary will repeatedly try to reconnect. Once the connection is re-established, the missing redo logs are automatically shipped over.

**Can the secondary syetem perform the takeover automatically?**

SAP HANA does not provide automatic failure detection and takeover with System Replication. However, hardware vendors may provide Cluster Management software which is able to do that. SLES for SAP's cluster solution may also be part of an integrated solution.

**What happens to a BW transaction during a takeover?**

The "Connectivity Suspend" feature of the SAP Application server supports real zero downtime maintenance; long running transactions are "suspended" and resumed after the takeover. For more information on the "Connectivity Suspend" feature please refer to SAP note [4]. Additional details regarding connections from database and web clients are in the SAP HANA Master Guide [5].

**Is system replication secure?**

Page-level encryption can be activated separately on the primary and secondary system. SPS7 also supports SSL-encrypted transfer between the primary and secondary system.


## Further Reading

[1] Introduction to High Availability for SAP HANA http://www.saphana.com/docs/DOC-2775

[2] SAP HANA Administration Guide http://help.sap.com/hana/SAP_HANA_Administration_Guide_en.pdf

[3] SAP Note 1730932 - Using backup tools with Backint for HANA

[4] SAP Note 1913302 - Suspend DB connections for short maintenance tasks

[5] SAP HANA Master Guide: http://help.sap.com/hana/SAP_HANA_Master_Guide_en.pdf

[6] Scalability - SAP HANA: http://www.saphana.com/docs/DOC-2022.

More information about SAP HANA can be found on http://help.sap.com/hana_platform/.

Chaim.Bendelac@sap.com
SAP HANA Team

V2.0, February 2014