

Considerations for Disaster Recovery and Automated Failover Clusters for SAP HANA Infrastructures

Q&A with SUSE's Peter Schinagl, Technical Architect, and Markus Gürtler, Technical Alliance Manager

Data corruption, component failures, and power outages are just a few examples of major issues that IT departments must be prepared for at any given moment. In the event of system downtime, organizations must have a proper disaster recovery (DR) and high availability (HA) strategy in place to minimize the business impact. But how can organizations ensure failover and fallback procedures run smoothly? During a recent SAPinsider Live Q&A, Peter Schinagl, Technical Architect, Business Development, and Markus Gürtler, Technical Alliance Manager SAP, Business Development, of SUSE answered questions from attendees regarding DR and HA scenarios for SAP HANA infrastructure, automation solutions for failover, system replication, and more.

This is an abridged transcript of the SAPinsider Live Q&A conversation where Peter and Markus discussed best practices for DR and automated failover clusters for SAP HANA infrastructures. The full transcript is available on SAPinsider Online at bit.ly/HAandDRscenarios.

Q: What are the top causes of system downtime to consider when planning for disaster recovery (DR) or high availability (HA)?

Peter Schinagl (PS): The major issues organizations need to guard against are still the ones we've been seeing for years, such as operating system (OS) crashes, software errors, operator errors, data corruption, disk crashes, component failures, host crashes, power outages, and so forth. These occurrences can never be totally eradicated, so the general idea is to eliminate the single point of failure.

Q: What are the best HA and DR scenarios for an SAP HANA infrastructure?

Markus Gürtler (MG): The best scenario would be an SAP HANA scale-up infrastructure in a performance-based scenario. This means a company would have a minimum of two SAP HANA systems (with one or more nodes per site) that are in an automated failover cluster.

Q: What solutions can companies use for DR?

MG: The best option is probably a three-tier system replication scenario that uses three SAP HANA systems. Two systems are in a failover cluster using SUSE's SAP HANA System Replication Automation solution and are in system replication mode "sync." A third system is located in a geographically different site (a DR location) and connected to the second system in system replication mode "async." With that setup, there are always at least two copies of the in-memory data, one copy in the same location (synchronous replication) and a second copy, with some older data (asynchronous replication), in a separate location.

Q: What happens in the case of an unexpected power outage?

MG: In a scale-up scenario, SUSE's SAP HANA System Replication Automation solution detects the node failure, and the cluster starts the failover to the second node.

Q: Are there DR options other than duplicate stand-by servers?

MG: The DR alternative to SAP HANA system replication is storage replication. This relies on the mirroring functionality of your storage systems (for example, SAN-based mirroring). There are several SAP-certified solutions on the market supported by various storage vendors running on SUSE Linux Enterprise servers (SLES) for SAP software.

Q: Is falling back to a primary node after a failover recommended?

PS: It is possible to have an automated fallback to a primary node, but we would not recommend that. If the primary node was breaking for some reason, research should be performed to identify the root cause before going back to it.

Q: What is the difference between a database restart and a replicated system takeover?

PS: A database restart could take a long time, as it requires reading back all the data from disk to memory. Think about how long a few terabytes would take. With a takeover, by contrast, only some internal pointers need to be recovered.

Q: In the case of a DR scenario, how quickly are end users reconnected after the system switches over to the DR site?

PS: SUSE's SAP HANA System Replication Automation solution provides a failover process that is fully automated. It has a mode where it does synchronous replication from the memory of machine 1 to the memory of machine 2, so the switchover only takes minutes.

Q: In the case of an SAP HANA installation in a virtualized environment using tailored datacenter integration (TDI), does SAP support SAP HANA replication for HA?

PS: Yes. The good thing about SAP HANA system replication is that it is hardware agnostic.

Q: If an organization is using virtual machines (VMs), will it need to use VM-based tools to keep both systems updated?

MG: Yes. SAP HANA just takes care of replicating data inside the SAP HANA database. The database software itself has to be upgraded on involved systems manually or by using other tools, such as SAP Landscape Management. The OS can be patched centrally using SUSE Manager, which is an OS lifecycle and patch distribution system for large SUSE landscapes.

Alternatively, companies can use SUSE Subscription Management Tool (SUSE SMT), which is free and takes care of providing and distributing patches and updates within a SUSE landscape. This process can also be automated. However, the functionality of SUSE SMT is limited when compared to SUSE Manager.

Q: "Zero downtime" suggests that there's no outages for maintenance such as OS upgrades, SAP S/4HANA upgrades, patches, or database upgrades. Is this now possible with SAP HANA?

MG: It entirely depends on each company's architecture. Combining live patching with SUSE's SAP HANA System Replication Automation solution achieves very minimal downtime with all of these things. Kernel live patching includes security patches without any downtime. All other patches or database upgrades would require downtime, but that can be minimized by a failover to a second SAP HANA node using SUSE's SAP HANA System Replication Automation solution. ■



Peter Schinagl is a senior technical architect with the SAP Global Alliance team at SUSE, where he has worked since 2000. He joined the SAP Global Alliance team as an SAP consultant in 2011. In 2016, he moved to business development in the Alliance team and is taking care of Cloud and Strategic Alliances.



Markus Gürtler is a technical alliance manager with the SAP Global Alliance Team at SUSE. Markus is responsible for the technical projects of the strategic SAP and SUSE partnership, including topics like SAP HANA and SAP NetWeaver running on SUSE Linux Enterprise Server for SAP applications and SUSE Open Stack Cloud and Cloud Foundry for SAP landscapes.