




Protecting SAP & SAP S/4HANA in Azure with SIOS Protection Suite





Executive Summary

There are good reasons for migrating SAP and SAP S/4HANA applications to the cloud. The cloud is agile, affordable and dependable— advantageous for mission-critical applications.

Business continuity requires provisions for both high availability and disaster recovery. The Microsoft Azure cloud has provisions for assuring high uptime and it also supports an extensive partner ecosystem of additional options.

This white paper offers some considerations for executives determining how to run mission-critical SAP and SAP S/4HANA applications in the Azure cloud on the Linux operating system. The insights result from long-standing partnerships SIOS Technology has built with both Microsoft and SAP to augment high availability and disaster recovery capabilities in Azure.

This white paper shows how organizations can benefit from a purpose-built high availability and disaster recovery solution for all applications running in a private, public, or hybrid cloud environment.

SIOS Technology and Microsoft

- Long-standing partnership for delivering industry-leading, general-purpose HA/DR protections.
- Gold and Silver Application Development status.
- Certification for Windows Server and Linux, as well as Azure.
- Certification for SAP and other applications, including SQL Server and Oracle.
- Numerous joint “marquee” customer accounts.



SIOS Technology and SAP

- Long-standing partnership for providing HA and HA and DR solutions.
- Recognized as a cost-effective and easy-to-use solution for both Windows Server and Linux.
- Certified by SAP for integration with SAP NetWeaver and SAP S/4HANA.
- Verified by SAP for use with HANA System Replication.





Business continuity via high availability and disaster recovery

The goal for both high availability (HA) and disaster recovery (DR) is the same: to keep the business running by minimizing downtime, eliminating data loss, and maintaining data integrity. But there are some important differences that must be understood before one is able to assess the options available for protecting SAP applications in the Azure cloud.

HA and DR fundamentals

The differences between HA and DR are rooted in the differences between “failures” and “disasters.” Failures are short in duration and small in scale, affecting a server, rack, or the power or cooling in a datacenter. Disasters have enduring impacts and are more widespread, affecting entire datacenters in ways that preclude rapid localized recovery. For example, a tornado, hurricane, or earthquake might knock out power and networks and close roads, making a datacenter inaccessible for days.

Perhaps the biggest difference involves the redundant resources (systems, software, and data), which can be local—on a Local Area Network (LAN)—to recover from a failure. By contrast, the redundancy required to recover from a disaster must be “long distance” across a Wide Area Network (WAN). For applications that require high transactional throughput performance, the ability to replicate the active instance’s data synchronously across the LAN enables the standby instance to be “hot” (in synch with the active instance), ready to take over immediately and automatically in the event of a failure. Such rapid, automatic response should be the goal of all HA provisions.

Because of latency inherent in the WAN, using synchronous replication could have an adverse impact on the throughput performance in the active instance. In this case, data is normally replicated asynchronously in DR configurations. This means that updates to the standby instance always lag behind updates being made to the active instance. Such replication lag makes the standby instance “warm” and results in the need for a manual recovery process that introduces an unavoidable delay in the recovery time. A delay in DR provisions is tolerable because disasters are rare, but can disrupt users.

Because data and information provide great value, low or zero Recovery Time Objectives (RTO) are common for both HA and DR purposes. Recovery Point Objective (RPO) is the maximum period during which data loss can be tolerated. If no data loss is tolerable, the RPO is zero.

There are significant differences, however, in the Recovery Time Objectives normally established for HA and DR purposes. RTO is the maximum tolerable duration of an outage. Mission-critical applications have low RTOs, normally on the order of a few seconds for HA purposes, and high-volume online transaction processing applications generally have the lowest. For HA, synchronous data replication makes it relatively easy to satisfy both a low or zero RPO and a low RTO of a few seconds. For DR, RTOs of many minutes or even hours are common owing to the extraordinary cost of implementing provisions capable of fully recovering from a widespread disaster in just a few minutes.

For the vast majority of applications, four-nines (99.99%) is generally accepted by both database and system administrators as constituting mission-critical high availability. For this white paper, “high availability” is defined as ensuring SAP services and databases are operating when and as needed. “When” takes into account the percentage of time the application is up and running, while “as” takes into account proper operation, with no data loss.



Protecting SAP in the Azure cloud

The options for protecting SAP in the Azure cloud can be organized into four categories:

- Infrastructure redundancy on the Azure cloud.
- Capabilities included with or available for the Linux operating system software.
- Complexities with open source HA/DR projects.
- Features available in the SAP software.
- Purpose-built third-party failover clustering software.

Infrastructure redundancy on Azure

The Azure cloud offers redundancy within datacenters, within regions and across multiple regions. Within datacenters, redundancy is provided by Availability Sets that distribute virtual machines (VMs) across different fault domains in different racks to protect against failures at the rack level. Availability Sets afford redundancy for most hardware failures, but provide no redundancy for a datacenter-wide failure.

For protection from single datacenter-wide failures, Azure has Availability Zones (AZs). Regions that support AZs have at least three datacenters that are networked with sufficiently high bandwidth and low latency to accommodate synchronous replication. To protect against major disasters, Azure offers Region Pairs, in which every region gets paired with another within the same geography, such as the US, Europe, or Asia. The pairs are strategically chosen to enable rapid recovery during widespread network or power outages, or major natural disasters. Azure also allows customers to choose and establish

replication mechanisms and streams between Azure regions of their choice.

Capabilities for Linux operating system

Options available for Linux fall into two categories: open source and commercial, with some of the latter often incorporating some of the former. Either approach must include, at a minimum, support for three related capabilities: data replication, server clustering, and resource management with heartbeat monitoring.

Data replication forms the foundation of high-availability clusters, putting it at the bottom of the HA stack. Distributed Replicated Block Device (DRBD) is an open source block-level data storage system that replicates data in a distributed manner. DRBD is the dominant choice because it is included with the Linux kernel in virtually all distributions.

The Corosync cluster engine is used to synchronize messaging, create replicated state machines, implement a quorum system, and provides features to restart application processes that have failed. Pacemaker is open source software for managing compute and storage resources in HA clusters. It provides the heartbeat monitor. Its daemon and cluster resource manager are used to handle application services to maintain high availability.

Complexities with open source HA/DR projects

Using these open source software components to build your own HA/DR framework for mission-critical applications like SAP does not make sense given both the complexity of building such frameworks and the criticality of the systems protected. Especially factoring in: the cost of maintaining, versioning, and management of any scripts along with the retention of and training required for people to deploy, maintain, and update these scripts. Commercial HA/DR solutions for Linux are available. All are easier to

use than DIY solutions, but most still require some customization, and specialist expertise. Especially for companies who use multiple different Linux distributions or a spread of different releases within a Linux distribution, a need for simplification and unification of their HA infrastructure is of high priority. Unified processes of operating and managing HA architectures across various Linux releases is another priority often named. The need for unified and simpler handling and setup of such HA architectures across different applications can be another reason to look for an alternative approach to create HA architectures on Linux.

SIOS Protection Suite for Linux is a commercial HA/DR solution that helps organizations eliminate DIY complexities and limitations using turnkey Application Recovery Kits purpose-built for SAP.

Features available in SAP

SAP has two of its own high availability features: HANA System Replication and Host Auto-Failover. As its designation implies, HANA System Replication (HSR) makes a copy of the SAP system consisting of both the services and the in-memory database. HSR can be configured to operate in one of three different modes:

- Synchronous with the primary system waiting until the secondary system has received data and persisted it to disk.
- Synchronous in-memory with the primary system waiting until the secondary system has received data.

- Asynchronous with the primary system proceeding without replicating data (unless the asynchronous log buffer is full and waiting is configured).

For HA, the two synchronous modes are ideal. For DR, the asynchronous mode is normally used to avoid adversely impacting the primary system's throughput performance. While HANA System Replication is necessary for HA and DR, it is not sufficient by itself. First, not all of the SAP services and data are replicated, Second, SAP, in the case of Linux, does not provide orchestration for an automatic failover from a failing active node to a passive node (the replication target). Instead SAP relies on frameworks offered for Linux.

The SAP services and non-database data stored on the different Azure storage solutions must be replicated by some other means—either between two VMs in an Azure Availability Set or Availability Zones, or in the DR case, even between different Azure Regions.

A significant disadvantage with all application-specific options, whether for SAP or any other application, is the need for administrators to use different HA and DR solutions for different applications. Having multiple solutions inevitably increases complexity and costs, making this another reason why application-agnostic third-party solutions are so popular. And these third-party HA/DR solutions have another major advantage: They are purpose-built to supplement options in Azure for SAP running on Linux.

Mini case study: Bonfiglioli

Bonfiglioli is a leading Italian design, manufacturing and distribution company specializing in industrial automation, mobile machinery, and wind energy. The company's 3600 employees depend on SAP-based enterprise resource planning applications to keep business operations running smoothly around the globe and around the clock. One of the goals established for migrating the SAP environment to the Azure cloud was to improve on the service levels achieved in the enterprise datacenter. To ensure success, the company enlisted the help of BGP Management Consulting, which recommended using a SIOS SANless failover cluster. A key reason for BGP's recommendation is the strategic partnerships SIOS has with both Microsoft and SAP, which have resulted in the solution being proven to work well in the Azure cloud.



SIOS Protection Suite for Linux

High availability requires two parts—application orchestration and data replication. Application orchestration is the monitoring of application health and, in the event of an issue, automatically failing over operation of the application to remote VM(s) in compliance with application best practices. Data replication, as discussed, ensures remote VMs have access to the most current data after the failover is completed to comply with near-zero RPO requirements.

SIOS Protection Suite for Linux provides a complete HA/DR clustering software solution that includes application-aware failover orchestration. The suite provides out-of-the-box, SAP-certified high availability protection. It includes Application Recovery Kits for SAP that automate and dramatically simplify the implementation, testing, and ongoing operation of HA and/or DR protections in full compliance with SAP-specific best practices. Note that customers can use the application orchestration software with HANA System Replication (HSR) or the real-time data replication included in the suite.

SIOS clustering software: A closer look

SIOS Protection Suite for Linux is a complete, high availability failover clustering solution that provides a tightly integrated combination of data replication, continuous application monitoring, and configurable failover/failback recovery policies. The suite is implemented entirely in software and utilizes a shared-nothing architecture to ensure full protection for the full SAP application stack—from the server, storage, and networking hardware through all system and application software.

The suite has been designed to deliver high availability in a way that preserves high performance. Data replication is provided by SIOS Protection Suite's DataKeeper software for all data stored in the SAP filesystem or in the storage of Azure's SANless infrastructure. Data is replicated synchronously across multiple Availability Zones within one Azure region for

While the focus in this white paper is on Linux, most of the capabilities offered by [SIOS Protection Suite for Linux](#) are also included with the SIOS DataKeeper Cluster Edition for Windows Server along with Microsoft Windows Failover Clustering.

SIOS DataKeeper for Windows Server is available in both a Standard Edition and a more robust Cluster Edition. The Standard Edition provides real-time data replication for DR protection in a Windows Server environment. The Cluster Edition provides seamless integration with Windows Server Failover Clustering—enabling customers to use WSFC for orchestration in a SANless clustering environment in Azure or other clouds.

HA protection, and can optionally be replicated asynchronously across Azure Regions for DR protection, independent of the Azure block storage chosen.

For the data stored in memory in SAP HANA, SIOS Protection Suite leverages HANA System Replication (HSR). In both cases, the SIOS LifeKeeper software is used for failure detection, alerting and alarming, and (automatic) failover and failback actions.

Figure 1 shows a simple two-node failover cluster containing a single standby node, which should be located in a different Availability Zone. It is important to note that the active instance of SAP's Enqueue Replication Server (ERS) runs in the standby instance, which is depicted in the diagram by giving it a white vs. grey background. For this reason, its data is replicated in the reverse direction from all other data in the SAP filesystem.

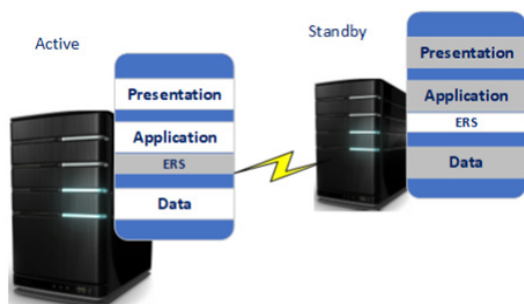


Figure 1. This simple two-node SANless failover cluster has a single standby instance, which should be located in a separate Availability Zone.

When a problem is detected, SIOS Protection Suite automatically initiates one or more of three configurable recovery actions: attempt to restart the application on the same server; fail over to a standby server; and alert a system administrator, optionally, to take action manually. Combinations are also possible, such as always alerting an administrator even when the automatic failover is initiated immediately, or failing over to a standby server only after first attempting a restart that fails. Automatic failovers can also be paused under

certain circumstances to enable an administrator to approve or reject the action.

Because high-availability solutions with sophisticated capabilities like these are often difficult to use, SIOS Protection Suite offers a simple wizard-driven implementation with an intuitive graphical user interface for “single pane of glass” monitoring and management. Right-click convenience makes it easy to allocate resources and specify recovery policies with automatic and manual failover/failback options. These options also help simplify HA and DR testing and enable planned maintenance to be performed with minimal downtime. Implementation and operation are made even easier with Application Recovery Kits purpose-built for SAP.

SIOS Application Recovery Kits

For most Linux high-availability solutions, achieving the desired operation is subject to the dreaded “some assembly required” caveat. This is often true for solutions using open source software, but is typically also the case for most commercial offerings. And owing to the enormous complexity in the scripts and configuration files in the full HA stack, the “assembly” is prone to errors that result in failover provisions failing when needed. SIOS Protection Suite for Linux avoids this common headache with turnkey recovery kits for popular applications.

With SIOS Application Recovery Kits (ARKs), the configuration wizard is pre-populated out of the box with appropriate choices, thereby minimizing the need to make changes, while also enabling changes to be made when desired. Each ARK is typically complete with configurations for the entire infrastructure, including for storage, physical and virtual server instances, network resources, application processes and services, and even those elements that are unique to the Azure cloud.

SIOS offers multiple Application Recovery Kits that can be used with SAP. The ARK for SAP NetWeaver deployments provides continuous monitoring

to assure that all databases are mounted and available, that all file shares, mounts and exports are available, and that clients are able to connect.

The specific resources monitored include the physical servers and virtual machines, the Linux operating system, SAP Central Services, ABAP SAP Central Services, the Primary Application Server, and the Enqueue Replication Server, as well as any back-end databases, including Oracle, Sybase, DB2, Max DB, MySQL, and PostgreSQL.

A separate ARK is provided for use with SAP S/4HANA deployments. This ARK leverages HANA System Replication to protect the in-memory database and uses the Suite’s DataKeeper data replication software to protect all other services and data in the SAP environment.

The ARKs for SAP NetWeaver and SAP S/4HANA both leverage other ARKs used to protect networked services commonly shared among multiple applications. Those pertinent to SAP include Network File System (NFS) mounts and exports, Logical Volumes (LVMs), and IP and Virtual IP addresses.

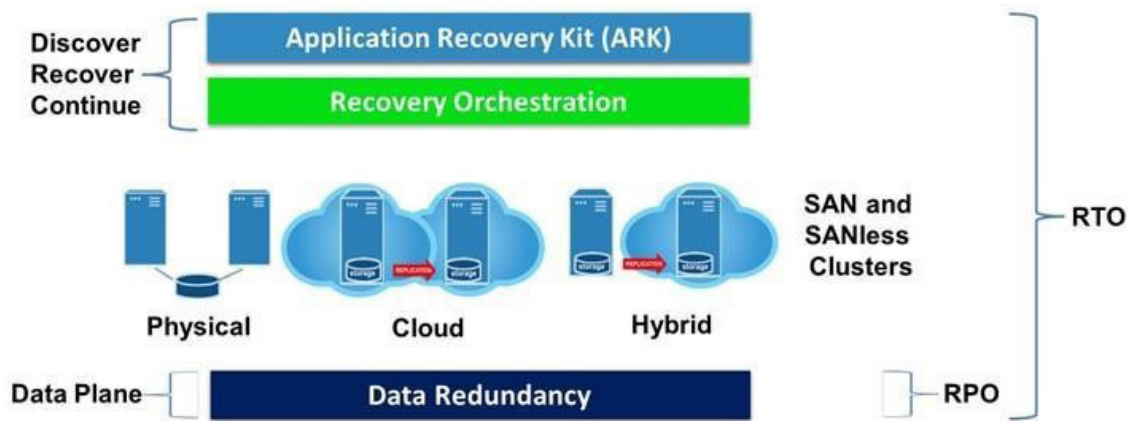


Figure 2. Application Recovery Kits provide templates that dramatically simplify configuring, testing, and maintaining failover clusters.



Putting the SIOS Protection Suite to work

While there are many different ways to configure SAP services and data, these three are the most popular:

- All-in-One SAP Services and Data
- Separate SAP Services and Data Tiers
- SAP Multi-node S/4HANA Database

All-in-One SAP Services and Data is the easiest to protect because it requires only a single standby instance for both the services and data. This is the configuration depicted in the simple two-node SANless failover cluster example in Figure 1.

Separate SAP Services and Data Tiers requires at least two active instances and two standby instances, one each for the services and data. In cases where the Network File System (NFS) is not integrated with the other data, a third active and a third standby instance are required. This is the configuration depicted in the example in Figure 3.

SAP Multi-node S/4HANA Database protection leverages HANA System Replication for the in-memory database, but requires separate provisions for all other SAP services and data stored on conventional media.

Table 1 compares the three configurations.

Comparing configurations

● Good ◐ OK ○ Poor

	All-in One SAP Services and Data	Separate SAP Services and Data Tiers (NFS Integrated)	SAP Multi-node (S/4HANA)
Management <i>The ease of installation and operation</i>	●	◐	○
Reliability <i>Likelihood the system is working as expected</i>	○	○	●
Availability <i>Probability that the system remains operational</i>	◐	◐	●
Serviceability <i>The ease of servicing/maintaining the system</i>	◐	◐	●
Cost <i>The Total Cost of Ownership (TCO)</i>	●	●	○
Performance <i>The ability to optimize for performance as needed</i>	○	◐	●
Scalability <i>The ability to expand the system as needed</i>	○	○	●

Table 1. The table provides a summary comparison of all three configurations.

Figure 3 shows a SANless failover cluster that provides both HA and DR protections for a tiered deployment that separates the SAP Central Services from the data, and utilizes separate physical servers for each. The servers in the middle are active; the servers on the left are standby in a separate Availability Zone for HA protection; and the servers on the right are standby in a separate region for DR protection.

When a problem is detected, SIOS Protection Suite takes advantage of synchronous replication to automatically and immediately fail over both the services and data to their standby nodes in

the other Availability Zone. You also have the option to attempt a restart and alert the system administrator. Should the region with the two Availability Zones experience a widespread disaster, a manual recovery process will be needed to failover the services and data to their standby nodes in the separate disaster recovery Azure Region.

Three-node configurations like this have another advantage: They make it possible to update the hardware and software one node pair at a time while still protecting the SAP environment by rotating the assignment of the active node pair.

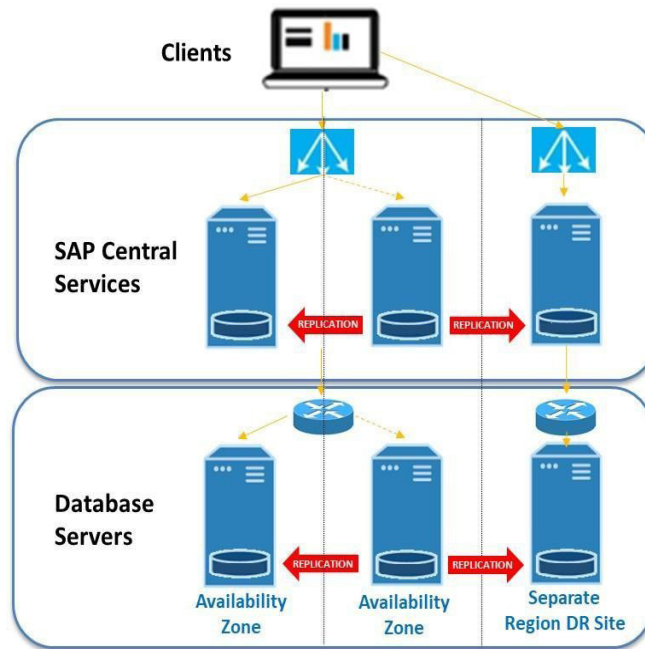


Figure 3. This SANless failover cluster for SAP NetWeaver, configured with separate tiers for SAP and the data, provides HA protection across two Azure Availability Zones, as well as DR protection with a third node located in a separate Azure Region.



Conclusion

SIOS Protection Suite for Linux provides robust high availability and disaster recovery protection for SAP applications running in the Azure cloud. The suite's SANless failover clustering is proven in practice to deliver carrier-class protection with failover across Availability Zones, and to do so without incurring a carrier-like total cost of ownership. The Application Recovery Kits make it even easier to configure clusters correctly

for dependable operation. And when needed, SIOS offers responsive support and a variety of professional services to assure your success.

For Windows environments, SIOS DataKeeper can be used with Windows Server Failover clustering to create a SANless clustering environment across Availability Zones for disaster protection.

To learn more about how SIOS Protection Suite for Linux or SIOS DataKeeper for Windows Server can help protect your SAP NetWeaver or SAP S/4HANA applications, please visit SIOS Technology on the Web at <https://us.sios.com/solutions/sap-high-availability>.

While there, be sure to check out how SIOS SANless failover clusters protect other applications, and take advantage of the Free Trial offer.