SIOS WHITEPAPER

Protecting SAP & SAP S/4HANA in Amazon EC2 with SIOS Protection Suite

High Availability and Disaster Recovery

us.sios.com

Executive Summary

Amazon Web Services (AWS) provides many attractive benefits for organizations looking to migrate an SAP or SAP S/4HANA landscape to the cloud. It's agile, affordable, and scalable. But no cloud environment is wholly immune from spot outages, catastrophic outages (though these may be rare) or even site-wide and regional disasters. While AWS has provisions for ensuring availability at the infrastructure level, the business-critical nature of SAP landscapes warrants more protection. That is, specialized attention needs to be paid to the application and database layers. There several options for providing this added protection.

This white paper looks at best practices for running an SAP or SAP S/4HANA on Linux distributions, including Oracle Linux, Red Hat, SUSE, and CentOS, on AWS. The recommendations reflect SIOS Technology's 20+ years of experience in providing expert HA/DR protection for critical applications, databases, and ERPs in on-premises, public cloud, and hybrid cloud environments.

Ensuring Business Continuity

Ensuring business continuity, that is, keeping the business running by minimizing system downtime, avoiding data loss, and ensuring data integrity, is what HA/DR solutions are designed to do. But a few critical distinctions need to be made to ensure SAP and SAP S4/HANA is sufficiently protected when running on AWS. First, HA at the infrastructure layer is not the same as HA at the application and database layers. Second most IT professionals define application (and database) HA as 99.99% uptime. That is, it is the ability to interact with the application, database, or ERP as well as their storage data 99.99% of the time. Third, 99.99% HA at the application and storage layers requires that disaster recovery (DR) is an intrinsic part of the landscape.

HA and DR: Interrelated Requirements

Think of the difference between HA and DR as the difference between an outage and a disaster. An outage is a failure that affects a discrete component or components within the service delivery infrastructure—a server, a rack, a network segment; perhaps the power or cooling system in a data center. A disaster is more expansive, enduring, and consequential — a tornado or earthquake, for example, that cuts power and network access to an entire data center. Recovery from an outage may take only take seconds or minutes; recovery from a disaster may require days, weeks, or months. Note that while AWS provides service level agreements (SLAs) for the availability of their infrastructure, there are many reasons that outages can happen at the application and storage layers that are not covered by these agreements.

For these reasons, to provide HA for SAP and SAP S4/HANA in the face of an outage requires the deployment of redundant resources (systems, software, and data) in a location close enough to be on the same local area network (LAN) that can be called into services immediately after the outage is detected. DR protection is also needed in which these resources are deployed at a great enough distance that the disaster affecting the primary infrastructure is unlikely to affect the redundant infrastructure. It's essentially on a wide area network (WAN).

In a typical configuration, a failover cluster is configured on AWS with each cluster node connected to local storage. Efficient replication is used to synchronize storage between the primary and secondary cluster nodes so that, in the event of an outage, the secondary node will take over and access identical data as the primary node. Ensuring such synchronization of storage when the components of an SAP landscape are actively executing transactions and processes requires replication of transactions and database updates—hence the importance of the high-speed LAN connecting the infrastructures in an HA solution.

For applications like SAP and S/4HANA that require high transactional throughput performance, an HA solution must be able to ensure that the standby instance in a multinode HA cluster is a is kept running with all the same software and data that the active instance has—so that it can take over immediately and automatically in the event of an outage.

Because of latency inherent in the WAN, a DR solution can't rely on synchronous replication between the primary and backup infrastructures. Instead, a DR solution replicates data asynchronously, which can result in a slight lag in the synchronization state of the two infrastructures.

Failover Across Availability Zones

A critical component of HA is ensuring redundant cluster nodes are able to fail over across both cloud availability zones (AZ) but also across regions. This enables geographic separation to enable critical SAP and HANA systems to continue to run through local, site-wide, and even regional outages and disasters.

Protecting SAP on AWS

A number of mechanisms exist to support the availability of an SAP landscape on AWS:

- Infrastructure redundancy on AWS.
- Capabilities included with or available for the Linux operating system.
- Features available in the SAP software.
- Purpose-built third-party failover clustering software.

Given the importance and complexity of an SAP landscape, some mechanisms, though valuable in other use cases, are sub-optimal when it comes to ensuring HA and DR of an SAP landscape on AWS.

Infrastructure Redundancy on AWS

AWS offers multiple levels of infrastructure redundancy—within data centers, within regions, and across regions. Within data centers, redundancy is provided by distributing virtual machines (VMs) across different Placement Groups in different racks to protect against failures at the rack level. This affords redundancy for many hardware failures, but provides no redundancy for a data center-wide failures.

For protection against failures affecting a single data center, AWS EC2 configurations can be distributed across one or more Availability Zones (AZs). An individual AZ comprises one or more data centers that are networked with sufficiently high bandwidth and low latency to accommodate synchronous replication between the data centers in the AZ. You may also choose and establish the replication mechanisms and streams between the AWS regions you select.

Capabilities for Linux Operating System

Regardless of the method you choose, to provide HA/DR for SAP and S4/HANA on AWS in a Linux environment (Oracle Linux, SUSE, Red Hat, and CentOS), your solution needs to support four related capabilities: efficient data replication at the storage layer, application health monitoring, application failover orchestration and a cluster quorum function. Let's compare open source HA options, SAP HA features, and purpose-build HA/DR solutions.

Open Source Options:

- Data replication lies at the heart of HA clusters. One replication solution, the Distributed Replicated Block Device (DRBD), is an open source block-level data storage system that replicates data in a distributed manner. DRBD is included in virtually all distributions of the Linux kernel.
- The Corosync cluster engine is also available in some open source distributions. It can 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 used in SUSE HAE, Red Hat, and other open source operating systems that can manage compute and storage resources in HA clusters. It can also orchestrate failover of application services to maintain high availability. Pacemaker also monitors the system heartbeat monitor and its daemon and cluster resource.

While each of these services offers clear value, they add a level of complexity and failover instability that should be considered carefully before using in SAP and SAP S/4HANA landscapes.

- **Error Prone Manual Scripting Required.** Open source solutions are add-ons to operating systems that are not specifically optimized for HA/DR. All of their set up, configuration, and resource creation steps require complex, manual scripting. None of these solutions validate inputs or prevent misconfiguration steps.
- **No SAP or SAP S/4 HANA Application Intelligence.** Open source solutions do not contain application-aware intelligence that guide configuration or ensure failovers happen according to SAP best practices.
- **Ongoing Management Complexity.** The time and cost of maintaining, versioning, and managing scripts, along with the expertise needed to deploy, maintain, and update these scripts, can be significant. More importantly, this complexity often results in delays and inefficiency in performing failover testing, upgrades and patches, and worst of all - unstable and unreliable failovers.
- **Issue Resolution.** Open Source vendors are OS-centric, lacking the expertise needed to resolve issues in complex SAP HA clustering including network, storage, and application..

SAP HA Features

SAP offers high availability features that cover two of the four critical application clustering areas: replication and failover via the HANA System Replication (HSR) and Host Auto-Failover. As its name implies, HSR makes a copy of the HANA database - including both the services and the in-memory database and can be configured to operate in one of three different modes:

- Synchronous replication, with the primary system acknowledging data writes only when the secondary system has written the data to disk.
- Synchronous in-memory replication, with the primary system acknowledging data writes when the secondary system has received data.
- Asynchronous replication, with the primary system writing data to disk without acknowledgement of the write status of the secondary node (unless the asynchronous log buffer is full and waiting is configured).

The two synchronous modes work well across multi-AZs, as inter-AZ networking can support the necessary bandwidth and network latency for synchronous replication. For HA within systems in a single AWS AZ, the two synchronous modes operate well. For DR spanning Amazon EC2 infrastructures in separate AWS Regions, the asynchronous mode is normally used to avoid slowing performance of the primary systems, which would otherwise have to wait for data write confirmations to be returned from the remote infrastructure.

While HSR is necessary in an HA or DR configuration that relies on SAP-native options, it is an incomplete solution. It does not replicate all of the SAP services and data and therefore, non-replicated SAP services and data must be replicated by some other means. In addition, HSR in the Linux environment does not provide the automated failover orchestration services that are required for an automatic failover from a failing active node to a passive node (the replication target). For these services, SAP relies on frameworks offered for Linux.

Purpose-built Failover Clustering Software

Most organizations have more than one mission-critical application to protect. A significant disadvantage with all application-specific and OS-vendor HA and DR options is the inability of those tools to extend HA and DR protection to other applications within your organization. Administrators must rely on different HA and DR solutions to support the other mission critical applications and having multiple solutions inevitably increases complexity and costs.

This is one reason why purpose-built HA and DR solutions are so popular. Certain third-party HA and DR solutions have another major advantage when it comes to deployment on AWS including configuration and management automation and ease-of-use that enable IT to implement, configure and maintain more stable HA clusters quickly and easily. They also enable failover/failback testing without disruption to business operations - another key element in meeting stringent 99.995 HA uptime requirements.

SIOS Protection Suite for Linux

One such purpose-built solution is SIOS Protection Suite for Linux. SIOS Protection Suite is a complete HA and DR clustering software solution that includes efficient application health monitoring, application failover orchestration; multiple cluster quorum options, and optional block level replication. It also SAP Netweaver and HANA certified and includes SAP and HANA specific Application Recovery Kits (ARKs) that automate and dramatically simplify the implementation, testing, and ongoing cluster operations

The SIOS Protection Suite includes SIOS LifeKeeper for application monitoring and failover orchestration. The SAP and HANA ARKS monitor the entire environment - application, OS, storage, and networking. The ARKs provide added intelligence SIOS LifeKeeper, enabling failover orchestration in compliance with SAP-specific best practices. Wizard-driven configuration self-validate inputs and prevent common mistaken entries. SIOS LifeKeeper leverages HANA System Replication (HSR).

Application Awareness

SIOS ARKs for SAP deployments provide continuous monitoring of

- Servers - both physical and virtual machines
- Linux operating system,
- SAP Central Services,
- ABAP SAP Central Services,
- Primary Application Server,
- Enqueue Replication Server.
- Any back-end databases—including SaP HANA, Oracle, Sybase, DB2, Max DB, MySQL, and PostgreSQL

SAP Certified

SIOS has a long-standing partnership with AWS and has numerous joint customer accounts. SIOS software has also undergone stringent testing to be certified by SAP for integration with SAP NetWeaver, and SAP S/4HANA. and Verified by SAP for use with HANA System Replication.

Advanced Automation

Automated, wizard-driven configuration validates inputs, and ensures clusters will failover according to SAP and S4/HANA best practices. Easy switchover and switchback enables easy failover testing and ongoing rolling upgrades and maintenance without disruptions to end users.

SAP® Certified
Integration with SAP NetWeaver®

SAP® Certified
Integration with SAP HANA®

Figure 1 below shows a tfailover cluster containing one primary and one standby node, which should be located in a different data center within an AWS AZ. It is important to note that the active instance of SAP's Enqueue Replication Server (ERS) runs in the standby instance, which is illustrated in the diagram with a white background rather than a grey background. For this reason, its data is replicated in the reverse direction from all other data in the SAP filesystem.

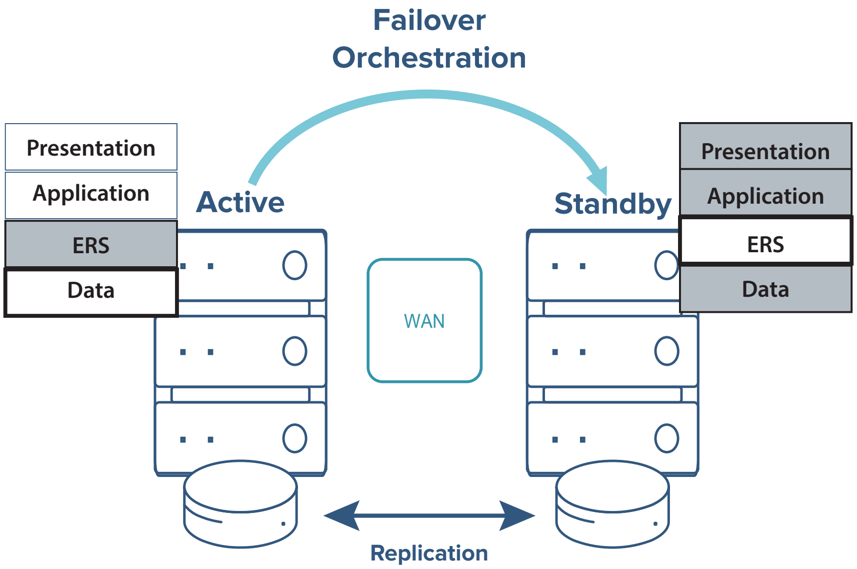

Figure 1. This two-node SANless failover cluster has a single standby instance, which is located in a separate data center within an AWS Availability Zone. The SAP ERS is active on the secondary node according to SAP best practices.

Built-in Intelligence Avoids Unnecessary Failovers

Unlike open source solutions that respond to all application issues will full failovers, SIOS Protection Suite provides multiple, configurable interim recovery actions:

- It attempts to restart the application on the same server
- It fails over to a standby server
- It alerts a system administrator, optionally, to take action manually.

Combinations of actions 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.

In a common scenario, an SAP out of memory alert would trigger a full failover in an open source solution. In contrast, SIOS clustering would flag the issue, initiate a restart, and provide clear reason for it.

SIOS Application Recovery Kits

SIOS Protection Suite includes several ARKs designed specifically to facilitate configuration and management of HA clusters to protect SAP, HANA, on AWS. SAP and HANA ARKs automate and dramatically simplify configuration and ongoing management of SIOS Protection Suite. 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.

With these ARKs, a configuration wizard is pre-populated out of the box with appropriate choices- speeding configuration while enabling changes to be made if desired. Each ARK is typically complete with configuration options supporting an entire application infrastructure, including storage, virtual server instances, network resources, application processes and services, and even those elements that are unique to AWS.

Several SIOS ARKs 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.

A separate ARK supports SAP HANA deployments, leveraging HANA System Replication to protect the in-memory database. The ARKs for SAP NetWeaver and SAP S/4HANA both leverage other built-in intelligence 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.

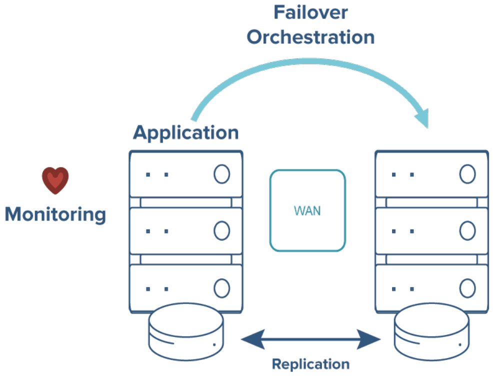

Figure 2. SIOS Protection Suite includes failover orchestration, monitoring, and Application Recovery Kits that dramatically simplify configuring, testing, and maintaining failover clusters.

Using SIOS Protection Suite with SAP in Amazon EC2

While there are many different ways to configure SAP services and data with instances on AWS, 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 in the Amazon cloud for both the services and data. This is the configuration depicted in the two-node failover cluster example in Figure 1.

Separate SAP Services and Data Tiers requires at least two active instances and two standby instances on AWS one each for the services and data. 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.

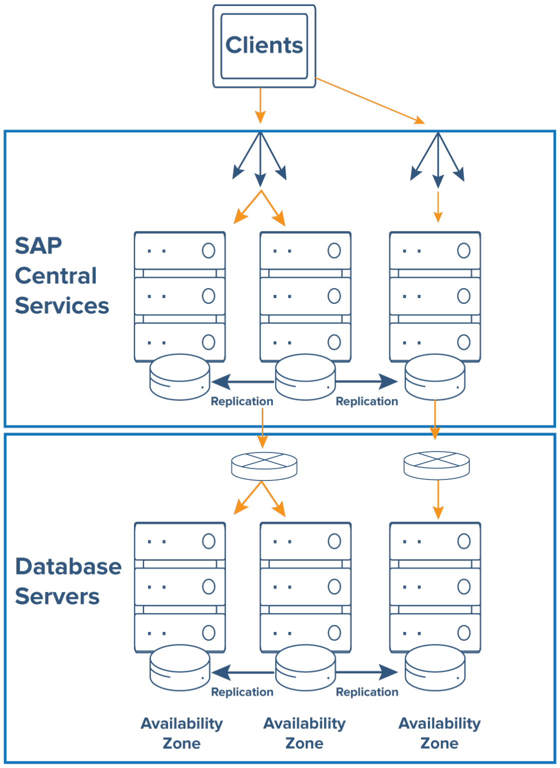

Figure 3. This SANless failover cluster for SAP NetWeaver, configured with separate tiers for SAP and the data, provides HA protection across two AWS Availability Zones, as well as DR protection with a third node located in a separate AWS Region.

A SANless failover cluster can be created (see Figure 3) that provides both HA and DR protections for a tiered deployment that separates the SAP Central Services from the data, and uses separate physical servers for each. The servers in the middle are active; the servers on the left are standby in a separate AWS 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, the 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 AWS region with the two AZs 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 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.

Efficient Replication

SIOS offers efficient replication in configuration where SAP is not used with HANA, such as SAP on Oracle Database in an Oracle Linux. SIOS DataKeeper, provides application-agnostic, block-level data replication services between storage on primary and secondary VMs on AWS. Data is replicated synchronously across multiple Availability Zones within one AWS region for HA protection, and can optionally be replicated asynchronously across AWS Regions for DR protection, independent of the AWS storage service chosen.

HA-Expert Technical Support

For more than 20 years, SIOS' customer experience team has been providing expert support to customers protecting mission critical applications from downtime and disasters. SIOS support professionals are not only experts in HA/DR on prem and in the cloud, they have a deep understanding of the entire IT infrastructure environment. They undergo rigorous onboarding and ongoing training and many hold advanced cloud certifications in AWS EC2 operations. They have clear lines of communication with SAP, operating system vendors, and AWS for fast, accurate issue resolution.

Comparing Configurations

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

Good ● Fair ◐ Poor/None ○	All in One SAP Services and Data	Separate SAP Services & Data Tiers (NFS Integrated)	SAP Multinode (S4/HANA)
Ease of Configuration	●	◐	○
Reliability of Failover	○	○	●
99.99% Availability	◐	◐	●
Ease of Maintenance	◐	◐	●
TCO	●	●	○
Performance Optimization	○	◐	●
Scalability	○	○	●
HA/DR Expert Support	◐	◐	●

Conclusion

SIOS Protection Suite for Linux provides robust HA and DR protection for SAP landscapes running on AWS. SIOS application-aware failover clustering delivers carrier-grade protection with failover across AZs—in an easy-to-implement, easy-to-manage solution for low TCO. SIOS Protection Suite ARKs make it even easier to support HA and DR configurations of an SAP landscape correctly for dependable operation. And when needed, SIOS offers responsive support and a variety of professional services to assure your success.

SIOS Technology Corp.
155 Bovet Road, Suite 476
San Mateo, CA 94402
Tel: 650-645-7000

info@us.sios.com
https://us.sios.com