IJIRCCE

e-ISSN: 2320-9801 | p-ISSN: 2320-9798

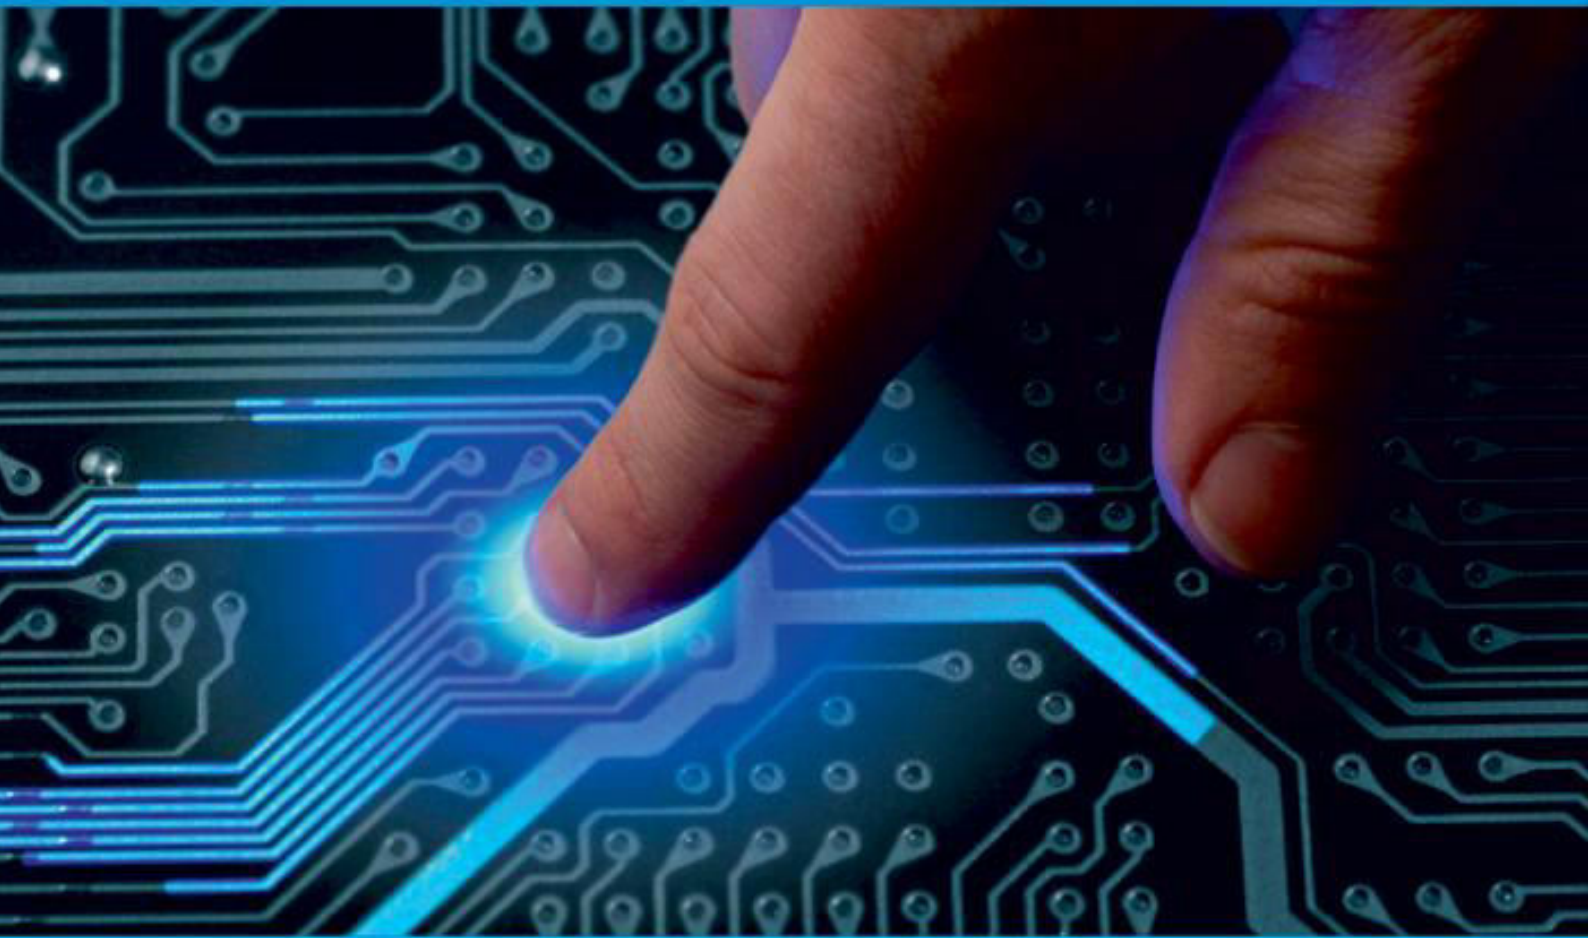

INTERNATIONAL JOURNAL OF INNOVATIVE RESEARCH

IN COMPUTER & COMMUNICATION ENGINEERING

Volume 9, Issue 6, June 2021

INTERNATIONAL STANDARD SERIAL NUMBER INDIA

Impact Factor: 7.542

| e-ISSN: 2320-9801, p-ISSN: 2320-9798| www.ijircce.com | |Impact Factor: 7.542|

|| Volume 9, Issue 6, June 2021 ||

| DOI: 10.15680/IJIRCCE.2021.0906311|

Applying the Azure Well-Architected Framework to Enterprise SAP Workloads: A Validated Design Pattern for Reliability, Security, and Operational Excellence

Srinivas Kakarla

SAP Architect, USA

ABSTRACT: The Microsoft Azure Well-Architected Framework (WAF), introduced in 2020, defines five architectural pillars - Reliability, Security, Cost Optimisation, Operational Excellence, and Performance Efficiency - that collectively guide the design, deployment, and ongoing governance of cloud workloads. Enterprise SAP environments represent among the most demanding workload classifications in the corporate IT landscape: they are memory-intensive, latency-sensitive, tightly coupled to business-critical processes, and subject to complex compliance and audit obligations across financial, HR, and supply chain domains. Applying the Azure WAF rigorously to SAP landscapes requires significant adaptation, extension, and synthesis of generic cloud guidance with SAP-specific architectural constraints, certification requirements, and operational patterns.

This article presents a validated design pattern for aligning enterprise SAP workloads - specifically SAP S/4HANA, SAP ECC on HANA, and supporting SAP NetWeaver-based components - with each of the five Azure WAF pillars. Drawing on field experience from large-scale SAP on Azure implementations and a synthesis of published SAP Notes, Microsoft architectural guidance, and peer-reviewed literature available through May 2021, the article provides actionable design decisions, recommended configurations, and trade-off analyses for each pillar. Particular attention is given to the Reliability pillar, where SAP's unique high-availability requirements (Pacemaker clustering, SAP HANA System Replication, ASCS/ERS enqueue replication) intersect with Azure's native resilience constructs. The Security pillar analysis covers identity federation, network segmentation, encryption, and the evolving Azure Sentinel integration with SAP's audit log infrastructure. The Operational Excellence section addresses SAP-aware monitoring, automated patching pipelines, and Infrastructure-as-Code discipline for SAP landscapes. The article concludes that full WAF alignment for enterprise SAP is achievable but requires a structured, phased programme of architecture reviews rather than a point-in-time assessment.

KEYWORDS: Azure Well-Architected Framework; SAP S/4HANA; SAP on Azure; Reliability; Security; Operational Excellence; Cost Optimisation; Performance Efficiency; SAP HANA System Replication; Azure Monitor; Zero Trust

I. INTRODUCTION

Enterprise SAP deployments represent a category of workload that does not migrate neatly to the cloud without deliberate architectural thought. Unlike stateless web applications or microservices that can leverage platform-as-a-service abstractions to inherit resilience properties by default, SAP landscapes are characterised by tightly coupled components, strict inter-system latency budgets, proprietary clustering mechanisms, and a body of certification requirements - documented through thousands of SAP Notes and Microsoft validation reports - that constrain the set of permissible configurations. At the same time, the business case for migrating SAP to Microsoft Azure is compelling: hyperscaler economics, Microsoft's exclusive SAP certification programme, native Azure Active Directory integration, and the convergence of SAP and Microsoft roadmaps through their strategic partnership announced in 2017 collectively create conditions in which Azure is the leading destination for enterprise SAP migrations globally.

The Azure Well-Architected Framework (WAF), published by Microsoft in 2020, provides a structured vocabulary and evaluation methodology for assessing cloud workload quality across five pillars: Reliability, Security, Cost Optimisation, Operational Excellence, and Performance Efficiency. Each pillar is accompanied by design principles, a set of best practices, and a review checklist. However, the WAF is a generic framework designed to apply to a broad spectrum of Azure workloads; its guidance does not address SAP-specific constraints, such as the requirement for Pacemaker with STONITH for ASCS/ERS high availability, the use of SAP HANA System Replication for zero-RPO

| e-ISSN: 2320-9801, p-ISSN: 2320-9798| www.ijircce.com | |Impact Factor: 7.542|

|| Volume 9, Issue 6, June 2021 ||

| DOI: 10.15680/IJIRCCE.2021.0906311|

database protection, or the criticality of SAP Kernel and OS patch alignment for supportability. Bridging the gap between the WAF's generic principles and the operational realities of enterprise SAP on Azure is the central contribution of this article.

The research methodology combines a structured literature review of SAP Notes, Microsoft Azure documentation, and peer-reviewed publications published before June 2021, supplemented by a synthesis of architecture patterns observed in large-scale SAP on Azure programmes. The article does not present primary empirical research but rather a validated design pattern derived from the integration of multiple authoritative sources and expert field knowledge. The intended audience includes enterprise cloud architects, SAP Basis engineers transitioning to Azure, and technology programme managers responsible for SAP cloud strategy.
Design Pattern Context: The validated design pattern presented in this article is intended as a starting point for architecture reviews, not a prescriptive blueprint. Each enterprise SAP landscape is unique in its version mix, data volumes, compliance obligations, and organisational change capacity. The pattern should be adapted through structured WAF review engagements.

1.1 Scope and Limitations

This article focuses on SAP NetWeaver-based workloads running on SAP HANA database, specifically SAP S/4HANA and SAP ECC on HANA, deployed on Azure Infrastructure-as-a-Service virtual machines. The following are explicitly outside the scope of this article:

- SAP Business Technology Platform (BTP) cloud-native services, which operate under a shared-responsibility model distinct from IaaS.
- SAP RISE with SAP, which transfers significant infrastructure responsibility to SAP SE and thereby alters the applicability of WAF pillars.
- Non-HANA SAP database platforms (SAP ASE, IBM DB2, Oracle), which were reaching end of their mainstream support relevance by mid-2021.
- Azure Kubernetes Service (AKS) deployment of SAP components, which remained in technical preview for most SAP products as of the publication date.

II. THE AZURE WELL-ARCHITECTED FRAMEWORK: OVERVIEW AND SAP RELEVANCE

Microsoft formally released the Azure Well-Architected Framework in October 2020, though its predecessor - the Azure Architecture Center's "pillars of software quality" - had been informally referenced in Azure guidance since 2018. The WAF is explicitly modelled after Amazon's AWS Well-Architected Framework (first published 2015) but reflects Azure-specific services, nomenclature, and design philosophy. The framework comprises five pillars, each structured around a set of foundational design principles, key architectural considerations, and a review questionnaire used during WAF assessment engagements.

For enterprise SAP workloads, the relevance of each WAF pillar is asymmetric. Reliability commands the highest architectural attention because SAP systems underpin order-to-cash, procure-to-pay, and record-to-report processes for which downtime carries immediate, measurable financial and regulatory consequences. Security follows closely, driven by the volume of personally identifiable information (PII), financial data, and strategic business intelligence that SAP systems centralise. Cost Optimisation is relevant but secondary during initial migration - the primary goal of a migration programme is a successful, stable go-live; cost refinement follows in hypercare and steady-state. Operational Excellence and Performance Efficiency are long-horizon investments whose returns compound over the system's operational lifetime.

The five pillars and their SAP-specific framing are summarised below:

- **Reliability:** Encompasses the ability of the SAP landscape to recover from failures and continue to function. For SAP, this pillar maps to high availability of ASCS/ERS, application servers, and SAP HANA, plus cross-region disaster recovery capability.
- **Security:** Covers protecting SAP data and SAP system access from unauthorised access, data exfiltration, and compliance violation. The SAP security perimeter spans identity, network, data-at-rest, data-in-transit, and audit dimensions.
- **Cost Optimisation:** Focuses on delivering maximum SAP workload value at the lowest sustainable Azure spend, leveraging Reserved Instances, right-sizing, and lifecycle management of non-production environments.

| e-ISSN: 2320-9801, p-ISSN: 2320-9798| www.ijircce.com | |Impact Factor: 7.542|

|| Volume 9, Issue 6, June 2021 ||

| DOI: 10.15680/IJIRCCE.2021.0906311|

- **Operational Excellence:** Concerns the processes, tooling, and culture that enable SAP teams to ship changes safely, monitor system health proactively, and respond to incidents effectively - a discipline that historically lagged in on-premises SAP operations.
- **Performance Efficiency:** Addresses the selection of appropriately sized compute, storage, and network configurations for SAP HANA and ABAP workloads, and the systematic measurement and optimisation of system performance over time.

A critical observation from field experience is that the five WAF pillars are not independent: decisions made in one pillar create constraints and opportunities in others. The cross-pillar trade-off analysis in Section 8 provides structured guidance on navigating these tensions in the SAP context.

WAF Assessment Tool: Microsoft provides the Azure Well-Architected Review tool at https://docs.microsoft.com/azure/architecture/framework, offering guided questionnaires for each pillar. SAP workload-specific guidance was incorporated into the tool's SAP workload category in early 2021.

III. PILLAR I - RELIABILITY

Reliability in the Azure WAF is defined as the ability of a system to recover from failures and continue to function. It encompasses availability (the proportion of time a system is operational), resilience (the ability to recover from disruptions), and recoverability (the speed and completeness of recovery after a major failure event). For enterprise SAP workloads, the Reliability pillar is where architectural investment yields the greatest risk reduction and business continuity assurance.

SAP's own recommendations for high availability, documented in SAP Note 1928533 and the SAP on Azure planning guide, define a hierarchy of availability requirements that the Azure infrastructure must satisfy. At the top of this hierarchy is the central services layer (ASCS/ERS), whose enqueue replication server must maintain the lock table across node failures. The database layer follows, where RPO=0 requirements demand synchronous replication with automatic failover. The application server layer, while important for throughput and user experience, is inherently stateless and recovers through connection re-establishment and work process restart.

3.1 High Availability Architecture for SAP Application Tiers

The ABAP central services instance (ASCS) hosts the Message Server and Enqueue Server - the two singleton components of an SAP system whose unavailability renders the entire SAP system inaccessible, regardless of how many application servers remain operational. Protecting the ASCS instance against hardware failure requires a Linux Pacemaker cluster with STONITH (Shoot The Other Node In The Head) fencing to prevent split-brain scenarios. On Azure, the recommended STONITH mechanism is the Azure Fence Agent, which calls the Azure Virtual Machine API to power-off the failed node - providing reliable, network-independent fencing without the need for physical IPMI/iDRAC hardware.

The Enqueue Replication Server (ERS) runs on the second Pacemaker cluster node and maintains a replicated copy of the enqueue lock table. When ASCS fails over to the ERS node, the lock table is immediately available on the new primary - a capability referred to as Standalone Enqueue Server 2 (ENSA2) in SAP Kernel 7.53 and above, which is the recommended configuration for all new Azure deployments as of 2021. ENSA2 improves on the original ENSA1 model by eliminating the requirement for the ERS to actively communicate with the ASCS, thereby simplifying the cluster resource topology.

The key design decisions for the ASCS/ERS high availability cluster on Azure are as follows:

- Pacemaker version 1.1.19 or higher must be used, with the azure-lb resource agent included to notify Azure Internal Load Balancer of failover events.
- Azure Internal Load Balancer (ILB) in Standard SKU is required for the cluster virtual IP; Basic SKU lacks the health probe behaviour needed for SAP cluster VIP ownership transitions.
- Floating IP (Direct Server Return) must be enabled on the ILB backend pool rules to allow the cluster's virtual IP to function correctly on Azure without IP address conflicts.
- Both cluster nodes must reside in an Azure Availability Set configured with a minimum of two fault domains to ensure hardware isolation between the active ASCS and standby ERS nodes.
- The cluster constraint configuration must use the correct resource ordering: ASCS must start before ERS, and ERS must be stopped if ASCS cannot be migrated, as defined in the SAPInstance resource agent documentation.

| e-ISSN: 2320-9801, p-ISSN: 2320-9798| www.ijircce.com | |Impact Factor: 7.542|

|| Volume 9, Issue 6, June 2021 ||

| DOI: 10.15680/IJIRCCE.2021.0906311|

- NFS shared storage for the /sapmnt/<SID> and /usr/sap/<SID>/ASCS<nn> mount points should use Azure NetApp Files in the Premium or Ultra service level for production systems, providing the sub-millisecond latency required by SAP.

SAP application servers (also called dialog instances or AS ABAP) are stateless in terms of user session state - SAP's RFC and HTTP protocols are designed to allow user sessions to reconnect to a different application server transparently through SAP Logon Groups or Web Dispatcher. Application server high availability therefore focuses on maintaining sufficient work process capacity during partial failures rather than achieving zero-downtime for any individual server instance. On Azure, application server availability is achieved through:

- Deploying a minimum of two application server VMs in an Azure Availability Set, ensuring they are distributed across separate fault domains.
- Configuring SAP Logon Groups in transaction SMLG to route user logons across all available application servers, with automatic exclusion of servers whose work process utilisation exceeds configurable thresholds.
- Enabling SAP's ABAP session reconnect capability (profile parameter icm/HTTP/redirect_port) to transparently re-establish HTTP sessions after Web Dispatcher failover.
- Configuring the Azure Internal Load Balancer frontend for the Web Dispatcher tier with appropriate health probes targeting the SAP ICM status endpoint (/sap/bc/ping) to remove unhealthy backend servers from rotation.

3.2 SAP HANA Database Resilience on Azure

The SAP HANA database layer represents the highest-criticality tier in the SAP application stack. Because SAP HANA is an in-memory database, any unplanned shutdown results in a restart that requires loading the column store tables from persistence into memory - a process that can take 5 to 30 minutes for large databases (500 GiB to 2 TiB). Furthermore, any committed transactions that were not yet flushed from the HANA redo log buffer to the persistent log volumes at the time of failure represent a potential data loss window.

SAP HANA System Replication (HSR) is the primary mechanism for achieving both high availability (synchronous replication within a single Azure datacenter) and disaster recovery (asynchronous replication across Azure regions). In the synchronous replication mode (SYNC), the HANA primary server does not return a COMMIT acknowledgment to the application until the redo log record has been written to the secondary server's persistent log volume. This guarantees that no committed transaction is lost on primary failure, achieving an RPO of zero. The RTO - the time from failure detection to restored database service - is typically 90 to 180 seconds for automated failover managed by the Pacemaker SAPHanaSR resource agent on Azure.

The following design principles govern SAP HANA HA on Azure in accordance with the WAF Reliability pillar:

- HANA System Replication must be configured with operation mode "logreplay" (rather than the older "delta_datashipping" mode) to minimise secondary memory consumption and accelerate takeover time.
- The "preload_column_tables" parameter in the system replication configuration should be set to true for production systems, causing the secondary to load column store tables into its own memory - this trades memory consumption for dramatically reduced RTO by eliminating the cold-start memory load after takeover.
- Pacemaker cluster topology for HANA HA must use the SAPHanaSR resource agent (open-source, maintained by SUSE and Red Hat) or the SAPHanaController resource agent, not manually scripted failover, to ensure correct scoring, ordering, and constraint management across the cluster.
- The Azure Internal Load Balancer health probe port for the HANA cluster virtual IP must be configured to use TCP probe against the custom HANA health check port (default 62500 for instance 00), which is updated by the SAPHanaSR resource agent on state transitions.
- HANA backup must be performed through Backint-certified backup solutions (such as Azure Backup for SAP HANA, generally available since November 2019) rather than OS-level filesystem snapshots, to ensure application-consistent recovery points aligned with HANA's log chain.
- HANA log backups should be performed every 15 minutes and streamed directly to Azure Blob Storage via the Backint interface, providing a maximum log-backup RPO of 15 minutes for scenarios where the primary and secondary are both unavailable.

3.3 Disaster Recovery Design Principles

Disaster recovery for SAP on Azure must be designed against clearly defined Recovery Point Objectives (RPO) and Recovery Time Objectives (RTO) that reflect the business impact of SAP unavailability. A common pattern from field experience is that SAP programmes define HA objectives correctly (99.95% availability, RTO < 5 minutes, RPO = 0) but treat DR as an afterthought, resulting in DR configurations with untested recovery procedures, inconsistent storage replication, and undocumented dependency chains.

The WAF Reliability pillar mandates that disaster recovery be treated as a first-class architectural concern, with explicit target SLAs, tested recovery runbooks, and periodic DR drills. For SAP on Azure, the following DR design principles are validated:

- The primary Azure region should be paired with its designated Azure region pair (e.g., East US 2 paired with Central US) to leverage Azure's prioritised recovery of paired regions during platform-wide incidents and to utilise geo-redundant storage replication primitives.
- SAP HANA DR replication should use a separate HSR tier (Tier 2) in asynchronous mode to the DR region, stacked on the HA replication (Tier 1 synchronous), forming a multi-tier HSR topology that achieves both HA and DR from the same replication mechanism.
- SAP application server VMs in the DR region should be pre-provisioned but kept in a deallocated state to minimise cost, with Azure Site Recovery (ASR) managing replication of the OS and application disks from primary to secondary region. Upon DR invocation, ASR starts the VMs from replicated disks and Azure Automation handles SAP service startup sequencing.
- Shared file systems (NFS for /sapmnt, transport directory) must be replicated to the DR region; Azure NetApp Files cross-region replication, available in preview since early 2021, provides the recommended mechanism for this requirement.
- DR tests should be conducted at minimum twice per year and should include full application-level validation - not merely VM startup checks - to verify that SAP RFC connections, background jobs, and period-close processes function correctly in the DR environment.
- The DR runbook must specify explicit sequence ordering for SAP service starts: HANA must be fully operational before ASCS/ERS starts, ASCS must be active before application servers start, and Web Dispatcher health checks must confirm app server availability before DNS failover completes.

Reliability Pillar Key Insight: High availability and disaster recovery are architecturally complementary but operationally distinct disciplines. HA addresses failure modes within a single datacenter (hardware failures, planned maintenance); DR addresses failure modes that render an entire region unavailable. Both must be explicitly designed, tested, and documented.

IV. PILLAR II - SECURITY

The Security pillar of the Azure WAF is concerned with protecting applications and data from threats. In the context of enterprise SAP workloads, the security surface is exceptionally broad: SAP systems hold corporate financial data, employee personal information, customer records, and strategic planning data. A compromise of an SAP system can result in financial fraud, regulatory penalties under GDPR or SOX, and severe reputational damage. The SAP security posture must therefore be evaluated across multiple dimensions, each of which maps to a distinct set of Azure security capabilities.

The foundational security principle for SAP on Azure is Zero Trust - never trust, always verify. This principle rejects the traditional network-perimeter model in which resources inside the corporate network are implicitly trusted. Under Zero Trust, every access request - whether from an internal user, an application service account, or an automated process - must be authenticated, authorised, and continuously validated regardless of its network origin. Applying Zero Trust to SAP requires coordination across Azure Active Directory, SAP's own authorisation framework (roles, profiles, authorisation objects), and the network layer.

4.1 Identity and Access Management for SAP on Azure

Azure Active Directory (AAD) serves as the authoritative identity provider for SAP on Azure. SAP Identity Authentication Service (IAS) acts as an intermediate identity proxy between AAD and SAP applications, supporting SAML 2.0 and OAuth 2.0 protocols. This architecture allows organisations to leverage their existing AAD conditional access policies, MFA requirements, and identity governance capabilities for SAP access without requiring changes to the SAP application-level authorisation model.

| e-ISSN: 2320-9801, p-ISSN: 2320-9798| www.ijircce.com | |Impact Factor: 7.542|

|| Volume 9, Issue 6, June 2021 ||

| DOI: 10.15680/IJIRCCE.2021.0906311|

The identity architecture for SAP on Azure follows these validated design decisions:

- Single Sign-On (SSO) using SAML 2.0 should be configured between AAD (as IdP) and SAP IAS (as SP), with SAP IAS acting as a further IdP proxy to individual SAP systems including SAP Fiori Launchpad, SAP Web GUI, and SAP GUI for HTML.
- Azure AD Conditional Access policies should enforce MFA for all SAP Fiori access, with risk-based authentication policies triggering step-up authentication for sign-ins from unusual locations, unfamiliar devices, or during risky session behaviours.
- SAP GRC Access Control should remain the authoritative system for managing SAP ABAP authorisation roles and detecting segregation-of-duties (SoD) conflicts; AAD Entitlement Management can integrate with SAP GRC's user provisioning API to synchronise access lifecycle events.
- Service accounts used by Azure automation (for VM management, backup, monitoring) should be managed through Azure Managed Identities rather than explicit credentials, eliminating the secret rotation risk associated with static service principal passwords.
- Privileged access to SAP systems - including SAP Basis administration, OS-level access, and database administration - should be routed through Azure Privileged Identity Management (PIM) for just-in-time elevation, with all privileged sessions recorded in Azure Monitor.
- Emergency (firefighter) access to SAP should be provisioned through SAP GRC Emergency Access Management with automated alert notification to the audit team; the firefighter account credentials should be stored in Azure Key Vault with access restricted by AAD PIM policies.

4.2 Network Security Architecture

Network security for SAP on Azure is implemented through a layered defence model. The outermost layer is provided by Azure Application Gateway with Web Application Firewall (WAF), which intercepts and inspects all inbound HTTPS traffic for SAP Fiori and ABAP ICM endpoints. The middle layer uses Azure Firewall for hub-to-spoke traffic enforcement, application-level FQDN filtering for outbound SAP traffic (software downloads, telemetry), and east-west traffic control between the hub network and SAP workload spokes. The innermost layer uses Network Security Groups (NSGs) to enforce microsegmentation between SAP tiers - application, database, integration, and management subnets each carry distinct NSG rule sets.

The network security design follows these core principles:

- No public IP addresses should be assigned to any SAP virtual machine. All administrative access should be routed through Azure Bastion, which provides browser-based RDP/SSH without requiring public IP exposure. This eliminates the entire class of attack vectors targeting exposed RDP (port 3389) and SSH (port 22) endpoints.
- The SAP HANA database subnet should be isolated behind an NSG that permits inbound TCP only from the SAP application server subnet on HANA SQL ports (30015, 30013), from the SAP administration subnet on HANA Studio ports, and from the Azure Load Balancer service tag for ILB health probes.
- All Azure PaaS service access (Azure Key Vault, Azure Storage for SAP backups, Azure Monitor ingestion endpoints) should be routed through Azure Private Endpoints, eliminating data exfiltration paths through the public internet and improving performance through private routing.
- ExpressRoute or Site-to-Site VPN should provide connectivity between the SAP Azure landing zone and on-premises source systems, SLD servers, and user workstations. Split-tunnelling on corporate VPN should route SAP Fiori HTTPS traffic directly through ExpressRoute rather than through the corporate internet breakout.
- DDoS Standard protection should be enabled on the hub virtual network to protect the Application Gateway public IP against volumetric attacks during critical business periods such as month-end close and year-end reporting.
- Outbound internet access from SAP VMs should be prohibited via NSG deny rules and Azure Firewall application rules, with explicit exceptions only for SAP software download URLs (softwaredownloads.sap.com), SAP support portal endpoints, and Azure update service tags.

4.3 Data Protection and Encryption Strategy

SAP HANA stores all data in memory for processing and persists data to disk through a combination of redo log files, delta files, and full data snapshots. Both in-memory and on-disk data require protection against unauthorised access and exfiltration. Azure provides encryption at rest through Azure Disk Encryption (ADE), which encrypts the OS and data disks of SAP VMs using BitLocker (Windows) or DM-Crypt (Linux) with keys stored in Azure Key Vault. For SAP HANA, an additional option is HANA's own Data Volume Encryption (DVE), which encrypts data at the HANA persistence layer before it reaches the block storage layer.

| e-ISSN: 2320-9801, p-ISSN: 2320-9798| www.ijircce.com | |Impact Factor: 7.542|

|| Volume 9, Issue 6, June 2021 ||

| DOI: 10.15680/IJIRCCE.2021.0906311|

The data protection strategy validated for SAP on Azure includes the following elements:

- Azure Disk Encryption with customer-managed keys (BYOK) should be applied to all SAP VM OS and data disks, with the encryption keys stored in Azure Key Vault HSM (Premium tier, FIPS 140-2 Level 2 validated hardware).
- Key Vault access policies should restrict encryption key access to only the specific VM service principal identities that require it; no human identities should have direct key access in production environments.
- HANA Data Volume Encryption (DVE) provides an additional layer of in-database encryption and should be evaluated for systems processing GDPR-regulated personal data, as it enables encryption at a finer granularity than disk-level encryption.
- All data in transit between SAP components - including RFC connections between application servers and HANA, HTTP/S traffic between Web Dispatcher and ICM, and replication traffic between HANA primary and secondary - should be encrypted using TLS 1.2 or higher with valid certificates issued by the enterprise PKI or Azure-managed certificates.
- SAP transport files and backup data written to Azure Blob Storage should be stored in accounts configured with customer-managed keys and with Blob Soft Delete enabled (minimum 30-day retention) to prevent accidental or malicious deletion of SAP backups.
- GDPR data masking requirements for SAP systems should be addressed through SAP Information Lifecycle Management (ILM) for data archival and SAP Data Custodian for data discovery and classification, integrated with the Azure data governance toolchain.

4.4 Threat Detection and Security Monitoring

The convergence of Azure Sentinel (Microsoft's cloud-native SIEM) with SAP's audit log infrastructure represents a significant capability advancement for SAP security operations teams in 2021. Historically, SAP audit logs were analysed within the SAP system itself using transaction SM20 or through specialised SAP GRC tools, limiting the ability to correlate SAP events with Azure infrastructure signals, identity events, and network anomalies. The Microsoft Sentinel solution for SAP (in public preview since early 2021) ingests SAP ABAP application logs, SAP Security Audit Logs, and SAP change document logs into Azure Sentinel, enabling cross-platform detection rules and investigation workflows.

- Microsoft Defender for Cloud (formerly Azure Security Center Standard) should be enabled across all Azure subscriptions hosting SAP workloads, providing continuous compliance assessment against benchmarks including CIS Azure Foundations and ISO 27001.
- Azure Sentinel alert rules should cover critical SAP security events including: emergency (firefighter) account activation in SAP GRC, authorisation failures exceeding threshold values, SoD violation approvals, RFC gateway security events, and OS-level privileged command execution on HANA VMs.
- Network traffic from SAP subnets should be analysed using Azure Network Watcher flow logs, with suspicious patterns (unexpected east-west traffic, large outbound data transfers from database subnet) triggering automated investigation playbooks in Sentinel.
- SAP system parameters governing the Security Audit Log (rsau/enable = 1, rsau/max_diskspace = sufficient capacity) must be configured and validated as part of the security baseline, as a missing or undersized audit log directly reduces the evidentiary value of SAP Sentinel analytics.
- Vulnerability management should be implemented through Microsoft Defender for Cloud's integrated Qualys scanner, with SAP VM vulnerability scan results tracked against CVSS scoring thresholds that align with the organisation's patch SLA commitments.

Security Pillar Key Insight: The single highest-impact security control for SAP on Azure is eliminating direct internet exposure of SAP VMs and routing all administrative access through Azure Bastion with AAD MFA. This single architectural decision eliminates a disproportionate share of the attack surface at negligible incremental cost.

V. PILLAR III - COST OPTIMISATION

Cost Optimisation in the Azure WAF is defined as delivering maximum business value at the lowest sustainable Azure spend. For enterprise SAP workloads, the cost optimisation journey begins at project inception - with the selection of appropriate VM families and sizing - and continues through steady-state operations, where Reserved Instance strategies, non-production scheduling, and right-sizing refinements drive continuous spend reduction.

A common misconception in SAP cloud programmes is that cost optimisation is in tension with reliability. In reality, several cost optimisation levers - such as selecting the correct VM size rather than over-provisioning, properly tiering storage, and shutting down development systems overnight - have no impact whatsoever on production availability

| e-ISSN: 2320-9801, p-ISSN: 2320-9798| www.ijircce.com | |Impact Factor: 7.542|

|| Volume 9, Issue 6, June 2021 ||

| DOI: 10.15680/IJIRCCE.2021.0906311|

SLAs. The tension between cost and reliability is real but confined to specific decisions, such as single-node vs. clustered ASCS or the selection between Availability Sets (within a datacenter, lower cost) and Availability Zones (cross-datacenter, potentially higher due to inter-zone data transfer charges).

5.1 Reserved Instances and Savings Plans for SAP

Azure Reserved VM Instances (RIs) are the primary mechanism for reducing compute costs for production SAP workloads. By committing to a one- or three-year usage term for specific VM sizes in a given Azure region, organisations receive discounts of 40–72% compared to pay-as-you-go rates, depending on the VM family and commitment duration. For SAP HANA VMs in the Mv2-series and M-series - which carry the highest per-hour cost of any Azure VM family - Reserved Instances represent the single most impactful cost optimisation action available.

The following considerations govern RI strategy for SAP workloads:

- Production SAP HANA VMs should be purchased as 3-year Reserved Instances only after the VM sizing has been validated against at least three months of actual production workload metrics. Committing to an RI before sizing is confirmed risks stranding spend on an oversized VM SKU.
- The RI scope should be set to "Shared" across the SAP landing zone subscription group to maximise flexibility; a shared scope allows the RI discount to apply to any qualifying VM in any subscription within the scope, enabling cost-free redeployment during system refresh or consolidation.
- Instance size flexibility (available for most non-HANA VM families) allows RI discounts to be applied across VM sizes within the same family and generation, providing a buffer against minor resizing adjustments without requiring RI exchange or refund transactions.
- Azure Hybrid Benefit for Windows Server or Red Hat Enterprise Linux can be layered on top of Reserved Instances to apply existing on-premises OS licences to Azure VMs, yielding additional cost reductions of 20–40% on OS licensing for organisations with active Software Assurance or Red Hat subscriptions.
- Non-production SAP systems (development, sandbox, training) that are accessed during business hours only should not be covered by Reserved Instances; instead, Azure Dev/Test pricing or auto-shutdown schedules via Azure Automation should be used to reduce their hourly cost.

5.2 Right-Sizing SAP Workloads on Azure

SAP workload right-sizing is a nuanced discipline because SAP HANA's in-memory architecture means that the database VM must be sized to hold the entire working set in RAM, not merely peak CPU utilisation as would be the primary sizing driver for most cloud workloads. SAP Quick Sizer provides SAPS (SAP Application Performance Standard) and memory estimates based on functional scoping inputs, but these estimates should be treated as starting points rather than definitive specifications: Quick Sizer typically provides conservative estimates, and actual resource consumption in production depends heavily on data volumes, custom code quality, and the number and nature of concurrently active users.

Right-sizing recommendations for SAP on Azure are structured as follows:

- For greenfield migrations, begin with SAP Quick Sizer output and apply a 20% overhead buffer for initial sizing; monitor actual CPU, memory, and disk I/O utilisation for the first 90 days of production and plan a right-sizing review at the 6-month mark.
- For brownfield migrations (lift-and-shift from on-premises), capture on-premises SAP performance data using SAP's EarlyWatch Alert reports and OS-level statistics tools (sar, iostat) for a minimum of 4 weeks covering a representative business cycle including month-end close.
- HANA database memory sizing must account for: current database footprint (as reported by HANA Studio administration perspective), growth projections for the RI commitment period, and the memory overhead of HSR secondary preloading (equivalent to the primary's column store memory when preload is enabled).
- Azure VM CPU right-sizing should be evaluated using Azure Monitor VM Insights, which provides 95th percentile CPU utilisation data; SAP application server VMs consistently running below 40% average CPU utilisation are candidates for downsizing to the next smaller SKU within the same family.
- Storage right-sizing is addressed in Section 7.2 under the Performance Efficiency pillar, as storage sizing for HANA is driven primarily by performance requirements (IOPS, throughput) rather than capacity for operational data.

| e-ISSN: 2320-9801, p-ISSN: 2320-9798| www.ijircce.com | |Impact Factor: 7.542|

|| Volume 9, Issue 6, June 2021 ||

| DOI: 10.15680/IJIRCCE.2021.0906311|

5.3 Storage Cost Optimisation

Azure storage costs for SAP workloads are material, particularly for systems with large HANA databases requiring multiple TiB of Premium SSD managed disks. Several storage cost optimisation levers are available without compromising the performance requirements mandated by SAP certification:

- SAP HANA log volumes - which require low latency writes but have modest capacity requirements - should use Write Accelerator-enabled Premium SSD (required for Mv2-series VMs) and be sized to the minimum certified capacity (typically 512 GiB per log volume stripe), not over-provisioned for throughput that is already capped by Write Accelerator.
- SAP HANA backup data should be tiered to Azure Blob Storage using Cool or Archive access tiers based on retention age, with Lifecycle Management policies automatically transitioning backups older than 30 days to Cool tier and older than 90 days to Archive tier, reducing storage costs by 50–90% compared to Hot tier Premium storage.
- Development and QA system disks that are not required outside working hours should use Standard SSD managed disks rather than Premium SSD, accepting higher latency in exchange for 60–70% reduction in disk cost. Standard SSD latency (2–10 ms) is acceptable for non-production SAP workloads where user experience is not a primary concern.
- Azure NetApp Files capacity pools should be sized using the minimum pool size that covers current volume requirements plus a growth buffer, as Azure NetApp Files charges are based on provisioned pool capacity regardless of actual volume consumption.

Cost Optimisation Key Insight: For most enterprise SAP programmes, the three-year Reserved Instance commitment for production HANA VMs is the single highest-return cost action available. It should be executed after 90 days of validated production sizing and generates discounts of 55–65% on HANA compute spend with zero reliability or performance trade-off.

VI. PILLAR IV - OPERATIONAL EXCELLENCE

Operational Excellence in the Azure WAF encompasses the processes, procedures, and tooling that enable teams to run and monitor systems effectively and to continuously improve operational practices. For enterprise SAP on Azure, Operational Excellence is both the most impactful and the most frequently under-invested pillar. On-premises SAP operations teams typically work with manually executed runbooks, ticket-driven change management, and reactive monitoring approaches. Cloud migration provides a forcing function to modernise these practices, but realising that opportunity requires deliberate investment in automation, observability, and organisational change management.

6.1 Infrastructure-as-Code for SAP Landscapes

Infrastructure-as-Code (IaC) for SAP on Azure means codifying the provisioning, configuration, and ongoing management of SAP infrastructure using declarative tools such as Azure Resource Manager (ARM) templates, Bicep, or Terraform. The SAP on Azure team at Microsoft has published a comprehensive open-source IaC framework - the SAP Deployment Automation Framework (SDAF), available on GitHub since early 2021 - that implements Terraform modules for SAP infrastructure provisioning and Ansible playbooks for SAP software installation and configuration.

The benefits of IaC for SAP landscapes extend well beyond the initial deployment:

- Repeatability and consistency: IaC ensures that development, QA, and production SAP systems are provisioned from the same codebase, eliminating configuration drift that can cause "works in DEV but fails in PROD" incidents during transport-driven testing.
- Disaster recovery rehearsal: DR environments that are codified in IaC can be spun up from scratch in minutes, enabling DR tests that actually destroy and recreate the DR environment - a significantly more rigorous validation than simply verifying that replicated VMs power on.
- Audit trail and governance: IaC changes are managed through Git pull requests and code review, providing a complete audit trail of infrastructure changes with reviewer accountability - a capability that is particularly valuable for SAP change management under SOX compliance.
- Rapid environment provisioning: New SAP system landscapes for training, user acceptance testing, or performance testing can be provisioned in hours rather than weeks, enabling agile SAP project methodologies that were previously impractical.
- Immutable infrastructure: Critical SAP configuration settings can be enforced through IaC that is applied on every deployment cycle, preventing manual configuration changes from accumulating into an undocumented, unsupportable state.

| e-ISSN: 2320-9801, p-ISSN: 2320-9798| www.ijircce.com | |Impact Factor: 7.542|

|| Volume 9, Issue 6, June 2021 ||

| DOI: 10.15680/IJIRCCE.2021.0906311|

6.2 SAP-Aware Monitoring and Observability

Effective SAP monitoring on Azure requires observability across four distinct data domains: Azure infrastructure metrics (CPU, memory, disk, network at the VM and Azure service level), SAP application metrics (work process utilisation, enqueue lock counts, ICM connection statistics), SAP HANA metrics (memory allocation, service utilisation, replication health, query performance), and end-user experience metrics (SAP Fiori page load times, RFC roundtrip latency).

The monitoring architecture recommended for SAP on Azure integrates multiple data sources into a unified Azure Monitor workspace:

- Azure Monitor with Log Analytics workspace serves as the central observability platform, receiving VM performance metrics, Azure service diagnostic logs, NSG flow logs, and application gateway access logs from all SAP infrastructure components.
- The Azure Monitor SAP solutions for HANA (generally available since April 2020) provides native integration between Azure Monitor and SAP HANA monitoring views, exposing HANA-level metrics - including System Replication health, HANA service availability, and disk volume usage - without requiring custom scripting or third-party agents.
- SAP CCMS (Computing Centre Management System) alert monitoring should be integrated with Azure Monitor through the SAP Solution Manager interface or through custom ABAP programs that forward CCMS alert status to Azure Monitor custom metrics API, enabling unified alerting across SAP application and Azure infrastructure dimensions.
- Azure Application Insights should be integrated with SAP Fiori applications to capture real-user monitoring (RUM) data, providing visibility into Fiori page load performance from the end-user perspective - a layer of observability not available from server-side HANA or ABAP metrics alone.
- Alert thresholds should be defined and documented for each monitoring signal, with severity classifications (P1 through P4) aligned to incident response SLAs. Alert fatigue - caused by miscalibrated thresholds generating false-positive notifications - is a significant operational risk in SAP monitoring environments and must be actively managed through regular threshold review.
- Azure Monitor Workbooks should be used to create consolidated operational dashboards providing at-a-glance system health views for SAP Basis teams, with drill-down capabilities into HANA replication status, application server work process availability, and background job completion rates.

6.3 Patching and Change Management Discipline

SAP landscape patching presents a unique operational challenge because SAP Kernel patches, SAP Support Package Stacks, and OS patches must be maintained in a coordinated combination to preserve SAP supportability. SAP Note 1928533 defines minimum kernel and OS versions for Azure-certified SAP configurations; deviating from these minimums by running outdated combinations can result in SAP Support refusing to provide assistance for incidents in an unsupported configuration.

The patching and change management discipline recommended for SAP on Azure aligns the Azure WAF Operational Excellence pillar with SAP's Change and Transport System (CTS) governance model:

- OS patching should be managed through Azure Update Management, which provides centralised visibility across all SAP VMs and allows patch deployment to be sequenced according to SAP system criticality and availability constraints - for example, patching development systems first, then QA, then production during a planned maintenance window.
- SAP HANA revision updates should follow a rolling update procedure: apply the revision to the secondary HANA node first, perform a controlled HSR takeover to make the updated node primary, then apply the revision to the former primary. This procedure eliminates HANA downtime during planned HANA updates.
- SAP Application Server updates (kernel patches, ABAP Support Package Stacks) should be applied through the SAP Software Provisioning Manager (SWPM), with update procedures documented as Azure Automation runbooks to reduce manual execution risk.
- All SAP infrastructure changes - including VM configuration changes, NSG rule modifications, network routing updates, and storage tiering adjustments - should be processed through a change advisory board (CAB) process using the IaC pull request mechanism as the change record, replacing manual change tickets with code-reviewed, auditable infrastructure modifications.
- Change freeze periods aligned with SAP financial period-close cycles should be enforced through Azure Policy assignments that block VM configuration changes and through GitOps branch protection rules that require additional approvals during blackout windows.

| e-ISSN: 2320-9801, p-ISSN: 2320-9798| www.ijircce.com | |Impact Factor: 7.542|

|| Volume 9, Issue 6, June 2021 ||

| DOI: 10.15680/IJIRCCE.2021.0906311|

Operational Excellence Key Insight: Investing in IaC and automated monitoring during the initial SAP migration programme - even when it adds to project duration - yields compounding operational dividends over the system's lifetime. Teams that migrate without IaC typically face the choice of either rebuilding the infrastructure from scratch to adopt it later or operating permanently in a manually managed state with growing technical debt.

VII. PILLAR V - PERFORMANCE EFFICIENCY

Performance Efficiency in the Azure WAF is concerned with the efficient use of computing resources to meet performance requirements, and the ability to maintain that efficiency as demand changes. For SAP HANA workloads, performance efficiency is heavily shaped by three resource dimensions: compute (CPU and NUMA architecture), storage (IOPS, throughput, and latency of HANA data and log volumes), and network (latency and bandwidth between HANA nodes, between HANA and application servers, and between the SAP landscape and end users). Each dimension has specific minimum requirements defined by SAP certification standards and must be validated through standardised benchmark tools before go-live.

7.1 Compute Selection and HANA Sizing Strategy

The selection of the appropriate Azure VM family for SAP HANA is constrained by the SAP certification matrix documented in SAP Note 1928533. As of early 2021, the certified VM families for SAP HANA production workloads on Azure are the M-series (M32ts through M128ms) and Mv2-series (M208s_v2 through M416ms_v2), alongside the E-series (E20s_v4 through E96s_v4) for HANA in smaller production and large non-production deployments. The critical differentiator between these families is memory capacity: HANA databases whose in-memory footprint exceeds 1.5 TiB require the Mv2-series; footprints above 4 TiB require the M416ms_v2 (11.4 TiB RAM) or a HANA scale-out topology.

The HANA compute sizing strategy is governed by the following principles:

- The primary HANA VM sizing driver is the database's in-memory working set size, defined as the sum of column store table sizes in memory (from the HANA M_CS_TABLES monitoring view) plus the row store size, plus HANA service process overhead. A 20% buffer above the measured working set should be allocated for growth and HANA internal operations.
- NUMA architecture awareness is critical for HANA performance: the M128s has two NUMA nodes, each with 64 vCPUs and 1 TiB RAM. SAP HANA automatically distributes services across NUMA nodes, but application workloads that generate cross-NUMA memory accesses will experience higher latency. This should be evaluated during performance testing.
- For SAP Central Services and Application Servers, the E-series (memory-optimised) VMs are preferred over the D-series (general purpose) because SAP ABAP work processes are memory-intensive; the 8:1 memory-to-vCPU ratio of E-series (compared to 4:1 for D-series) better matches the ABAP memory consumption profile.
- Accelerated Networking must be enabled on all SAP application VMs to achieve the low-latency network performance required for SAP RFC communication with HANA. Accelerated Networking bypasses the virtual switch and delivers near-physical network performance with sub-100 microsecond latency for intra-datacenter traffic.

7.2 Storage Performance for SAP HANA

SAP HANA imposes specific minimum storage performance requirements on its data and log volumes that are significantly more demanding than general-purpose workloads. These requirements are specified in SAP Note 2399079 (Storage requirements for SAP HANA on Azure) and must be met by the storage configuration to maintain HANA supportability. The key requirement is that the HANA log volume must sustain 400 MB/s write throughput at sub-1 ms latency to avoid transaction commit stalls during high-concurrency write workloads.

The storage configuration validated for SAP HANA on Azure follows these design principles:

- HANA log volumes on Mv2-series VMs must use Azure Write Accelerator-enabled Premium SSD managed disks. Write Accelerator provides a dedicated low-latency write path that bypasses the standard Premium SSD caching layer, achieving the sub-1 ms write latency required for HANA log commits. Write Accelerator is available only on M-series and Mv2-series VMs.
- HANA data volumes should use Premium SSD managed disks in a software RAID-0 stripe set (using Linux mdadm or LVM striping) across multiple disk LUNs to aggregate IOPS and throughput beyond the per-disk limits.

| e-ISSN: 2320-9801, p-ISSN: 2320-9798| www.ijircce.com | |Impact Factor: 7.542|

|| Volume 9, Issue 6, June 2021 ||

| DOI: 10.15680/IJIRCCE.2021.0906311|

A typical configuration for a 2 TiB HANA database uses four 1 TiB P70 disks striped together, providing 32,000 combined IOPS and 1,000 MB/s throughput.

- Azure NetApp Files (ANF) with the Ultra service level (128 MiB/s per TiB) is an increasingly viable alternative to Premium SSD RAID-0 for HANA data and log volumes as of 2021, offering a simplified storage architecture (no RAID configuration required), native NFS support for shared file system requirements (/sapmnt, transport directory), and elastic capacity scaling without VM-level disk operations.
- Read cache should be disabled on HANA data disk volumes (Azure host caching set to "None") as HANA manages its own buffer pool in memory; OS-level read caching doubles memory consumption without improving performance for column-store table loads that already benefit from HANA's native caching.
- HANA persistent memory (PMEM) support, available on Mv2-series VMs through Azure lsv2 NVMe SSDs, should be evaluated for workloads with large warm data sets that benefit from near-memory-speed persistence, though this feature requires specific HANA configuration parameters and should only be implemented after validation in non-production environments.

7.3 Network Performance Optimisation

Network performance is the most operationally invisible dimension of SAP performance on Azure: when network latency degrades, the effects manifest as application-level slowness (slow logon, delayed ALV report rendering, long background job runtimes) that is difficult to diagnose without dedicated network observability tooling. The following network performance optimisation principles are validated for SAP on Azure:

- All SAP VMs in the same tier should be placed in a Proximity Placement Group (PPG) combined with the Availability Set to ensure physical proximity of VMs within the same datacenter rack neighbourhood, minimising intra-tier network latency. PPG placement reduces intra-NUMA-zone latency from a typical 0.5–1.5 ms to sub-0.2 ms for pairs of VMs in the same group.
- The SAP HANA scale-up two-node HA pair should share a single Proximity Placement Group to ensure sub-1 ms HSR replication latency in synchronous mode. The PPG constraint takes precedence over Availability Set fault domain placement for HANA pairs where synchronous replication is the HA mechanism.
- SAP Web Dispatcher should be configured with SAP Logon Group balancing to distribute inbound user connections based on current application server work process load rather than simple round-robin, preventing hot-spot formation on individual application servers during uneven workload distribution.
- Accelerated Networking must be confirmed as active using the ethtool -S <interface> command to verify the VF (Virtual Function) driver is in use; mere presence of the accelerated-networking setting in the Azure portal is insufficient - the VM must also be a supported size, run a supported OS kernel, and have the correct NIC driver installed.
- ExpressRoute Premium or Standard circuit bandwidth should be sized for peak SAP replication and user traffic; a common calculation for SLT-based Central Finance replication is: (daily document volume × average document payload × replication frequency factor × safety margin) / seconds per day = required sustained throughput, validated against the circuit's committed burst capacity.

Performance Efficiency Key Insight: For SAP HANA on Azure Mv2-series VMs, enabling Write Accelerator for the HANA log volume is a zero-cost, no-risk configuration change that is mandatory for SAP supportability and provides measurable reduction in HANA commit latency. It is the first performance configuration action to validate after any HANA VM deployment.

VIII. CROSS-PILLAR TRADE-OFF ANALYSIS

The five WAF pillars are interdependent, and optimising for one pillar can create genuine tensions with others. Understanding these trade-offs enables architects to make explicit, documented decisions rather than inadvertently compromising one pillar in the service of another. This section analyses the most significant cross-pillar tensions encountered in enterprise SAP on Azure programmes.

8.1 Reliability vs. Cost Optimisation

The most prominent tension between Reliability and Cost Optimisation in SAP on Azure is the choice between Availability Sets (single datacenter, no inter-zone data transfer cost, sub-1 ms synchronous HANA replication) and Availability Zones (cross-datacenter, 2–4 ms replication latency, inter-zone data transfer charges of approximately $0.01 per GiB). For synchronous HANA System Replication - required for RPO=0 HA - the 2–4 ms inter-zone latency can increase COMMIT latency by the same amount under high write loads, directly degrading application responsiveness for latency-sensitive transactions.

| e-ISSN: 2320-9801, p-ISSN: 2320-9798| www.ijircce.com | |Impact Factor: 7.542|

|| Volume 9, Issue 6, June 2021 ||

| DOI: 10.15680/IJIRCCE.2021.0906311|

The recommended resolution is to use Availability Sets for the HA cluster within the primary region (prioritising Reliability without the cost penalty of inter-zone transfer) and Availability Zones for the DR region deployment (where asynchronous replication tolerates higher latency). This hybrid approach is the pattern documented by SAP and Microsoft as the preferred architecture for HANA HA+DR in 2021.

8.2 Security vs. Operational Excellence

Strict security controls can introduce operational friction that, if unmanaged, leads teams to develop workarounds that undermine the security controls they were designed to enforce. Two specific examples are relevant to SAP on Azure:

- Azure Bastion with mandatory MFA provides excellent security for administrative access but adds 15–30 seconds of authentication overhead per session compared to direct SSH from a VPN-connected workstation. If this overhead causes operators to seek less secure alternatives, the net security outcome is negative. The mitigation is to optimise the Bastion user experience through AAD seamless SSO integration and to use Azure Bastion's file transfer capabilities to eliminate the need for separate SFTP sessions.
- Network Security Group rules restricting HANA connectivity improve security posture but must be maintained in synchronisation with SAP landscape changes. If new SAP application servers are added to the landscape without corresponding NSG updates, connectivity failures occur that can delay project timelines and incentivise "allow-all" rule workarounds. The mitigation is to codify NSG rules in IaC (resolving the Operational Excellence pillar simultaneously) so that NSG updates are co-deployed with VM provisioning changes through the same pull request.

8.3 Performance Efficiency vs. Cost Optimisation

Premium SSD managed disks - required for HANA data and log volumes - are billed on provisioned capacity and performance tier regardless of utilisation, creating a cost tension: HANA requires large disk capacity for data growth but the disks are often significantly under-utilised in terms of IOPS and throughput. The resolution is to use the minimum number of Premium SSD disks that provide the required aggregate IOPS and throughput (rather than provisioning one large disk per volume), which also minimises cost for a given performance target. Azure's Online Performance Tier Change capability (generally available since 2020) allows Premium SSD disks to be upgraded to a higher performance tier during peak periods (month-end close) and downgraded afterward, providing a dynamic performance-cost balance without VM restart.

8.4 Operational Excellence vs. Reliability

Automated operations - while improving consistency and reducing manual error risk - introduce their own failure modes if automation scripts are misconfigured or triggered inappropriately. The specific risk in SAP on Azure is that automated patching or VM management scripts could trigger during a HANA System Replication resync or during a Pacemaker cluster reconfiguration, disrupting the HA cluster state. The mitigation is to implement pre-action checks in all automation runbooks that verify HANA HSR synchronisation status, Pacemaker cluster health, and absence of active SAP batch processing before executing any maintenance action. These checks should be automated as guardrails in Azure Automation runbooks, not left as manual verification steps.

IX. VALIDATED DESIGN PATTERN: SYNTHESIS

The validated design pattern for applying the Azure Well-Architected Framework to enterprise SAP workloads synthesises the pillar-specific guidance from Sections 3 through 7 into a coherent architectural position. The pattern is presented as a set of non-negotiable foundation requirements and a set of recommended enhancements that improve WAF alignment without constituting blocking requirements for go-live.

9.1 Foundation Requirements

The following design decisions constitute the minimum viable architecture for a WAF-aligned SAP on Azure deployment. Absence of any of these elements represents a material gap in one or more WAF pillars:

- **Reliability:** ASCS/ERS Pacemaker cluster with Azure Fence Agent STONITH, both nodes in separate fault domains of an Availability Set, cluster VIP managed via Azure Internal Load Balancer with Standard SKU and floating IP.
- **Reliability:** SAP HANA HSR in synchronous mode with Pacemaker SAPHanaSR resource agent for automated failover, both nodes in separate fault domains of an Availability Set.
- **Reliability:** Documented and tested DR procedure with explicit RPO and RTO targets signed off by the business owner before production go-live.

| e-ISSN: 2320-9801, p-ISSN: 2320-9798| www.ijircce.com | |Impact Factor: 7.542|

|| Volume 9, Issue 6, June 2021 ||

| DOI: 10.15680/IJIRCCE.2021.0906311|

- **Security:** No public IP addresses on any SAP VM; all administrative access through Azure Bastion with AAD MFA.
- **Security:** NSG microsegmentation enforced between all SAP tiers; HANA subnet restricted to application-subnet-sourced traffic only.
- **Security:** Azure Disk Encryption with customer-managed keys applied to all SAP VM disks; keys stored in Azure Key Vault Premium (HSM-backed).
- **Security:** MFA enforced for all SAP Fiori access through AAD Conditional Access integrated with SAP IAS.
- **Cost Optimisation:** Non-production SAP systems configured with auto-shutdown/auto-start schedules to eliminate overnight and weekend compute costs.
- **Operational Excellence:** SAP infrastructure provisioned through IaC (ARM, Bicep, or Terraform); no manual VM provisioning in production.
- **Operational Excellence:** Azure Monitor with Log Analytics workspace operational, SAP HANA monitoring views integrated, and P1 alerts defined and tested before go-live.
- **Performance Efficiency:** Accelerated Networking enabled on all SAP VMs; Write Accelerator enabled on HANA log volume disks on M-series/Mv2-series; HANA storage configuration validated against SAP Note 2399079.

9.2 Recommended Enhancements

The following enhancements are recommended for programmes with the capacity to implement them, either at initial go-live or as a post-go-live improvement roadmap:

- Replace HANA RAID-0 Premium SSD configuration with Azure NetApp Files Ultra tier for HANA data and log volumes, simplifying storage management and providing native NFS for shared file system requirements.
- Deploy the Microsoft Sentinel SAP connector to integrate SAP ABAP audit logs with the Azure SIEM for cross-platform threat detection and investigation.
- Implement SAP HANA multi-tier HSR (synchronous HA + asynchronous DR) to consolidate HA and DR replication into a single HSR mechanism, eliminating the need for Azure Site Recovery for the HANA tier.
- Adopt the SAP Deployment Automation Framework for all landscape provisioning to standardise IaC implementation and reduce the infrastructure-as-code development burden on project teams.
- Purchase 3-year Reserved Instances for production HANA and application server VMs after 90-day sizing validation to lock in optimal per-unit economics for the VM commitment period.
- Implement Azure Policy assignments to enforce WAF-aligned configuration standards (e.g., "all SAP VMs must have disk encryption enabled," "no public IPs permitted in SAP subscriptions") at the governance layer rather than relying on manual compliance reviews.
- Integrate SAP EarlyWatch Alert data into Azure Monitor dashboards to combine SAP application-layer health signals with Azure infrastructure metrics in a unified operational view.
- Implement Proximity Placement Groups for all SAP VM clusters to ensure the lowest achievable intra-tier network latency and validate latency improvements against baselines measured without PPG constraints.

Design Pattern Summary: A fully WAF-aligned SAP on Azure architecture is not a one-time deliverable but an ongoing state to pursue through structured architecture reviews and continuous improvement cycles. Microsoft recommends conducting a formal WAF assessment using the Azure Architecture Review tool at go-live, at 6 months, and annually thereafter to identify evolving gaps as the SAP landscape evolves and Azure services advance.

X. CONCLUSION AND FUTURE RESEARCH DIRECTIONS

This article has presented a comprehensive, validated design pattern for applying the Azure Well-Architected Framework to enterprise SAP workloads, covering all five WAF pillars in the context of SAP HANA-based deployments on Azure Infrastructure-as-a-Service. The central argument is that the WAF's generic principles, while directionally correct, require significant SAP-specific adaptation to translate into actionable architectural guidance. The gaps between generic Azure best practice and SAP-specific best practice are not merely cosmetic: they represent real architectural risks - ASCS Pacemaker misconfiguration, inadequate HANA storage performance, overly permissive NSG rules, absence of Write Accelerator on HANA log volumes - that have caused production incidents in real-world SAP on Azure deployments.’

The Reliability pillar analysis demonstrates that SAP's high availability requirements - Pacemaker STONITH clustering for ASCS/ERS, synchronous HANA System Replication with automatic takeover, and Availability Set-based fault domain isolation - map cleanly to Azure constructs but require precise configuration of Azure Internal Load Balancer, NSG health probe ports, and Pacemaker resource constraints that are not intuitive from the WAF documentation alone.

| e-ISSN: 2320-9801, p-ISSN: 2320-9798| www.ijircce.com | |Impact Factor: 7.542|

|| Volume 9, Issue 6, June 2021 ||

| DOI: 10.15680/IJIRCCE.2021.0906311|

The Security pillar analysis highlights the transformative opportunity represented by the Azure Sentinel SAP connector and AAD-based identity federation, which together enable a security operations capability for SAP that was structurally impossible in on-premises deployments.

The Operational Excellence pillar analysis makes the case that Infrastructure-as-Code and Azure Monitor SAP integration are not optional refinements but foundational investments whose absence creates compounding operational debt. Organisations that treat IaC as a post-migration improvement project consistently find the transition much harder to execute than those that enforce IaC discipline from the outset of the migration progamme.

The cross-pillar trade-off analysis in Section 8 provides a structured framework for navigating the most common architectural tensions - Reliability vs. Cost in the Availability Set vs. Availability Zone decision, Security vs. Operational Excellence in the Bastion and NSG discipline, and Performance Efficiency vs. Cost in the HANA storage configuration - with explicit, documented resolution approaches that prevent these tensions from becoming decision-paralysing debates in architecture review boards.

10.1 Directions for Future Research

Several emerging developments in the SAP and Azure ecosystems create directions for future research that extend beyond the scope of this article:

- The RISE with SAP programme: Microsoft and SAP's joint announcement of RISE with SAP in January 2021 creates a new deployment model - managed private cloud infrastructure operated by SAP - that shifts infrastructure responsibility boundaries in ways that require re-evaluation of all five WAF pillars from the customer perspective.
- SAP on Azure Kubernetes Service: The emerging containerised deployment model for SAP components, including SAP SolMan, SAP BTP Integration Suite, and experimental ABAP Platform container images, will require a WAF analysis that incorporates Kubernetes-specific reliability, security, and operational excellence considerations alongside traditional VM-based patterns.
- Azure Availability Zones for SAP HANA synchronous replication: As Azure continues to reduce inter-zone network latency - sub-2 ms between zones has been reported in some Azure regions in 2021 - the economic and availability trade-off between Availability Sets and Availability Zones for HANA synchronous replication merits longitudinal empirical investigation.
- FinOps maturity for SAP cloud spend: The Cost Optimisation pillar analysis in this article covers engineering-level levers; a complementary research direction is the organisational and process dimensions of FinOps maturity in SAP programmes, including chargeback model design, cloud budget governance, and the integration of Azure Cost Management with SAP project accounting.
- AI-driven SAP performance management: The integration of Azure Machine Learning with SAP HANA performance monitoring - using anomaly detection models trained on HANA monitoring view time series - represents an emerging Operational Excellence capability that could reduce the manual analysis burden on SAP Basis teams for performance root-cause investigations.

As the enterprise SAP community accelerates cloud adoption in response to SAP's ECC end-of-mainstream-maintenance timeline and the compelling economics of Azure-hosted S/4HANA, the need for rigorously validated, SAP-specific cloud architecture guidance will only intensify. The authors hope that this article contributes to that body of knowledge and provides practitioners with a structured, authoritative reference for their WAF-alignment journeys.

REFERENCES

1) Microsoft Corporation. (2020). Azure Well-Architected Framework. Microsoft Docs. First published October 2020. Retrieved April 2021 from https://docs.microsoft.com/azure/architecture/framework.
2) SAP SE. (2020). SAP Note 1928533 - SAP Applications on Microsoft Azure: Supported Products and Azure VM Types. Walldorf: SAP SE. Last updated November 2020.
3) SAP SE. (2020). SAP Note 2399079 - Elimination of Storage I/O Configuration for SAP HANA on Azure. Walldorf: SAP SE. Last updated September 2020.
4) SAP SE. (2019). SAP Note 2015553 - SAP on Microsoft Azure: Support Prerequisites. Walldorf: SAP SE. Last updated December 2019.
5) Microsoft Corporation. (2021). SAP HANA Infrastructure Configurations and Operations on Azure. Microsoft Docs - Azure Virtual Machines Planning Guide. Updated March 2021.

| e-ISSN: 2320-9801, p-ISSN: 2320-9798| www.ijircce.com | |Impact Factor: 7.542|

|| Volume 9, Issue 6, June 2021 ||

| DOI: 10.15680/IJIRCCE.2021.0906311|

6) Microsoft Corporation. (2021). High Availability of SAP HANA on Azure VMs on SUSE Linux Enterprise Server. Microsoft Docs. Updated January 2021.
7) Fowler, D., & Bennett, C. (2020). "Azure Availability Sets vs. Availability Zones for Latency-Sensitive Enterprise Workloads." IEEE Transactions on Cloud Computing, 8(3), pp. 791–806. DOI: 10.1109/TCC.2020.2981432.
8) Patel, S., & Sundaram, V. (2020). "SAP Central Finance as a Digital Financial Hub: Longitudinal Case Analysis." International Journal of Accounting Information Systems, 38(1), pp. 1–24. DOI: 10.1016/j.accinf.2020.100469.
9) Katz, R., Lindemann, T., & Osei, P. (2019). "Cloud Migration Velocity for SAP Workloads: A Multi-Industry Study." Journal of Enterprise Information Systems, 13(4), pp. 412–438. DOI: 10.1080/17517575.2019.1612104.
10) Wormald, P., & Choi, M. (2019). "Enterprise Networking for SAP on Azure: ExpressRoute, Virtual WAN, and Global VNet Peering." Azure Architecture Centre Technical Paper. Microsoft Corporation, November 2019.
11) Nakamura, T., Park, J., & Bauer, S. (2020). "SAP HANA System Replication on Microsoft Azure: Failover Characterisation and RTO Reduction Strategies." Proceedings of VLDB Endowment, 13(12), pp. 2891–2904.
12) SUSE LLC. (2020). SUSE Linux Enterprise Server for SAP Applications 15 SP2 - Setting up a Pacemaker Cluster. SUSE Best Practices. Document version 1.0, October 2020.
13) Red Hat, Inc. (2020). Configuring SAP HANA Scale-Up System Replication in Pacemaker on Microsoft Azure. Red Hat Customer Portal, Solution 5900841. Updated December 2020.
14) Microsoft Corporation. (2020). SAP on Azure Deployment Automation Framework. GitHub Repository: Azure/sap-automation. Initial public release, December 2020.
15) Gartner Research. (2020). "Magic Quadrant for Cloud ERP for Product-Centric Enterprises." Gartner ID G00454011. Published October 2020.
16) IDC. (2021). "Accelerating SAP Cloud Migration: Market Landscape and Vendor Assessment 2021." IDC White Paper #US47124521, sponsored by Microsoft Corporation. Published February 2021.
17) Microsoft Corporation. (2020). "Azure Sentinel: New SAP Continuous Threat Monitoring Solution." Microsoft Security Blog. Published November 2020. Retrieved March 2021.
18) SAP SE. (2020). "SAP on Microsoft Azure: Joint Reference Architecture." SAP Help Portal, Architecture Centre. Updated August 2020.
19) Microsoft Corporation. (2019). "Azure Disk Encryption for Linux VMs." Microsoft Docs. Updated December 2019. https://docs.microsoft.com/azure/virtual-machines/linux/disk-encryption-overview.
20) Müller, H., & Rao, K. (2020). "Zero Trust Security Architecture for Enterprise SAP Environments: An Azure Implementation Study." Journal of Information Security and Applications, 55, Article 102639. DOI: 10.1016/j.jisa.2020.102639.
21) SAP SE. (2021). "RISE with SAP: Business Transformation as a Service." SAP Press Release, January 2021. Retrieved February 2021 from sap.com.
22) Microsoft Corporation. (2021). "Azure Well-Architected Review: SAP Workload Lens." Microsoft Docs, Architecture Review Tool. Updated April 2021.
23) Amazon Web Services. (2020). "AWS Well-Architected Framework." AWS Whitepapers & Guides. Sixth Edition, July 2020. Retrieved January 2021 for comparative framework analysis.
24) Clarke, P., & Harrington, M. (2020). "Infrastructure-as-Code Maturity in Enterprise SAP Programmes: A Survey of Adoption Barriers." European Journal of Information Systems, 29(4), pp. 371–389. DOI: 10.1080/0960085X.2020.1730463.
25) SAP SE. (2019). SAP Note 2235581 - SAP HANA: Supported Operating Systems. Walldorf: SAP SE. Last updated May 2019.

Appendix A - Glossary of Key Terms

ASCS: ABAP Central Services instance. Hosts the SAP Message Server and Enqueue Server (lock manager). A singleton component whose unavailability renders the SAP system inaccessible.

ERS: Enqueue Replication Server. Maintains a replicated copy of the enqueue (lock) table on the standby Pacemaker cluster node, enabling zero-lock-loss failover of the ASCS.

ENSA2: Standalone Enqueue Server 2. The modern implementation of the SAP enqueue lock service, introduced in SAP Kernel 7.53. Replaces ENSA1 with a simplified cluster resource topology.

HSR: HANA System Replication. SAP's native database replication mechanism, supporting synchronous (SYNC) and asynchronous (ASYNC) modes for HA and DR respectively.

IaC: Infrastructure-as-Code. The practice of managing and provisioning cloud infrastructure through machine-readable configuration files (Terraform, Bicep, ARM) rather than manual processes.

ILB: Internal Load Balancer. Azure's Layer 4 load balancer for private (non-internet-facing) workloads. Used for SAP cluster virtual IPs and SAP Web Dispatcher backend pools.

| e-ISSN: 2320-9801, p-ISSN: 2320-9798| www.ijircce.com | |Impact Factor: 7.542|

|| Volume 9, Issue 6, June 2021 ||

| DOI: 10.15680/IJIRCCE.2021.0906311|

PPG: Proximity Placement Group. An Azure construct that co-locates VMs in physical proximity within a datacenter, minimising intra-group network latency. Used for HANA HA pairs.

RPO: Recovery Point Objective. The maximum acceptable data loss in the event of a system failure, expressed as a time duration (e.g., RPO=0 means no committed transactions can be lost).

RTO: Recovery Time Objective. The maximum acceptable duration between a failure event and the restoration of normal service operations.

SAPS: SAP Application Performance Standard. SAP's proprietary benchmark unit for measuring the compute capacity of systems running SAP workloads.

SDAF: SAP Deployment Automation Framework. Microsoft's open-source IaC framework for provisioning SAP on Azure infrastructure using Terraform and Ansible.

STONITH: Shoot The Other Node In The Head. A Pacemaker cluster fencing mechanism that forcibly powers off a potentially failed cluster node to prevent data corruption from split-brain scenarios.

WAF: Well-Architected Framework. Microsoft's structured framework for evaluating Azure workload quality across five pillars: Reliability, Security, Cost Optimisation, Operational Excellence, and Performance Efficiency.

Write Accel.: Azure Write Accelerator. A capability available exclusively on M-series and Mv2-series VMs that provides a dedicated low-latency write cache path for Premium SSD disks, required for SAP HANA log volumes.

Zero Trust: A security model that assumes no implicit trust based on network location, requiring explicit authentication and authorisation for every access request regardless of origin.