



CONA optimizes SAP landscape on Azure by moving SAP BW on HANA to Azure Virtual Machines

A Microsoft technical customer story

February 2021

When CONA Services LLC moved its SAP landscape to a modern cloud infrastructure on Azure almost two years ago, the migration represented the largest of its kind at the time—and a leap forward for the company. As early adopters of then-new SAP options on Azure, CONA trusted Microsoft with its cloud strategy and knew its plans would evolve as Azure evolved. The good news is that the move brought big cost savings almost immediately for the company, which serves the 12 largest Coca-Cola North American bottling partners. Now they have the SAP solution they need to collaborate as one company.

But this story begins with what happened next, when CONA moved SAP Business Warehouse (BW) on SAP HANA off of physical servers and onto high-throughput Azure M-series v2 (Mv2-series) virtual machines (VMs). The migration improved application performance and reduced total cost of ownership by an additional 40 percent, while giving CONA the flexibility and scalability of a virtual infrastructure on Azure.

"CONA realized significantly improved query performance, flexibility, and cost optimization following migration of our SAP BW on HANA environment to Azure Virtual Machines scale-out cluster with Azure NetApp Files."

*—Uday Reddy, Director – Cloud Engineering,
CONA Services LLC*

A company committed to innovation

CONA is the IT services organization that supports North America's largest beverage bottlers and producers of Coca-Cola products. To do this, the company runs a 48-terabyte (TB) deployment of SAP BW on HANA—the world's largest scale-out implementation on Azure. The bottling partners use the SAP solution to run their schedules and supply chains and to generate daily batch reports on trades and pricing. Since the move to Azure, the bottlers can now gather vital business intelligence in near real time based on queries and analytics from the data stored in the SAP HANA database.

For many enterprises, the journey to Azure starts with a lift-and-shift migration. This was CONA's approach in 2019 when it moved nine SAP landscapes from the company's on-premises datacenter and another cloud service provider to Azure. Everything—from the sandbox to the production servers to the disaster recovery site—was moved to Azure. The goal was to migrate as much as possible as quickly as possible using the best available options.

CONA chose to work with Microsoft, the only hyperscale vendor at the time, to provide an SAP-certified environment known as the Tailored Datacenter Integration (TDI). SAP and Microsoft have a long history of partnership. SAP named Azure the preferred cloud platform of [Embrace](#), a joint commitment between Microsoft and SAP to simplify and modernize the cloud journey for their customers. This partnership also offers unified technical support for the two platforms, making it easy for customers to get the help they need.

The result was a successful SAP migration by all accounts. It gave the bottling partners the platform they needed to improve operations, and it provided CONA the cost value it was looking for.

Before: A physical SAP HANA environment

The lift-and-shift started with SAP BW for HANA, a full-featured data warehouse system optimized for online transaction processing (OLTP). At the time, it was a 20+ TB database with more than 22,000 users.

The target was a unique hardware offering that Azure provides for running SAP HANA databases. Called SAP HANA on Azure (Large Instances) (also known as HANA Large Instances, or HLI), these special-purpose physical servers meet TDI standards and, at the time, were the only option big enough to handle CONA's migration. One HANA Large Instances stamp can support up to a 24-TB single instance of memory for SAP HANA workloads (and scale out up to 120 TB).

In the database tier of the "before" architecture on Azure (Figure 1), SAP HANA is deployed in a high-availability configuration that supports disaster recovery. The primary

region (Azure US East) is backed by a secondary region (Azure US West) that runs the production environment. Azure ExpressRoute gateways connect the regions.

Disaster recovery replication is provided through SAP HANA system replication. The traffic destined for the secondary site is routed through Azure ExpressRoute circuits, which create private connections between Microsoft datacenters and infrastructure to avoid the public internet.

“When we did this, it was the biggest SAP HANA instance of its kind ever to be migrated to Azure,” says Uday Reddy, Director – Cloud Engineering, CONA Services LLC. “We partnered with Microsoft knowing that our cloud strategy would evolve as Azure offerings for SAP also evolved.”

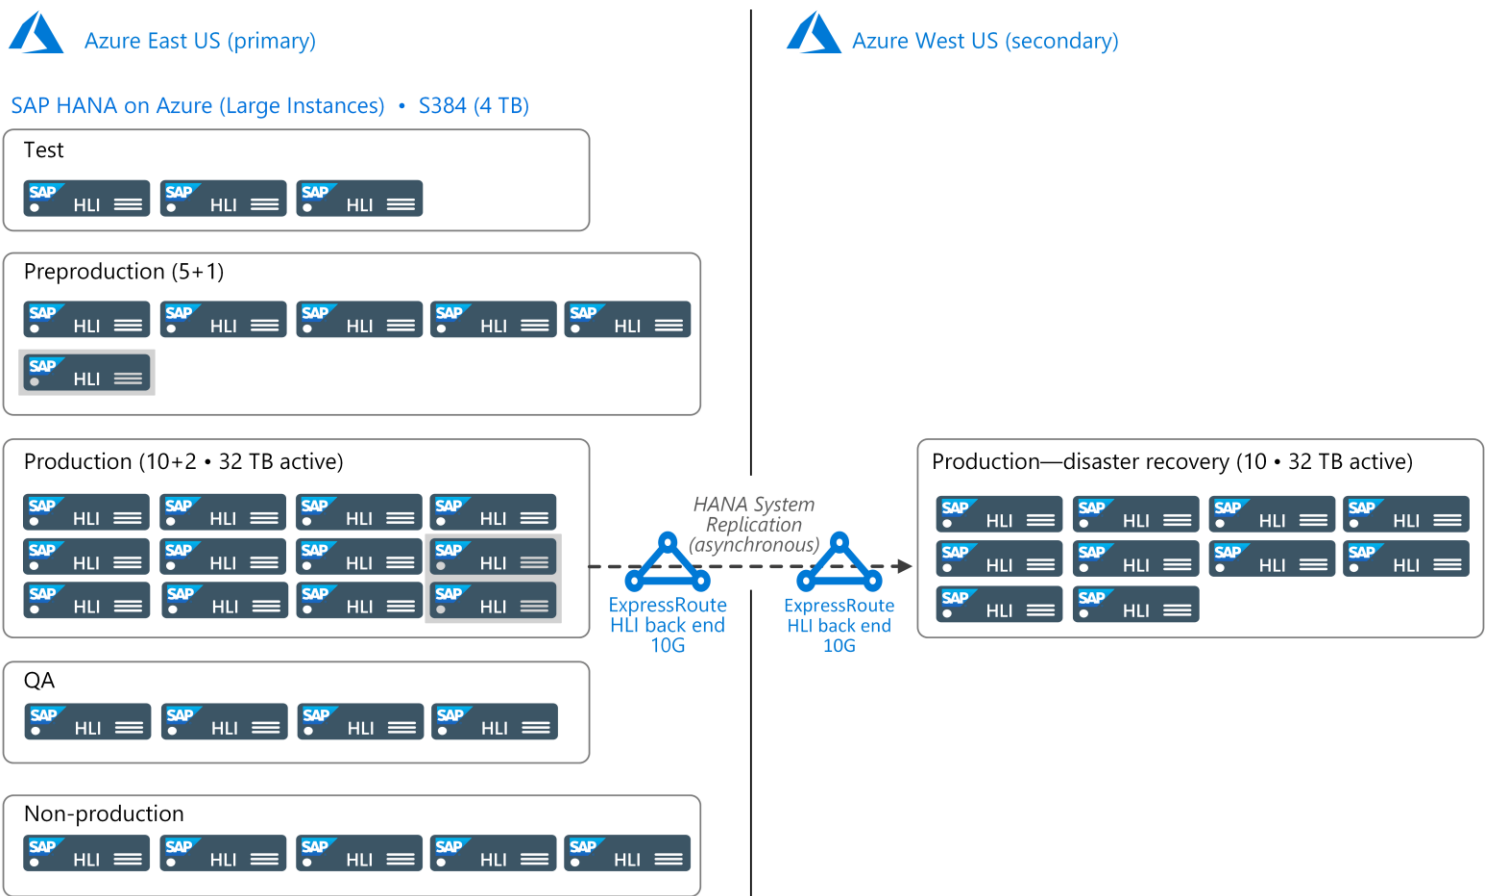


Figure 1. Before: Close-up of the highly available database tier before the move to Mv2-series VMs. The infrastructure uses physical servers—SAP HANA on Azure (Large Instances)—in the primary and secondary regions.

📘 When to use HANA Large Instances

Azure makes it possible to run and deploy SAP HANA on bare-metal servers that are dedicated to your organization. HANA Large Instances are physical servers based on the Intel Broadwell and Cascade Lake CPU architecture and configured in a large instance stamp. Within a large instance stamp, servers are not shared—they are dedicated to running one customer's deployment of SAP HANA.

These bare-metal server units are supported to run SAP HANA only. The SAP application layer, or workload middleware, runs in VMs. A variety of units (SKUs) are available, supporting up to a 24-TB single instance (120 TB scale-out) of memory for SAP BW/4HANA or other SAP HANA workloads. Capabilities and deployment requirements vary by type, and availability varies by region.

After: Virtual compute power for SAP HANA

For approximately two years, CONA ran in the physical HANA Large Instances environment. However, the company's data footprint continued to grow as the business grew, and the task of manually managing such a large physical infrastructure became time-consuming.

"As the business expanded, we had so much data growth that we needed to increase the number of nodes every three months," Reddy explains. "In a physical infrastructure, that takes time. We wanted to go to virtual to increase the speed of scaling up and down."

When Microsoft announced new TDI-certified VMs with up to 6 TB of capacity, CONA was ready to embrace the new. By using memory-optimized Mv2-series VMs, CONA could achieve even higher scalability and gain more flexibility for the company's large datasets—without sacrificing its highly available scale-out architecture.

"Our strategy was simple—running fewer servers would mean lower runtime costs for our big SAP investment," Reddy says.

Mv2-series VMs offer the highest vCPU count (up to 416 vCPUs) and largest memory of any VM in the cloud. With 3 TB, 6 TB, and 12 TB options, the Mv2-series is a high-throughput, low-latency platform that runs on a hyper-threaded Intel Xeon Platinum 8180M 2.5 GHz (Skylake) processor. Skylake provides an all-core base frequency of 2.5 GHz and a maximum turbo frequency of 3.8 GHz. Mv2-series VMs are also certified by SAP for SAP HANA OLTP and online analytical processing (OLAP) production workloads.

Teams from CONA, Microsoft, and the bottlers rehearsed the cutover to make sure they discovered any issues before the go-live. "Part of that was using mocks—that is, test runs in production hardware," notes Shibli Subhani, an Azure Principal Program Manager who worked with CONA. "That's how we made sure we were prepared and weren't going to find problems after the migration."

Replication of the 48-TB database began on a Monday. The cutover started on Friday night, and by Saturday afternoon the migration was complete. “We wanted this to be so seamless that the business didn’t notice the migration,” says Reddy. “And we did it. Microsoft was really helpful from day one.”

A virtual architecture on Azure

CONA successfully switched its SAP BW database from HANA Large Instances to 48 TB of highly available 6-TB VMs with two nodes used for hot standby in the event of failure. The new, more resilient configuration uses a smaller number of larger VMs to run SAP HANA.

The VMs’ Skylake processors give CONA faster performance than the Haswell processors used by the HANA Large Instances servers. “We call these eight-socket machines *beast nodes*. They helped with some of the performance gains in terms of the query processing and data nodes,” explains Subhani.

The move to Mv2-series VMs also made it possible to greatly simplify the architecture using Azure NetApp Files, an Azure-native, high-performance file storage service. Not available at the time of the initial lift-and-shift, Azure NetApp Files was the perfect solution for the large SAP HANA data and log files. To optimize costs while meeting performance requirements, this architecture uses 90 TB of Azure NetApp Files capacity spread across all three available service-level tiers—Standard, Premium, and Ultra.

The team also saved on costs by moving the preproduction cluster. In the “before” architecture, it ran in the Azure US East region with the other CONA environments. Now it runs in the Azure US West 2 region with the secondary production cluster. “CONA reduced resource consumption and optimized costs with this disaster recovery solution,” notes Subhani.

Choosing the right VM for your SAP landscape

The CONA architecture reflects best practices for running SAP HANA in a high-availability, scale-up environment that supports disaster recovery on Azure. Microsoft and SAP jointly certify a range of VM sizes for SAP HANA workloads. For example, smaller deployments can run on an M-series VM with 192 GiB of memory. To support the largest SAP HANA memory sizes on VMs—up to 11.4 TiB of usable memory—you can use the Azure Mv2-series VMs. The M208 virtual machine types achieve approximately 260,000 SAPS, and the M416 virtual machine types achieve approximately 488,000 SAPS.

For details about SAP support for Azure VM types and throughput metrics (SAPS), see [SAP Note 1928533](#). (To access the SAP notes, you must have an SAP Service Marketplace account.) The [SAP Certified and Supported SAP HANA Hardware Directory](#) has a list of certified Azure VMs for the HANA database.

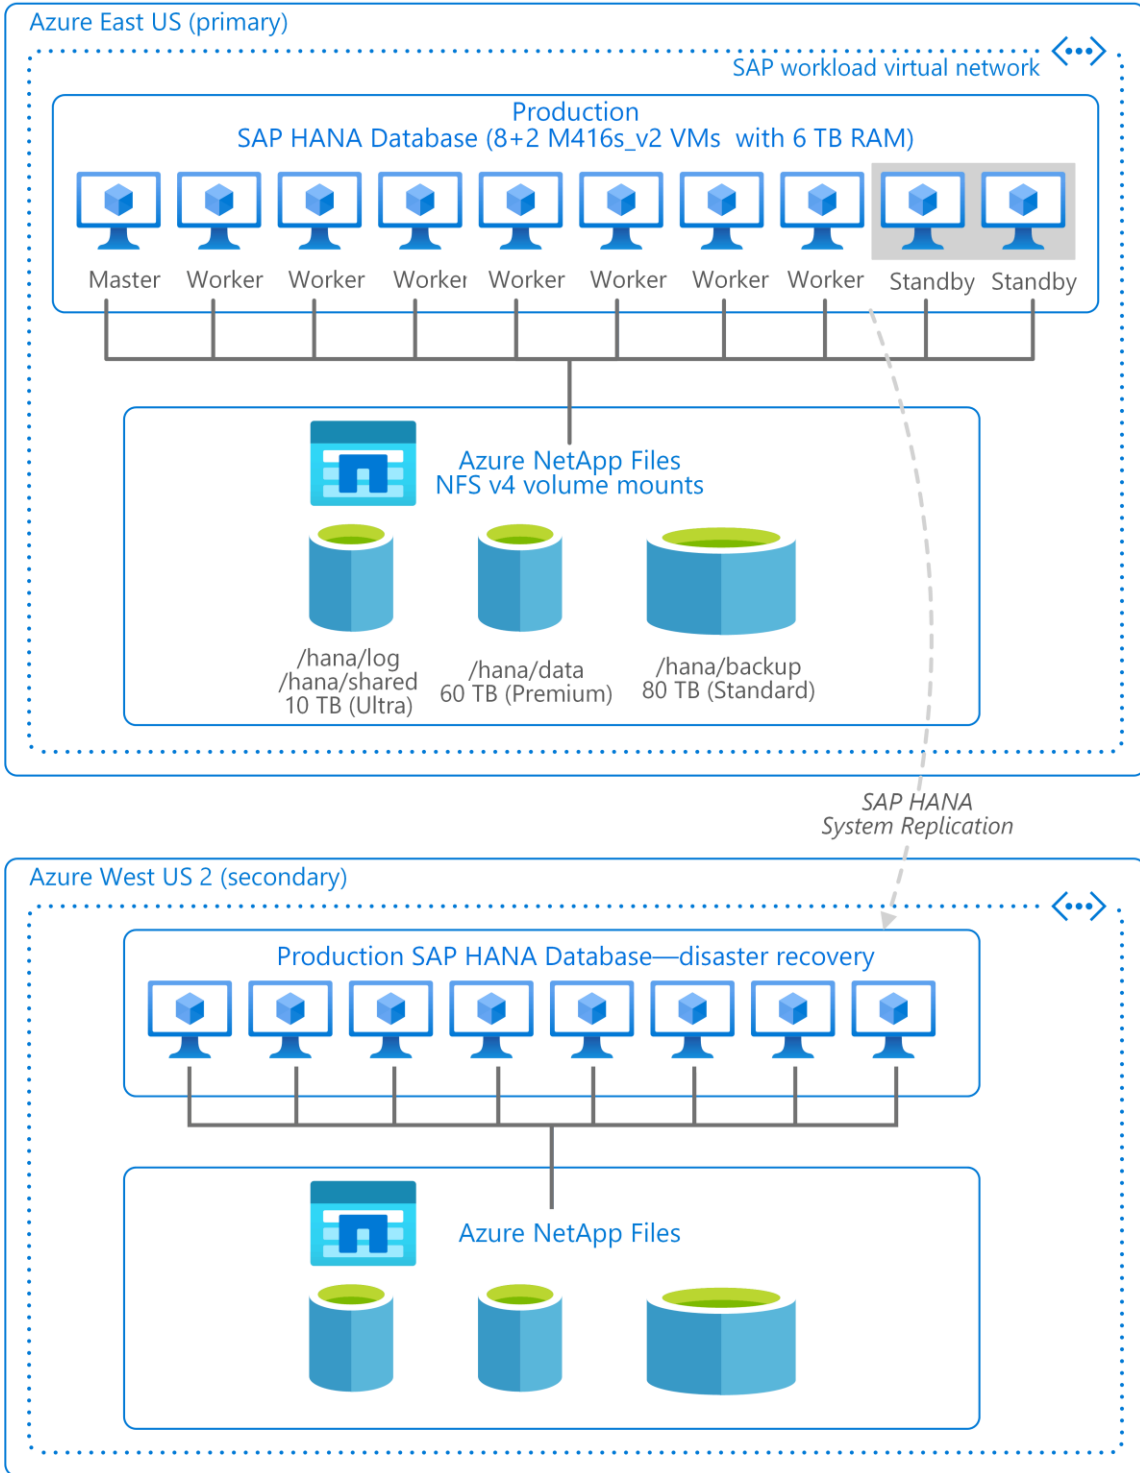


Figure 2. After: The more cost-effective configuration uses fewer Mv2-series VMs to run the database tier. Azure NetApp Files provides a highly available NFS, eliminating the need for an NFS cluster.

Getting the best performance at scale

Moving the database tier to VMs improved performance for users. For example, the increase in SAP HANA heap memory means that the bottlers' reports run smoothly. The move also reduced network latency. SAP application servers carry on constant communications with the database servers. A best practice is to place the application servers in close physical proximity to the database servers. When the database tier ran in a colocation facility, as was the case before the migration to Mv2-series VMs, CONA had to find other ways to meet performance requirements.

Other factors can affect performance. For example, the CONA load jobs were complicated, featuring complex queries on large datasets. In the original lift-and-shift architecture, network traffic traveled back and forth between the HLI servers, the application servers, and the on-premises network. Each of these hops added to network latency.

The move to VMs coincided with the rollout of powerful new gateway options. CONA implemented ExpressRoute Ultra Performance circuits, which support three times the connections per second, as compared to ExpressRoute Standard.

"We learned that the network teams need to be engaged early," Reddy adds. "It's vital to validate Azure connectivity and performance requirements."

Storage is another factor in performance. The move to Azure NetApp Files provided an impressive boost in input/output operations per second (IOPS), compared to traditional NFS mounts. In a team effort involving CONA, Microsoft, NetApp, and Capgemini, a Microsoft Azure Expert Managed Services Provider, more than 500 benchmark tests were run until the performance requirements were met. Transfer speeds also improved up 200 percent.

"During the cutover, we worked with Microsoft and NetApp to get the performance we needed," notes Reddy. "We also use Azure NetApp Files snapshot and deduplication copies to improve backup speed."

We wanted this to be so seamless that the business didn't notice the migration. And we did it. Microsoft was really helpful from day one.

—Uday Reddy, Director – Cloud Engineering, CONA Services LLC

Meeting availability and recoverability goals

In designing a highly available system, the usual strategy is to set up redundant resources and to meet the Recovery Time Objective (RTO) and Recovery Point Objective (RPO). For CONA, the RTO for the SAP HANA server is ASAP—as a mission-critical system, the business cannot tolerate downtime or data loss. In the old architecture, a failover of the physical servers meant that the reporting systems were offline until the secondary rebooted and the database was reloaded—a manual process that took up to 90 minutes.

With the new architecture, the standby nodes provide high availability. SAP HANA Host Auto-Failover minimizes RTO when a node fails by using a standby node to replace it. For a 5.7-TB VM, this process takes less than five minutes. The data is fully loaded into memory within about 15 minutes, making the combined RTO less than 30 minutes.

The new architecture also adds resiliency across regions through SAP HANA system replication, which minimizes RTO by replicating data at regular intervals. In this architecture, SAP HANA system replication replicates the database to a database instance in the disaster recovery region.

For even faster RTO, CONA plans to implement snapshot copy-based backup through Azure NetApp Files. The feature replaces the need to export a full backup copy of the database—a slow process for a database as large as CONA's. Instead, copies are created on the storage system, so they don't consume database resources. Creating a snapshot copy takes only seconds, because no data needs to be copied on the storage system.

Running a tight ship: Monitoring and backup

CONA was an early adopter of Azure Monitor for SAP Solutions. The service provides a central location for SAP telemetry, such as CPU, memory, disk, and network utilization, and the status of the host, SAP HANA system replication, and backups. Customized dashboards provide CONA with an end-to-end view of its SAP landscape. Operators can now easily recognize patterns and correlate data between various components—something that was a challenge in the original lift-and-shift architecture.

“Before we announced a private [preview of Azure Monitor for SAP Solutions](#) in September 2019, we heard from customers that they relied on complex and unmanageable disparate tools and dashboards,” explains Sameeksha Khare, an Azure Compute Program Manager who announced the preview. The new solution makes it easy to collect and consolidate Azure telemetry data and to create customized dashboards. It's also open source—[the Azure Monitor for SAP Solutions project welcomes contributions](#).

“We prepared dashboards that combine data from apps and logs,” says Reddy. CONA monitors the infrastructure health and usage and then correlates that data with the SAP database to identify patterns and to proactively tackle any issues. “The dashboards give us custom HANA monitoring, and that really helps us with troubleshooting.”

Next steps

To save on costs, CONA made an early commitment to Azure while looking for ways to optimize its massive SAP landscape. As Azure added capabilities, so did CONA, gaining the scalability and flexibility that it needed as its data footprint grew. The move also brought a 40 percent cost savings on top of improved business continuity with faster recovery.

Meanwhile, bottlers and business users now have the benefits of self-service reporting with near-real-time access to the SAP HANA data through SLT, an SAP tool for loading and replicating data. Query performance has also improved by about 30 percent, which has translated into fewer support calls for the IT team.

The move also represents a shift in day-to-day thinking for the IT team as it continues to move more components to the cloud. “Most of our business is still in our datacenter, and we’re slowly moving to cloud. That’s in progress right now,” Reddy explains.

The next step is to bring more of the SAP ERP Central Component (ECC) data to Azure for analytics and reporting. Another consideration is the 11.4-TB Mv2-series VMs. “Looking at larger nodes will tell us if it’s practical to migrate SAP ECC,” Reddy says. His team also plans a proof of concept to evaluate SAP S/4HANA on Azure. CONA and Microsoft continue to work together to drive efficiencies for the business and its network of bottlers.

The Azure NetApp Files advantage for SAP workloads

SAP shared file systems benefit from the high performance and low latency of an Azure-native NFS. Azure NetApp Files supports shared storage when using a cluster or when you need high-performance storage. CONA uses it to host the SAP HANA data and log files, which has the added benefit of supporting this type of HANA scale-out deployment model with standby nodes. Azure NetApp Files represents the culmination of a deep partnership between Microsoft and NetApp, combining NetApp’s proven ONTAP technology with the scale, reach, and enterprise capabilities of Azure.

Azure NetApp Files also supports high availability of SAP Central Services (ASCS). Azure NetApp Files can create volumes using Network File System (NFS) version 3 and 4.1, Server Message Block (SMB) 3.0 protocol, or dual protocols. Azure NetApp Files volumes, such as data or logs, can be dynamically grown or shrunk without downtime. Azure NetApp Files offers many other features for disaster recovery scenarios, such as snapshot capabilities and cross-region, block-level replication.

Learn more

Azure has a large variety of compute instances that meet TDI standards. To learn more, get started at [SAP on Azure](#).

- [SAP on Azure Architecture Guide](#)
- [SAP HANA hardware directory](#)
- [What is Azure NetApp Files](#)
- [Deploy a SAP HANA scale-out system with standby node on Azure VMs](#)

© 2021 Microsoft Corporation. All rights reserved. This document is for informational purposes only. Microsoft makes no warranties, express or implied, with respect to the information presented here.