SAP HANA on Azure 101: An Introduction to Running SAP HANA in Microsoft Azure

Abstract

This document is intended to provide an introduction to SAP HANA running in Microsoft Azure for those interested in moving to a full cloud or hybrid cloud solution. The paper includes considerations for configuration, deployment, and the scaling of your SAP HANA-based SAP system in Azure.

The intended audience for this whitepaper includes:

- **IT architects**: Who determine computing requirements for SAP landscapes and supporting systems.
- **IT professionals/administrators**: Responsible for the infrastructure that supports SAP systems.
- **SAP administrators**: Responsible for installing, configuring, and maintaining SAP landscapes.
- **IT and network administrators**: Those interested in deploying SAP applications based on SAP HANA in Microsoft Azure.
- **Technical decision makers (TDM)**: Those considering Microsoft Azure to support their SAP technical environment and influencing business decision makers (BDM) in their organization.

Contributors: Cameron Gardiner (Microsoft), Takayuki Hoshino (Microsoft), Bryan McCutchan (foursquared, Inc.), Matt Ordish (Microsoft), Juergen Thomas (Microsoft), Troy Shane (foursquared, Inc.)

***NOTE**: Certain recommendations contained herein may result in increased data, network, or compute resource usage and influence your license or subscription costs.*

Version 1, Published May 2016

Table of Contents

Executive Summary

Microsoft and SAP have partnered to offer certified SAP solutions to enterprises for the migration and deployment of their SAP applications to Microsoft Azure providing hybrid and full cloud solutions. This partnership enables customers to take advantage of on-demand resources in the cloud and maximize the value of their Microsoft and SAP investments. Microsoft Azure offers a reliable and secure cloud infrastructure platform that enables businesses to quickly deploy SAP solutions in the cloud to simplify infrastructure management, improve time to market, and lower costs.

The launch of the new SAP HANA on Azure offerings based on the Azure GS5 virtual machine type and new large instances for SAP HANA on Azure brings new levels of performance to an SAP cloud environment previously not available. This certified for production SAP HANA landscape integrates bare metal and virtualized components to achieve the best of both worlds while still offering hybrid integration with existing environments.

The new SAP HANA on Azure solution includes the following benefits:

- **Performance**: Provides up to 32-cores of CPU and 448 GB of memory for the Azure GS5 virtual machine type and 3 TB of memory for SAP HANA on Azure (Large Instances).
- **Azure Premium Storage support**: Azure Premium Storage enables up to 2,000 MB/sec of storage throughput, more than double any other public cloud provider.
- **Highest bandwidth**: The GS5 virtual machine type offers more than 20 GB/s of network bandwidth, also more than double the network throughout provided by any other public cloud provider.
- **SAP certified**: SAP has certified SAP Business Warehouse (BW) on SAP HANA running on Azure GS5 as well as SAP HANA One. In addition, for production use, SAP has certified SAP BW, SAP Business Suite, and S/4HANA running on SAP HANA on Azure (Large Instances). (S/4HANA running on Azure GS5 is also supported in controlled availability by SAP for those customers currently in the S/4HANA on Azure private preview.)

All Roads Lead to Cloud

SAP reengineered their entire business suite's roadmap to move to their SAP HANA in-memory platform to provide real-time analytics and transactions on a consolidated in-memory database optimized for both. This is a new paradigm of computing for the future enabling companies to drive their entire business in real time for unprecedented value. The roadmap calls for a phased consolidation of the entire SAP business suite beginning with business intelligence while vastly increasing performance and simplifying management with each phase. Most recently, SAP's "next-generation business suite", S/4HANA, was built on the SAP HANA platform and provides customers with simplified functional solutions with a goal of deriving more value more quickly from their SAP investment.

Just as SAP HANA has come a long way in the last several years, so has cloud computing. Initially,

SAP HANA was only certified as an on-premises installation in an appliance format and then SAP HANA Tailored Data Center Integration (TDI) certification was later released (but not as a cloud service per se). Public cloud providers running SAP HANA TDI hardware is now endorsed by SAP and has become a viable option for both production and non-production landscapes. Microsoft Azure is now seeing an exponential increase in its hosting of large SAP production environments as legacy systems are reaching end of life and decisions have to be made.

SAP HANA is a disruptive technology because legacy systems cannot run on it. As well, with SAP making SAP HANA its platform to support their entire SAP software landscape moving forward, companies must consider a new infrastructure for supporting their SAP landscapes. While SAP HANA appliances and SAP HANA TDI have always been options, both hybrid and full cloud options are an undeniable value proposition that are likely to become predominant in the near future. Regardless of an organization's initial decision on method, given the cloud value proposition, it is likely that all roads will ultimately lead to cloud as well.

Understanding Cloud Computing

To understand SAP HANA as a cloud service, it is important to understand the evolution of cloud computing and have a basic understanding of how it is offered. This section will clarify the differences between a traditional physical data center and infrastructure, platform, and software as a service.

The adoption of computing as a service, rather than as a product, is an evolution in the IT industry where vendors have worked with customers and recognized the customers' desire to manage only the minimal amount needed for a specific business purpose. Products like e-mail, collaboration, enterprise resource planning, customer relationship management etc., have evolved where customers desire a specific degree of control over the resources dedicated to each system or service, therefore governing the performance, cost and capabilities curve in managing them.

In the broad view of the industry, cloud computing has come about in three general types of services – IaaS, PaaS, and SaaS. Each of these respective services provide a sliding scale in terms of access to the underlying mechanics of the computing resources supporting the service and the required management effort.

Figure 1: The evolution of IT infrastructure.

With the evolution of IT Infrastructure over the last 15 years, we have seen customers move from fully physical (an operating system installed on dedicated hardware) to a virtualized model in the data center, taking advantage of the advances in hardware to support a denser, easily manageable, and more flexible data center.

After the shift from physical to virtual, the cloud evolution of computing began, with offerings such as Microsoft Azure virtual machines (IaaS), Microsoft Azure Web Sites (PaaS), and Microsoft Office 365 (SaaS). These services are built on top of the same infrastructure as Microsoft's other leading cloud services including Microsoft Bing, Windows Live, and Xbox Live.

This reshaping of the computing landscape means that customers have (more now than ever before), a broader choice in what makes up their computing landscape from a blend of on-premises, physical, virtual and cloud-based service offerings. Microsoft has built its cloud, Microsoft Azure, and its cloud-based services to work together to offer customers the full range of choice in how they run their Enterprise software.

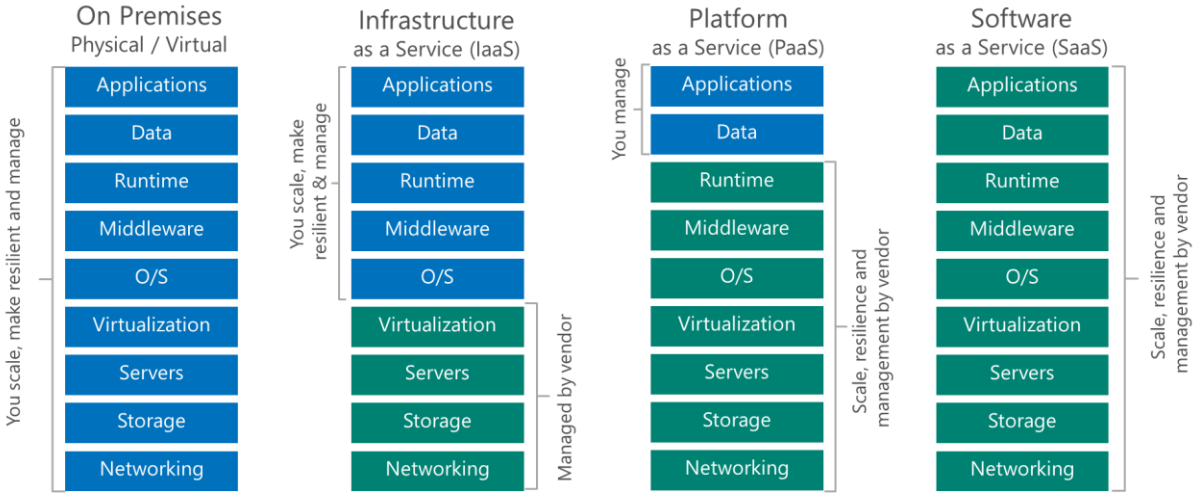

Figure 2: The division of management on the various cloud offering types.

The differences and relationship between on-premises, IaaS, PaaS, and SaaS are illustrated above in more detail, showing the demarcation in responsibilities between on-premises computing (whether physical or virtual) and cloud computing.

With the move into the cloud, infrastructure planning responsibilities is eliminated to various degrees. With on-premises computing, the customer is responsible for budgeting, planning, and roll-out of the entire computing stack including everything from physical servers and network topology, to the operating system, applications and data.

Infrastructure as a Service (IaaS) offers the customer virtual machines, whereby they are not responsible for the planning or implementation of servers, storage, virtualization, or networking, but rather assume control of the operating system and the applications running on it. This allows simple, expedient provisioning of virtual machines for development, test, quality assurance, or production.

Going one step further with Platform as a Service (PaaS), the vendor controls all of the servers,

storage, virtualization, networking, but also the operating system, the runtime and to a degree, the middleware used for messaging. The customer is responsible only for developing the applications running on the service, and the data it requires/produces. This was the first offering from Microsoft Azure before IaaS and remains a very capable and comprehensive service today.

With Software as a Service (SaaS), the customer simply logs in to an application and uses the service while the vendor manages all underlying computing resources including servers, storage, networking, the configuration of the OS, runtime, middleware and the applications and data.

Microsoft Azure Overview

To understand the scope of Microsoft Azure and the strength of the infrastructure that supports its SAP HANA cloud services, a little history is in order. Microsoft Azure was created in Microsoft's first data centers in 1989 on Microsoft's Redmond, Washington campus. These online services grew rapidly with the explosion of the Internet in the '90's and eventually came services such as Bing, MSN, Skype, Office 365, outlook.com, Xbox Live, OneDrive and the Microsoft Azure Platform; there are over 200 online services that Microsoft currently offers 24x7, 365 days of the year.

This service-driven approach has led to a consolidation and huge growth in Microsoft's data centers across the globe, run by Global Foundation Services and optimized for cloud-scale. This growth has been planned and managed with a focus on high reliability, operational excellence, cost-consciousness, and environmental sustainability.

Microsoft Azure has been and continues to be a very significant investment (+$15B) for Microsoft. Only a handful of technology companies have the ability to invest as much towards online, enterprise-level services capabilities available to customers. Along with this commitment, Microsoft Azure also provides the ability to innovate in data centers built around the world and the scale of such an investment and innovation provide significant benefits for customers. Azure is currently available in [30 online regions worldwide with an additional two planned](#).

Figure 3: Microsoft Azure data centers span across 24 regions worldwide and is still growing.

Microsoft Azure offers services in 90 countries around the world, and owns or leases over 100 data centers. These data centers are built on three core principles including:

- Reliability
- Security and compliance
- Environmental sustainability

These core principles have had to be re-thought when engineering services on a global scale have a demand like Azure, and considerable advancements have been made in the quest to provide a robust, enterprise-grade service infrastructure. This includes a new way of looking at traditional computing elements such as storage, CPU/RAM, scalability of PaaS services, securing core infrastructure, applications, data and communications across the entire Microsoft Azure offering.

Microsoft's investment has not stopped at the >1 million servers it has purchased, or the more than 1.5 million network operations it processes per second. Microsoft has gone to great lengths to demonstrate that it takes operational efficiency and business and privacy control very seriously, through significant industry standard certifications.

Understanding Microsoft Azure

Microsoft Azure is Microsoft's application platform for the public cloud providing Infrastructure as a Service (IaaS), Platform as a Service (PaaS), Software as a Service (SaaS). It is a growing collection of integrated cloud services—analytics, computing, database, mobile, networking, storage, and web—for moving faster, achieving more, and saving money. SAP HANA on Azure is provided as a IaaS solution providing a SAP certified and optimized infrastructure on which the client can deploy their own landscape. The optimized infrastructure integrates virtualized and bare metal components to achieve the best of both worlds.

Microsoft Azure Services

Microsoft divides its Azure offerings in different ways - the most obvious distinction being by class of service: IaaS, PaaS, or SaaS. Microsoft Azure services that comprise each of those are further broken down into types of service that are more closely aligned with today's data center computing concepts:

- **Compute**: Virtual machines, Web sites & mobile services
- **Data services**: Storage, SQL Azure databases, HDInsight, cache, backup & site recovery
- **App services**: Media services, service bus, notification hubs, scheduler, BizTalk Services, Visual Studio Online, Active Directory, multi-factor authentication, automation, Content Delivery Network (CDN), Azure RemoteApp & API Management
- **Network services**: Virtual networks, virtual private networks, ExpressRoute

An important factor to understand with Microsoft Azure is that these types of services are building

blocks for consuming each of the IaaS, PaaS and SaaS overall services. An example of this is that to utilize Infrastructure as a Service (IaaS) you require virtual machines from the compute section, as well as storage from data services and virtual networks from network services.

By modularizing the components that are the building blocks for Azure services, Microsoft offers customers a cost-effective way to leverage exactly the assets and services that they need to meet their computing and business needs, and also allows architectural innovation in the manual or programmatic assembling of those services to accomplish specific goals.

Accessing and Managing Azure

Azure is a subscription model that is administered through Azure subscriptions and managed through management interfaces such as the Azure Resource Manager and the underlying REST API.

Azure Subscriptions

The Microsoft Azure Portal is the web-based interface to Microsoft Azure that allows the customer to create and administer their Microsoft Azure Subscriptions, account access, storage accounts, and other Azure resources.

The Azure accounts are organized under subscriptions, which represent a specific customer, and all the customers' Azure resources are allocated underneath that subscription, such as storage accounts, containers, disks and disk images.

The subscriptions are tied to a Microsoft Account (formerly known as a Windows Live ID) which can have multiple subscriptions associated with them. This allows an individual to administer multiple Azure subscriptions, for example, finance and manufacturing, or personal and work.

From the Azure Portal, administrators can also manage users through steps like Active Directory syncing from an enterprise on-premises Active Directory.

Once the Azure subscription is accessible through the customers' Microsoft account, they get access to all the resources available within that subscription, including virtual machines, virtual disks, networks, and worker roles etc.

Creation and management of resources is open to each user that is given access to the Portal, so planning should include who will be given access to each subscription and how Azure components are deployed should be looked at on a per-subscription basis.

This portal can be accessed at: http://portal.microsoft.com.

Azure Resource Manager

The infrastructure for an application is typically made up of many components – for example, a virtual machine, storage account, and virtual network. These components are not seen as separate entities, but as related and interdependent parts of a single entity that can be deployed, managed, and monitored as a group. Azure Resource Manager enables working with the resources in a solution as a group. Customers can then deploy, update or delete all of the resources for their solution in a single, coordinated operation and then use a template for deployment and that template can work for different environments such as testing, staging and production. Resource

Manager provides security, auditing, and tagging features to help manage resources after deployment.

***NOTE**: Azure has two different deployment models for creating and working with resources: [Resource Manager and classic](). To simplify the deployment and management of resources, Microsoft recommends that you use Resource Manager for new resources, and, if possible, re-deploy existing resources through Resource Manager.*

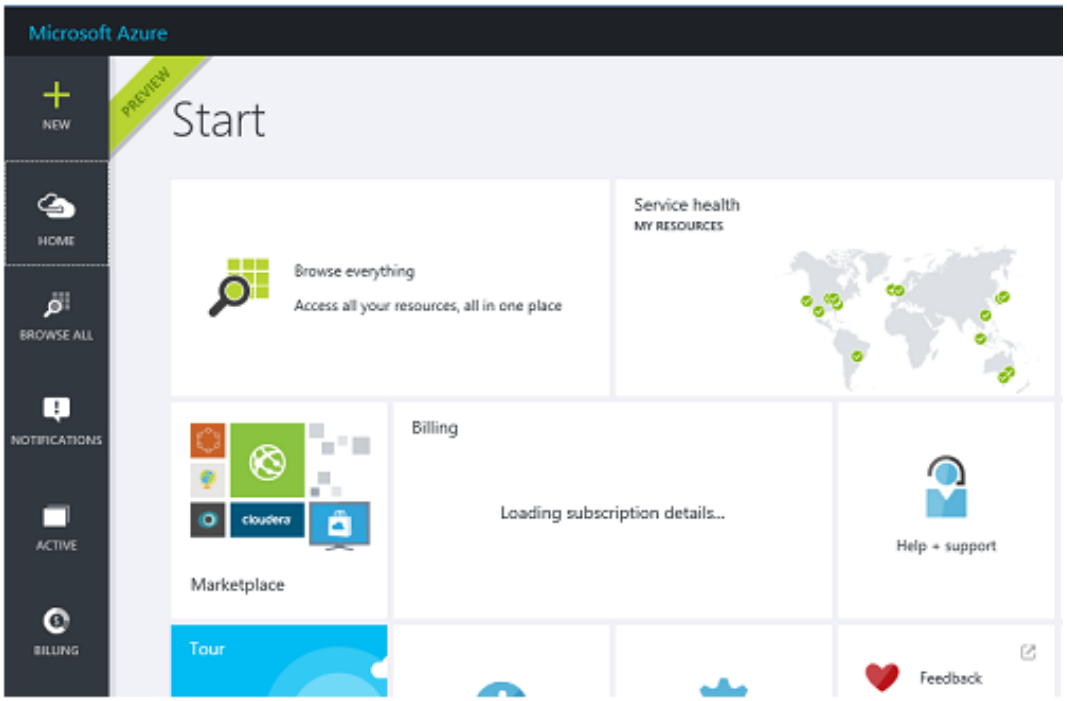

Figure 4: An Azure Resource Manager screenshot.

The Resource Manager deployment model provides a new way to deploy and manage the services that make up your application. This new model contains important differences from the classic deployment model, and the two models are not completely compatible with each other.

REST API

The REST API allows for full-featured programmatic control of Azure Services, meaning that customers can quite simply script a task in Azure through PowerShell, for example, that increases the PaaS worker roles, or creates new IaaS virtual machines from a gallery image or the customers' own template.

This REST API underlies the Azure Management Portal and exposes some functions that are not explicitly made available within the portal.

For information on getting started with Azure CmdLets and for downloading the latest PowerShell modules (as a local installation), please see:

http://msdn.microsoft.com/en-us/library/azure/jj554332.aspx

Infrastructure as a Service (IaaS)

Microsoft Azure Infrastructure as a Service is Microsoft's offering that consist of three separate parts:

- **Virtual machines**
- **Storage, backup, and recovery**
- **Big compute**

Moving forward in this paper, we will focus only on Microsoft Azure virtual machines and Azure storage with respect to SAP and SAP HANA running on Azure. Many, but not all, of the component services underlying IaaS are utilized to support SAP solutions running on Microsoft Azure.

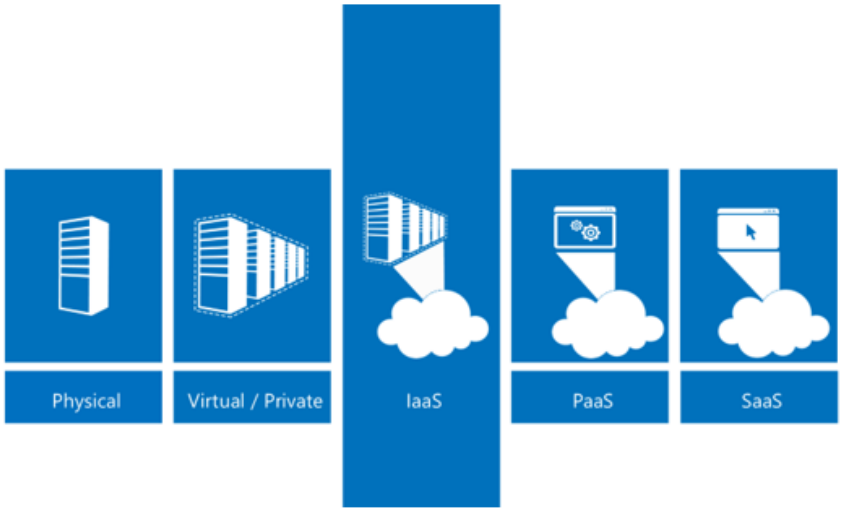

Figure 5: SAP HANA on Azure is an IaaS offering.

Microsoft Azure Virtual Machines

Azure virtual machines offer near-instant deployment capabilities that drastically reduce the barrier to entry and necessary capital planning and budgeting for small to large organizations. With built in Virtual Networking and load-balancing, compatibility with Hyper-V virtual machines running on premise, flexibility, and a low cost to run preconfigured virtual machine images offered through the Azure gallery, Azure virtual machines are designed to be easy to deploy and offer operational value for the SAP landscape.

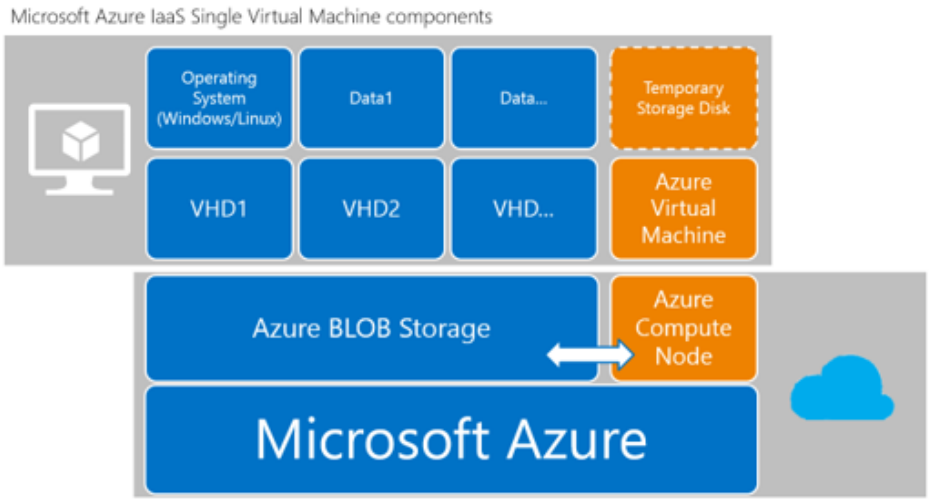

Figure 6: Microsoft Azure IaaS single virtual machine components.

The diagram above illustrates Microsoft Azure supporting services that comprise a single IaaS virtual machine. When configuring a virtual machine on Microsoft Azure, you utilize both Azure compute nodes to deploy the virtual machine to, as well as Azure Storage to deploy the attached virtual hard drives (both operating system, data partitions, and BLOB objects).

These 2 major components, along with an attached virtual network, are the underlying foundation for all virtual machines in Microsoft Azure.

The virtual machine, when deployed (or 'allocated'), is established on an Azure compute node – a server that supplies it with both some temporary storage, as well as the actual CPU and RAM required for computing. This is the start of the actual charges for using the virtual machine, which is billed by the minute. When you 'deallocate' your virtual machine (shutdown through the Azure Management Portal or API), the virtual machine is removed from the compute node and the only remaining cost is the virtual hard disks, which persists in Azure BLOB storage.

***NOTE**: It is important to note that issuing a shutdown to the operating system inside the virtual machine does not deallocate the virtual machine from the compute node and the hourly virtual machine costs will continue to be incurred.*

Virtual machine hard disks are attached to virtual machines as VHDs from the Azure BLOB Storage service (including both the operating system (hard disk (C:) and other disks attached that have executables, databases, etc. When allocated to a compute node, the virtual machine operating partition is copied from Azure Blob storage onto the compute node and resides there for the duration of the virtual machine allocation. All other VHDs stay resident in Azure Blob Storage and are communicated with over the network.

In addition to the OS and data partitions, each virtual machine is provided with one temporary virtual hard disk that exists only on the Azure compute node. When the virtual machine is first allocated, this drive will show up as the D: drive, but can be reassigned a different drive letter (e.g.

Z:\) at the users' discretion (which will persist through the lifetime of the virtual machine; including through allocation / deallocation).

***NOTE**: Due to the temporary nature of this virtual hard disk, it will not persist any data stored there, so it is unwise to utilize it for storing anything of importance. One possible use could be the operating system swap / page file.*

Microsoft Azure Virtual Machines in a Virtual Network

When deploying multiple virtual machines within an Azure data center, you most often will connect your virtual machines (deployed across various compute nodes around the data center, through the use of a virtual network), allowing them to communicate with each other. Different than a VPN, the virtual network equates to both the physical connection of a network to the virtual machine, as well as the subnet definition for the virtual machines attached to that virtual network, providing services from several layers in the OSI model.

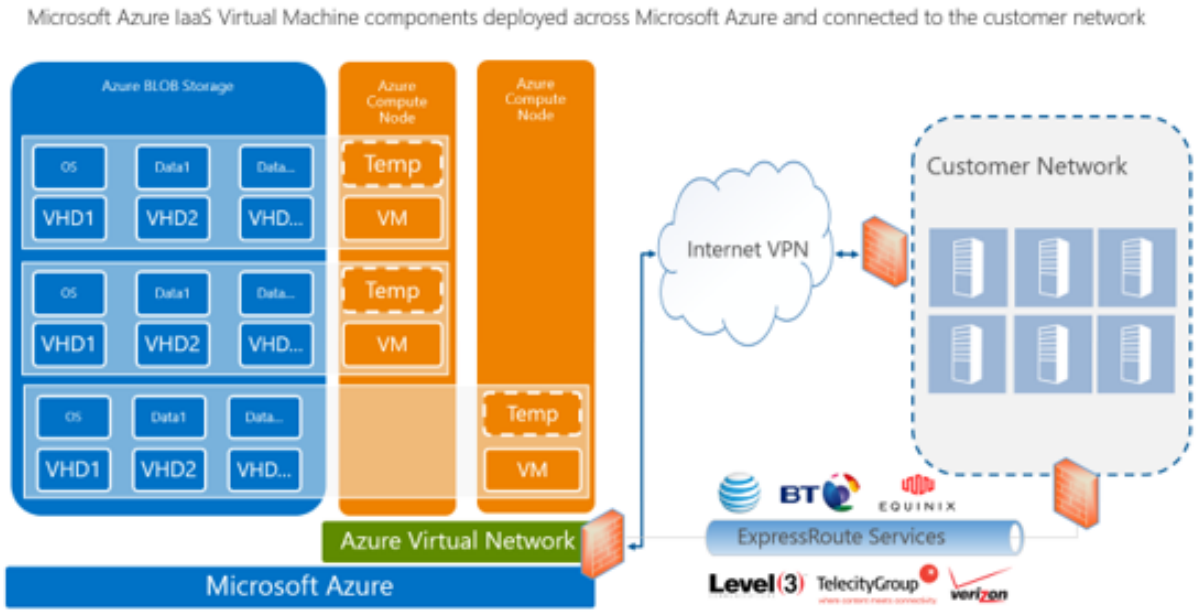

Figure 7: Microsoft Azure IaaS virtual machine component architecture.

Virtual networks are also used when connecting point-to-site VPN (single computer to Azure virtual network), site-to-site VPN (single VPN connection from customer network to Azure virtual network), multi-site VPN (multiple VPN connections from customer network into Azure) and ExpressRoute (dedicated networking connections through telco providers), to allow communication with the running virtual machines.

Microsoft Azure Virtual Machine Types

Microsoft Azure virtual machines are offered in predefined sizes: A-Series that offers a wide range of computing from a shared CPU core with access to 768 MB of memory ranging up to 16 CPU cores and 112 GB of memory; these compute nodes utilize SATA hard disks drives for local storage.

The release of the D-Series of virtual machines offer an even more attractive price-to-performance ratio. The D-Series of virtual machines are the next generation of hardware and use SSDs on the compute nodes, along with bringing a more cost-effective platform on which to run SAP.

The virtual machines are priced according to size, and also have dedicated levels of service with respect to how many data disks you can attach, the storage I/O quotas applied, network quotas, and as well, the size of the temporary virtual hard disk increases with the larger virtual machine sizes.

Later, we will examine the relationship between the virtual machine size and the SAP System architecture, performance and management implications.

Virtual machines are available in two tiers: basic and standard. Both types offer a choice of sizes, but the basic tier doesn't provide some capabilities available in the standard tier, such as load-balancing and auto-scaling. The standard tier of sizes consists of different series: A, D, DS, G, and GS. Considerations for some of these sizes include:

- **D-series virtual machines**: Designed to run applications that demand higher compute power and temporary disk performance. D-series virtual machines provide faster processors, a higher memory-to-core ratio, and a solid-state drive (SSD) for the temporary disk. For details, see the announcement on the Azure blog, New D-Series Virtual Machine Sizes.
- **G-series virtual machines**: Offer the biggest size and best performance and run on hosts that have Intel Xeon E5 V3 family processors.
- **DS-series and GS-series virtual machines**: Can use Premium Storage, which provides high-performance, low-latency storage for I/O intensive workloads. These virtual machines use solid-state drives (SSDs) to host a virtual machine's disks and also provide a local SSD disk cache. Premium Storage is available in certain regions. For details, see Premium Storage: High-performance storage for Azure virtual machine workloads.
- **SAP HANA on Azure (Large Instances):** SAP HANA on Azure (Large Instances) provide purpose-built SAP certified hardware and single-tenant bare-metal performance to handle the largest SAP HANA workloads (up to 3 TB of RAM) of any hyper-scale public cloud provider.

The size of the virtual machine affects the pricing. The size also affects the processing, memory, and storage capacity of the virtual machine. Storage costs are calculated separately based on used pages in the storage account. For details, see Virtual Machines Pricing Details and Azure Storage Pricing. For more details about storage for virtual machines, see information about disks and VHDs for virtual machines.

***NOTE**: See SAP Support Note #1928533 – "SAP Applications on Azure: Supported Products and Azure virtual machine Types" for most current information on supported configurations.*

Table 1: Current virtual machine types with respective sizes, number of SAPs supported, v-disks, and price.

Tier	VM Type	VM Size	Temp SSD	SAPS	# of v-disks
2-tier	A5	2 CPU, 14 GB		1,500	4
2-tier	A6	4 CPU, 28 GB		3,000	8
2-tier	A7	8 CPU, 56 GB		6,000	16
2-tier	A8	8 CPU, 56 GB		11,000	16
2-tier	A9	16 CPU, 112 GB		22,000	16
2-tier	A10	8 CPU, 56 GB		11,000	16
2-tier	A11	16 CPU, 112 GB		22,000	16
3-tier	A5 x 10	20 CPU, 140 GB		12,000	
3-tier	A6 x 10	40 CPU, 280 GB		25,000	
3-tier	A7 x 10	80 CPU, 560 GB		50,000	
2-tier	D11 / DS11	2 CPU, 14 GB	100 / 28 GB	2,325	4
2-tier	DS11v2	2 CPU, 14 GB	100 / 28 GB	3,530	4
2-tier	D12 / DS12	4 CPU, 28 GB	200 / 56 GB	4,675	8
2-tier	DS12v2	4 CPU, 28 GB	200 / 56 GB	6,680	8
2-tier	D13 / DS13	8 CPU, 56 GB	400 / 112 GB	9,350	16
2-tier	DS13v2	8 CPU, 56 GB	400 / 112 GB	12,300	16
2-tier	D14 / DS14	16 CPU, 112 GB	800 / 224 GB	18,770	32
2-tier	DS14v2	16 CPU, 112 GB	800 / 224 GB	24,180	32
2-tier	GS1	2 CPU, 28 GB	384 GB	3,580	4
2-tier	GS2	4 CPU, 56 GB	768 GB	6,900	8
2-tier	GS3	8 CPU, 112 GB	1,536 GB	11,870	16
2-tier	GS4	16 CPU, 224 GB	3,072 GB	22,680	32
2-tier	GS5	32 CPU, 448 GB	6,144 GB	41,670	64

Microsoft Azure Storage Services

Azure storage service is a highly scalable, locally and geographically redundant service that offers up the storage required by Azure virtual machines in order to provide virtual hard disks (VHDs) for operating system and data partitions, storage locations for various workloads (e.g. database files; application backup files) and secure, reliable, decentralized backup and recovery services.

Microsoft Azure Accounts are organized hierarchically, and it is important to understand the relationships and restrictions between the different elements in order to properly plan an implementation of business applications on Microsoft Azure.

Microsoft Azure storage offerings are designed to be durable, scalable and cost-effective. This is accomplished through several types of storage, suited for different uses:

- **Block blob storage**: Designed to store large amounts of text and binary data.
- **Page blob storage**: Designed for high-performance, frequent updates, durability and reliability.
- **Table and queue storage**: Offering NoSQL Tables for storage of semi-structured/unstructured data and reliable messaging Queues that allow messaging and scheduling of asynchronous tasks.

Azure storage is also replicated for high availability, and there are three options offered:

- **Locally Redundant Storage (LRS)**: Replicated three times within a data center.
- **Geo-redundant Storage (GRS)**: Replicated three times within a data center, and also three times at a second data center (six copies total).
- **Read-Access Geo-Redundant Storage (RA-GRS)**: Same features as GRS, plus allowing read-access at the second data center.

Page Blob Storage

Persistent virtual machine VHDs are stored in Azure Page Blob Storage, attached at virtual machine configuration time and exist through start/stop/restart and allocation/de-allocation of virtual machines to compute nodes. Azure storage accounts are allocated Azure subscriptions (multiple storage accounts can exist within a subscription) and are organized in a hierarchical fashion, with VHD's being stored in individual containers within the storage account. Each Azure subscription can have up to 50 Azure storage accounts.

VHD's are 'Page BLOB' objects (optimized for frequent read/write) are up to 1 TB in size, provide up to 500 IOPS each, need to be fixed (not dynamic) and are the location where you install the OS, applications or databases and store data.

***NOTE**: Azure Storage accounts have a 20,000 IOPS limit per account, so it's important to manage your SAP system deployments across storage accounts to achieve optimal performance.*

When Azure virtual machines are allocated to compute nodes, they have access to two different types of VHD storage, the multiple, persistent VHDs that reside in Azure BLOB Storage, and 1 non-persistent VHD that resides on the compute node itself. One exception there is that, during allocation, the virtual machine operating system VHD gets copied to the Azure compute node that hosts the virtual machine for the duration of the allocation. The non-persistent local VHD that originally becomes available as the D: drive on the Azure virtual machine is not designed to contain important information, as it is destroyed between virtual machine allocations. This means that any data placed on it becomes unavailable from allocation to allocation, and for that reason, it is best used for non-critical data, such as the OS page file.

Table 2: Page blob storage pricing model.

Storage Type	Max IOPS Per Disk	Price
Locally Redundant Storage (LRS)	500	$0.05 /GB /month
Geographically Redundant Storage (GRS)	500	$0.095 /GB /month

Premium Storage

Premium Storage delivers high-performance, low-latency disk support for I/O intensive workloads running on Microsoft Azure virtual machines. Several Premium Storage disks can be attached to a virtual machine. With Premium Storage, applications like SAP can have up to 64 TB of storage per virtual machine and achieve 80,000 IOPS (input/output operations per second) per virtual machine and 2,000 MB per second disk throughput per virtual machine with extremely low latencies for read operations. Microsoft highly recommends using Premium Storage for any production SAP landscape.

Table 3: Page blob storage pricing model.

Disk Type	Disk Size	IOPS/Disk	Throughput per Disk	Price
P10	128 GB	500	100 MB/s	$19.71 /month
P20	512 GB	2,300	150 MB/s	$73.22 / month
P30	1024 GB	5,000	200 MB/s	$135.17 /month

Business Considerations

As with any new technology, there's an evaluation period when organizations must determine whether this new technology is right for their organization. This process typically starts from a technical perspective evaluating capabilities, functionality, administration and support requirements, as well as cost/benefit analysis. This period is typically referred to by analysts like Gartner as "emerging" within an overall maturity cycle of software. During this time, proof-of-concepts (POCs) are often conducted and initial customers deploy the software in cooperation with the software vendor. In this case, SAP HANA and even more specifically, SAP HANA running on Microsoft Azure fits in the emerging phase of software maturity.

Whether it's SAP's traditional applications like SAP Business Warehouse (BW) and SAP Business Suite (i.e. SAP ECC) deployed on the SAP HANA database or their new S/4HANA platform, Microsoft Azure provides you with the capability to deploy your SAP HANA-based system quickly and at less cost than what an on-premises SAP HANA deployment would be – for example, installation of a SAP HANA appliance. Then, after thorough evaluation of SAP HANA's capabilities and Microsoft Azure's ability to host and manage your installation of SAP HANA, you can look forward to long-term deployment and management of your system in Microsoft Azure.

Value Proposition

Running SAP HANA on Microsoft Azure provides a number of customer benefits:

- **Hardware cost savings:** No longer do you have to procure expensive hardware (such as an SAP HANA Appliance) via your CapEx budget and then have IT provision over a relatively long period of time. The cost of Microsoft Azure is subscription-based and thus, OpEx from a budget perspective.
- **Operational cost savings:** Environments like development and test can be used only when they're necessary. When they're not in use, they can be turned off (deallocated) and therefore, minimizing operational costs.
- **Quick deployment:** The ability for your business users to evaluate new technology becomes easier as new SAP HANA environments can be deployed in literally minutes from established virtual machine images already developed.
- **More flexibility:** By way of configuring Microsoft Azure or using available 3rd-party tools, starting and stopping SAP HANA instances based on a time schedule for training, demo, or other purposes is feasible and provides much more flexibility than physical hardware environments.
- **Integration:** Integrating with other solutions such as Internet-of-Things (IoT) Device Gateway, Hadoop, Predictive Analytics, and Business Productivity is more feasible with deployment of SAP on Microsoft Azure.

In addition, the following points are based on actual customer conversations regarding running SAP on Microsoft Azure:

- **Customers with SAP project implementations are deploying an entire landscape (development, QA, and production) on Microsoft Azure**: Those that aren't yet utilizing Azure for production are putting their SAP sandbox, development, QA, and DR systems on Azure integrated with their production environment on-premises.
- **Development, QA, and temporary project systems:** For upgrades and support packs are the most common use case for SAP on Microsoft Azure.
- **Customers are moving old "legacy" archived systems from UNIX/Oracle or DB2 to Windows and SQL Server on Microsoft Azure**: The customer will not typically buy a SQL Server license (especially where such a system might only run a few hours per month), but will use a pay-per-minute SQL Server virtual machine image. These "legacy" archive systems are typically old SAP implementations that are now no longer active, but must be kept for tax, compliance, and audit purposes. Many of the legacy systems may run on expensive IBM or HP UNIX servers that use a lot of data center space, electricity and cooling. If legal requirements mandate that the data must be resident in a particular country, then the customer can take a compressed backup of the legacy, archived system and store the

backup on-premises. In many jurisdictions, this may meet respective legal requirements.

- **Disaster recovery systems for SAP on Microsoft Azure:** Another popular use case is to use Azure as a disaster recovery data center and use Azure Site Recovery as a cost effective solution to provide DR benefits.

Since the release of Microsoft Azure Premium Storage and larger virtual machines, more and more customers are deploying production systems onto Microsoft Azure.

Analysts

An important part of evaluating a particular technology is often engaging third-party analysts and research firms to understand their perspective and recommendations.

In particular, one the most influential firms, Gartner, publishes an annual "Magic Quadrant" (MQ) that ranks vendors based on "completeness of vision" and "ability to execute". Microsoft, and in particular Microsoft Azure, continues to rank as a leader in the MQs for Cloud IaaS, Cloud Storage, Enterprise App PaaS, x86 Server Virtual, and Operational DBMS Systems.

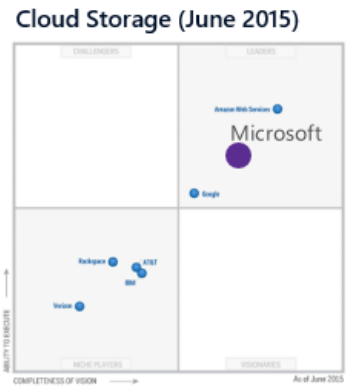

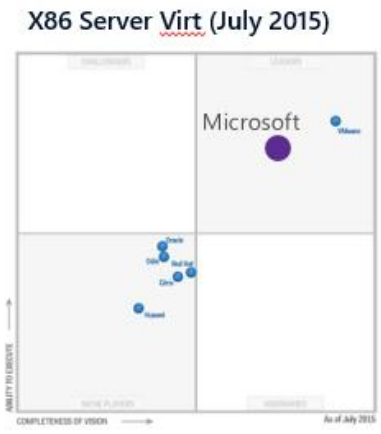

> *"Microsoft's comprehensive hybrid story, which spans applications and platforms as well as infrastructure, is highly attractive to many companies, drawing them towards the cloud in general."*
> *- Lydia Leong, Gartner*

Figure 8: Microsoft Azure is consistently a leader in Gartner's Magic Quadrants.

Benchmarking

Another aspect of building trust towards a particular technology is benchmarking and related certification by vendors – in this case, benchmarking performed by Microsoft and certified by SAP. The latest benchmarking information from SAP is available at http://www.sap.com/benchmark.

SAP's Standard Application Benchmarks are a hardware-independent measurement of processing capacity, the most common of which is the 2-tier Sales and Distribution (SD) benchmark. In this benchmark, 100 SAPS = 2,000 fully processed order line items per hour and is a direct measurement of a single instance's ability to run both the SAP application server components and the database on a single machine (physical or virtual) and provide a measure of computing capacity. When scaling the SAP system however, adding application server's results in a 3-tier configuration leading to considerable additional processing capability. While full 3-tier measurements have not yet been certified, the capacity exhibited by the 2-tier measurements demonstrates considerable scalability.

In addition, while the benchmark performed against Microsoft Azure's A9 virtual machine results in a higher SAPS rating (see above table) than the benchmark performed on the D14 virtual machine by roughly 20%, the D14 virtual machine's cost is just above half of the cost of the A9 virtual machine. That means about 50% of the cost for 80% of the performance.

This performance-to-cost differential holds true through the similar A-series to D-series virtual machine sizes, suggesting that the D-series virtual machines should be adopted over the A-series for price-to-performance benefit.

Use Cases

In planning for SAP HANA on Azure, it is important to determine your best use case as well as a roadmap for the future.

Test Drive on SAP HANA Developer Edition

In 2014, Microsoft and SAP announced support of running SAP HANA Developer Edition on Microsoft Azure. Since then, many customers have deployed SAP HANA Developer Edition via the availability of a SUSE Linux-configured virtual machine in the SAP Cloud Library (http://cal.sap.com). While this whitepaper is not focused on this developer edition, you can find more information on setting this scenario up at http://azure.microsoft.com/blog/2014/07/22/step-by-step-guide-for-

deployment-of-sap-hana-developer-edition-on-microsoft-azure/. (Keep in mind that SAP HANA Developer Edition is a limited use license of SAP HANA specifically for use by developers in a non-production environment.)

SAP HANA One

Build and deploy on-demand applications faster than ever before, with SAP HANA One – the community supported SAP HANA database. Now supported for production and non-production use on Azure DS14v2, Azure customers can build and deploy on-demand applications that leverage SAP HANA's real time, in-memory computing capabilities.

Deploying SAP HANA for Non-Production Use

As mentioned previously, SAP HANA is an emerging technology platform and deploying on Azure provides the ability for companies to evaluate and then run development, test, sandbox, and/or training environments for SAP Business Suite, SAP Business Warehouse, or even the new SAP S/4HANA solutions. For example, with SAP's new release of SAP Simple Finance based on S/4HANA, many customers are looking to effectively evaluate the solution. Instead of going down the traditional path of procuring hardware and then installing SAP for being able to properly evaluate the business aspects of running Simple Finance, you can do the same on Azure for less cost and have more flexibility as previously discussed.

Deploying SAP HANA for Production Use

SAP has certified the deployment of SAP BW on SAP HANA and S/4HANA (in controlled availability) for the Azure GS5 virtual machine type (allowing up to a 448GB in-memory database) and SAP BW and SAP Business Suite on large instances for SAP HANA on Azure (up to 3TB in-memory). This provides the most flexibility and cost effective approach of any public cloud provider for SAP HANA solutions.

Planning Deployment

Before deploying SAP software on Azure it is absolutely essential to completely read and understand the SAP on Azure documentation: https://msdn.microsoft.com/library/dn745892.aspx.

The Microsoft best practice guidance for deploying SAP on Microsoft Azure is contained in these documents and the provided configurations have been tested, proven, and should be followed.

Another important aspect at the start of your SAP HANA on Microsoft Azure deployment is to document and plan the configuration of virtual machine types, networking, storage, and disk layout before deploying. A comprehensive landscape diagram in Visio with respective inventory of Microsoft Azure object should exist before the actual objects are deployed.

Furthermore, it's important to develop a clear naming convention for all Microsoft Azure objects. Two examples of naming conventions are below.

Table 5: Naming convention example 1 – A virtual machine.

Name	Note
vmtrnecphdb01	A training SAP Business Suite on HANA database server virtual machine with a SID "ECP"
vm	denotes the Azure object is a virtual machine
trn	denotes the object belongs to the training landscape
ecp	identifies the SAP SID the virtual machine belongs to
hdb	identifies the fact it is a database server
01	identifies the object is the first in a series (02 would be the HA pair for example)

Table 6: Naming convention example 2 – A storage account.

Name	Notes
satrngeoneu01	A Storage Account used for non-SAP workloads
sa	denotes the Azure object is a Storage Account
trn	denotes the object belongs to the training landscape
geo	identifies the storage account is geo-redundant
neu	identifies the storage account is in North Europe
01	identifies the object is the first in a series

The above examples are just that and customers should create a naming scheme that best suits their organization and policies; however, it is essential to have a clear naming standard.

***NOTE**: Unless a clear and coherent naming scheme is established and implemented it will be difficult to clearly identify and troubleshoot even a medium sized Azure deployment. Furthermore, Azure storage account names are exposed as URLs and so they must be unique not only within the organization, but globally.*

For more information on naming conventions, please see

https://azure.microsoft.com/en-us/documentation/articles/virtual-machines-infrastructure-services-implementation-guidelines/.

Additional Deployment Planning Resources

In order to properly plan for deploying your SAP HANA-based solution on Microsoft Azure, please find the following resources and information from Microsoft and SAP.

Microsoft Resources and Information

- Microsoft Azure: http://azure.microsoft.com/en-us
- Microsoft Azure and SAP alliance site: https://azure.microsoft.com/en-us/campaigns/sap
- Microsoft Solution Developer Network (MSDN) SAP on Microsoft Azure site: https://msdn.microsoft.com/en-us/library/azure/dn745892.aspx
- MSDN Blog site for SAP on the Microsoft platform: http://blogs.msdn.com/b/saponsqlserver/
- SAP on Microsoft Azure Documentation: https://msdn.microsoft.com/library/dn745892.aspx

SAP Resources and Information

- SAP Community Network (SCN) SAP on Microsoft site: http://scn.sap.com/docs/DOC-56602
- SAP on Microsoft Azure FAQ: http://scn.sap.com/docs/DOC-56406
- SAP HANA Help Portal: http://help.sap.com/hana_platform
- SAP HANA on the SAP Service Marketplace: https://service.sap.com/hana
- SAP Standard Application Benchmarks in Cloud Environments: http://global.sap.com/campaigns/benchmark/appbm_cloud.epx
- SAP Benchmarking: http://www.sap.com/benchmark
- SAP Information on 'SAPS' (Unit of Measure): http://global.sap.com/campaigns/benchmark/measuring.epx
- SAP Information on Sales and Distribution (SD) Benchmark:
- http://global.sap.com/campaigns/benchmark/appbm_sd.epx
- SAP HANA Quicksizer: http://service.sap.com/hanaqs
- SAP Quicksizer: http://service.sap.com/quicksizer
- SAP HANA Storage Requirements: http://scn.sap.com/docs/DOC-62595
- SAP HANA Administration Guide: http://help.sap.com/hana/SAP_HANA_Administration_Guide_en.pdf

- SAP HANA Installation Guide: http://help.sap.com/hana/SAP_HANA_Server_Installation_Guide_en.pdf
- SAP HANA Master Guide: http://help.sap.com/hana/SAP_HANA_Master_Guide_en.pdf
- SAP HANA High Availability Whitepaper: http://scn.sap.com/docs/DOC-65585
- SAP HANA System Replication Guide: http://scn.sap.com/docs/DOC-47702

Table 4: Key SAP Note topics.

Topic	SAP Note Number
SAP Key Note for Public Cloud Environments	#1380654: "SAP support in public cloud environments"
Windows in Public Cloud Environments	#1409604: "Virtualization on Windows: Enhanced Monitoring" #1612283: "Hardware Configurations Standards and Guidance"
Linux in Public Cloud Environments	#1102124: "SAPOSCOL on Linux: Enhanced function"
SAP with Microsoft Azure	#1928533: "Supported Products and Azure VM Types" #2015553: "SAP on Microsoft Azure: Support prerequisities" #2178632: "Key Monitoring Metrics for SAP on Microsoft Azure" #1999351: "Troubleshooting Enhanced Azure Monitoring for SAP" #2035875: "SAP on Microsoft Azure: Necessary Adaption of your SAP License"
SAP HANA Sizing Considerations	#1637145: "SAP BW on HANA: Sizing SAP In-Memory Database" #1736976: "Sizing Report for BW on HANA" #1855041: "Sizing Recommendation for Master Node in BW on SAP HANA" #2121330: "FAQ: SAP BW on HANA Sizing Report" #2086829: "SAP HANA Dynamic Tiering Sizing" #1793345: "Sizing for SAP Suite on HANA" #1872170: "Sizing Report for SAP Business Suite on HANA" #1698281: "Assess the Memory Consumption of SAP HANA"
SUSE Linux Deployment Information	#1984787: "SUSE Linux Enterprise Server 12: Installation Notes" #1310037: "SUSE Linux Enterprise Server 11: Installation Notes" #1824819: "SAP HANA DB: Recommended OS Settings for SLES 11" #1944799: "SAP HANA Guidelines for SLES Operating System"
SAP HANA Deployment Information	#1523337: "SAP HANA Database Central Note"

Licensing

From an SAP licensing perspective, any deployment of SAP HANA on Microsoft Azure is considered a 'BYOL' (bring your own license) as you will procure your own SAP HANA licenses directly from SAP.

Also of note, you will need to apply a new license key each time SAP HANA is installed in a new virtual machine (even if copying or cloning from an existing SAP HANA instance).

Reference Architecture

In order to enable Microsoft partners and initial customers to deploy SAP HANA on Microsoft Azure, we have several reference architecture scenarios. While there are clearly more configurations and use cases to consider, these are what we believe the majority of our customers will initially deploy.

Single Node / Scale-Up

This is the deployment of the SAP Business Suite, SAP Business Warehouse, or S/4HANA on the SAP HANA database platform in a 3-Tier configuration including a database and application server. The SAP environment performance can be scaled as needed by adding additional application servers to the environment.

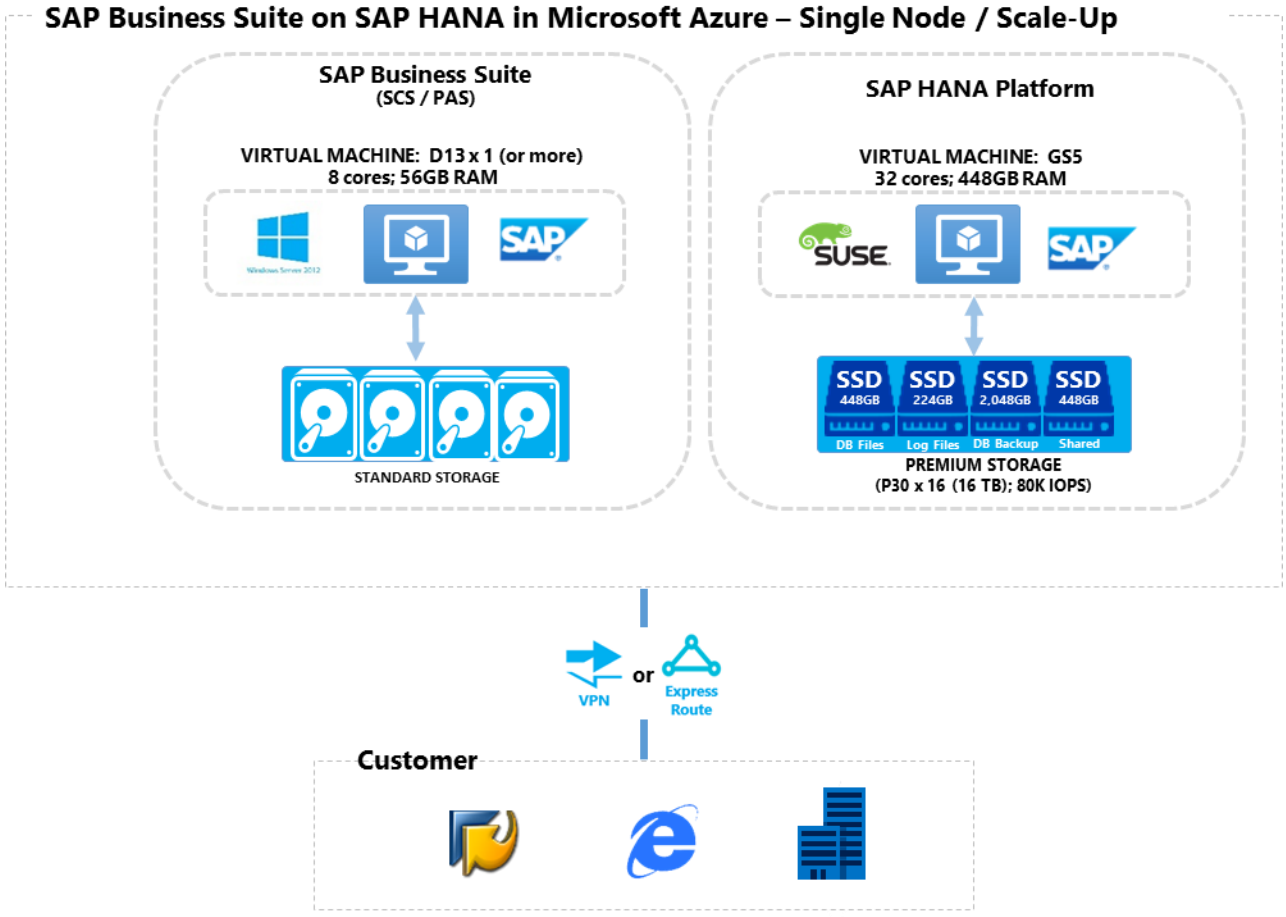

Figure 9: Single Node / Scale-Up Architecture.

Multi-Node / Scale-Out

The deployment of SAP Business Warehouse powered by SAP HANA with a master and respective worker nodes to support a scale-out architecture. For failover configuration of a scale-out scenario, see SAP's HANA Administration Guide.

***NOTE**: SAP recommends at least three worker nodes be used for a scale-out architecture.*

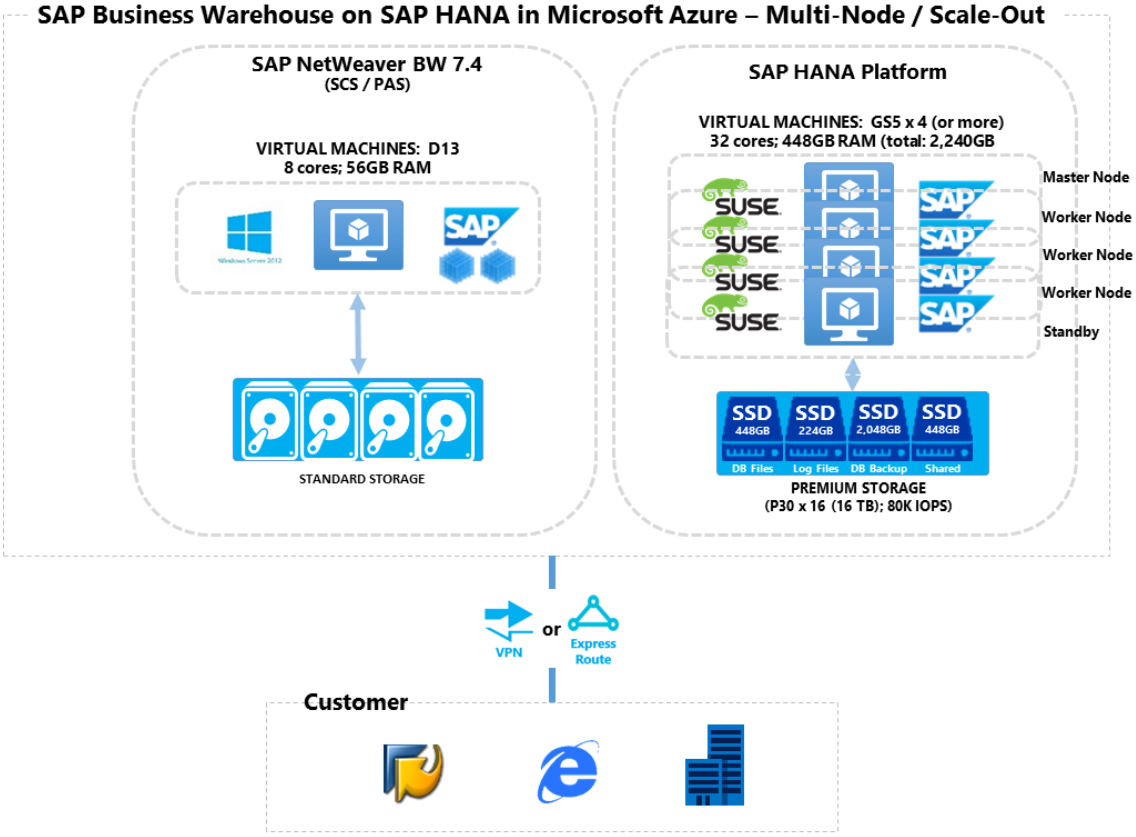

Figure 10: Multi-Node / Scale-Out Architecture.

System Replication

High availability for systems is an important consideration for enterprise customers that are concerned about business continuity. SAP HANA provides the ability to replicate itself within the same or over two data centers and this method can be applied to SAP HANA in Azure.

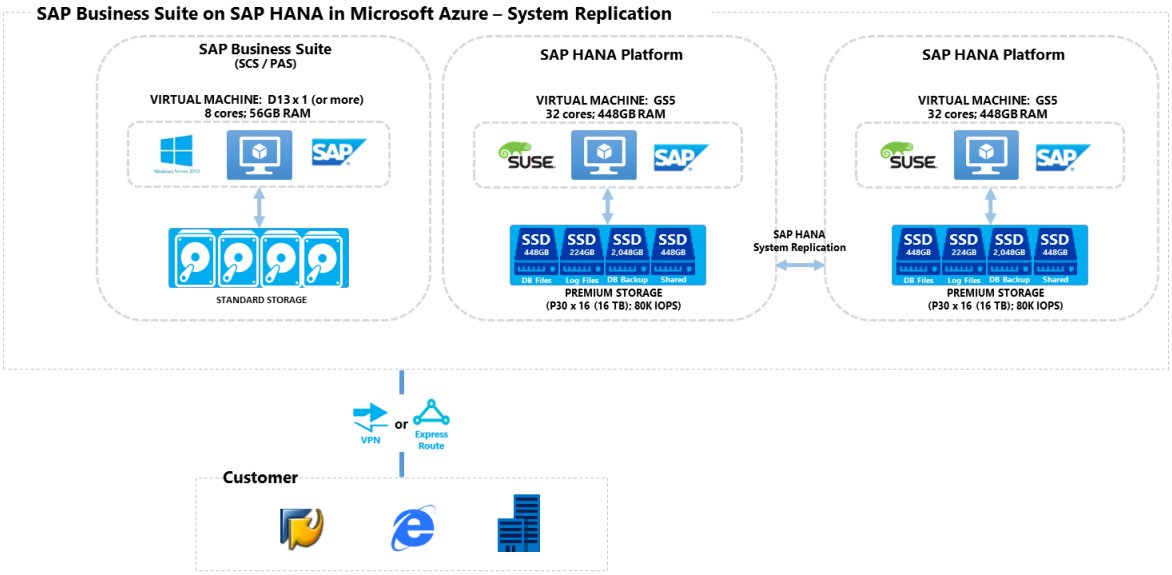

Figure 11: System replication architecture.

Storage

There are at least three types of Microsoft Azure storage relevant to SAP on Microsoft Azure. Specific to this whitepaper, this can be applied to SAP application servers running on Windows Server. For Linux-specific information, please see below.

Standard Storage

Microsoft Azure Blob Storage is suitable for almost all small and medium-sized SAP systems. Standard Microsoft Azure Blob Storage has four variants:

- **Locally redundant**: Keeps three data replicas within a single facility within a single region. Use this type of Storage Account for all SQL Server or DBMS storage.
- **Zone redundant**: Zone Redundant Storage (ZRS) stores 3 replicas of your data across 2 to 3 facilities. It is designed to keep all 3 replicas within in a single region, but may span across two regions.
- **Geo-redundant**: Data is replicated between fixed pairs of facilities asynchronously, keeping six replicas of the data. SAP application directories (such as the 'Transport' directory) can use this storage type. By default, GRS is selected and as such, it must be disabled for DBMS systems.
- **Geo-redundant Read Access**: A variation of geo-redundant storage that allows read-only access.

***Note**: Do not use Geo-Redundant storage accounts for DBMS workloads.*

Specific to SAP application servers, it's strongly recommended to install the \usr\sap\<SID> directory onto the file system (root) drive and not create a new Azure disk just for the SAP installation path.

Cache Settings for Standard Storage

Cache settings can be configured for each disk stored on a specific storage account. Write caching should be disabled for all disks on conventional Blob storage other than the OS boot disk and all virtual machine types supported for SAP applications can use standard Blob Storage. It's recommended to create disks at the maximum size (currently 1TB) as Microsoft Azure billing only charges for allocated and used space for Standard Blob Storage.

Non-Persistent Local Storage

Each Microsoft Azure virtual machine has a specific amount of local non-persistent (meaning certain operations will initialize the disk) storage. The amount of non-persistent storage depends on the virtual machine type and larger virtual machines generally have more non-persistent storage. This should be used for the swap space that SAP uses. See https://azure.microsoft.com/en-us/documentation/articles/virtual-machines-linux-suse-create-upload-vhd/ to setup swap space on the non-persistent local storage (via the Azure Linux Agent. Of note, the swap size recommended for SAP HANA on Linux is minimal – typically 2,048MB).

Premium Storage

Azure Premium Storage is backed by SSD disks and is currently only available on DS and GS series virtual machines. Leave the default "Read Only" cache set for most workloads except disks that are predominantly write intensive. Write intensive workloads such as the DBMS transaction log switch off caching altogether. Premium Storage is particularly useful for DBMS transaction logs.

For more information on Premium Storage see:

https://azure.microsoft.com/en-us/documentation/articles/storage-premium-storage-preview-portal/

Additional Storage Configuration Information

For more information about storage configuration specifics, please see the following information below.

- **Virtual machine performance optimization**: For configuration of Microsoft Azure storage for optimal performance with virtual machines, see: http://blogs.msdn.com/b/mast/archive/2014/10/14/configuring-azure-virtual-machines-for-optimal-storage-performance.aspx
- **Linux storage considerations**: Further Microsoft Azure and Linux storage considerations can be found at: http://blogs.msdn.com/b/igorpag/archive/2014/10/23/azure-storage-secrets-and-linux-i-o-optimizations.aspx
- **RAID on Linux**: To configure software RAID on Linux, see: https://azure.microsoft.com/en-us/documentation/articles/virtual-machines-linux-configure-raid/

- **Linux I/O optimizations**: For more Microsoft Azure storage tips including Linux I/O optimizations, see the following blog post: http://blogs.msdn.com/b/igorpag/archive/2014/10/23/azure-storage-secrets-and-linux-i-o-optimizations.aspx

---

***NOTE**: For deployment of SAP HANA on Linux, the best practice installation process for setting up software RAID on Linux is to use MDADM and not LVM.*

Sizing Considerations

SAP sizing and specifically, SAP sizing for Microsoft Azure is a complex subject. The following sizing checklist and SAP sizing resource recommendation will help you understand this topic in depth. In addition, please reference the whitepaper for Sizing SAP Solutions on Azure Public Cloud.

Sizing Checklist

A quick "checklist" regarding SAP sizing on Microsoft Azure is as follows:

1. **Catalog existing landscape**: Obtain accurate information about the current SAP landscape. For example, find the number of SAPS by running an ABAP report on the existing SAP instance. This can be compared to the respective SAP sizing support notes listed above.
2. **Database vs. application server SAPS**: Carefully separate database SAPS (for scale up scenarios only) from SAP application server SAPS (which can scale out). Dedicated database SAPS are typically less than 20% of the total number of SAPS. Do not directly map the total number of SAPS for the database/central instance to a dedicated database server on Microsoft Azure.
3. **Compare to SAP Note #1928533**: Check if the maximum number of database SAPS are larger than listed in SAP Note #1928533 - "SAP Application on Azure: Support Products and Azure VM Types".
4. **Plan and size Microsoft Azure Storage Accounts**: Consider each storage account a "virtual SAN" with a limited number of IOPS and throughput and also consider separate Storage Accounts for development, QA, and production environments for any sized landscape. Also consider dedicated storage accounts and/or Premium Storage for medium or larger systems - especially for QA and production.
5. **Carefully plan Microsoft Azure Virtual Networks**: Keep in mind that multiple network adapters inside a virtual machine do not increase the network throughput.
6. **Test on-premises to Microsoft Azure latency**: Use the Azure Speed Test for ExpressRoute maximum bandwidth and individual gateway bandwidth. Please note that Express Route Premium increases the routing table size.
7. **Right-sizing the environment**: On-premises deployments are sized with large amounts of buffer to allow an infrastructure to run for 3-5 years or more. However, Microsoft Azure deployments do not require a large sizing buffer as more resources can be added as the business volume increases such as when acquisitions occur or at peak load during year-end.

8. **Autoscale considerations**: Consider using Autoscale to run SAP application servers only when required and based on CPU utilization. Remember to set the Autoscale down increment to 0.

Sizing Methods

Specific to obtaining the information required to select the Microsoft Azure virtual machine for your SAP HANA instance, there are two methods for SAP sizing:

- **SAP HANA Quicksizer tool**: See - http://service.sap.com/quicksizer
- **Reference sizing**: Sizing that's based on the comparison of current SAP data (that then can be used in the following formulas and via execution of SAP ABAP reports) and other actual customer performance data and hardware configuration.

Sizing Resources

For more information about SAP for Azure sizing, please see the following information below.

- **SAP HANA Quicksizer tool**: For sizing SAP HANA, please use the SAP QuickSizer for HANA at:
  http://service.sap.com/hanaqs
- **SAP on Azure sizing blog**: For a general discussion of sizing SAP applications on the Microsoft Azure platform, see:
  http://blogs.msdn.com/b/saponsqlserver/archive/2015/06/19/how-to-size-sap-systems-running-on-azure-vms.aspx
- **SAP benchmarking and certification**: A detailed explanation of benchmarking and subsequent SAP certification is covered at:
  http://blogs.msdn.com/b/saponsqlserver/archive/2015/10/05/world-record-sap-sales-and-distribution-standard-application-benchmark-for-sap-cloud-deployments-released-using-azure-iaas-vms.aspx
- **VM type selection**: For an updated listing of supported SAP products and Azure virtual machine types, see:
  SAP Note #1928533

SAP Recommended Sizing Formulas

The following SAP recommended RAM and disk sizing formulas can be used on either a new SAP BW on SAP HANA system or when migrating to SAP HANA.

Table 7: SAP recommended RAM sizing formula.

SAP BW on SAP HANA – RAM Sizing Formula	
RAM	(Source data footprint - 60gb) * 0.5 * c1 + 90gb
Disk persistent	4 * RAM
Disk log	1 * RAM
c = the specific compression factor of the source database, when applicable.	

Table 8: SAP recommended disk sizing formula.

SAP BW on SAP HANA – Disk Sizing Formula	
Disk persistence	4 * RAM
Disk log	1 * RAM

For more information, see: SAP Note #1736976, "Sizing Report for BW on HANA"

***NOTE**: It is important to note that a more accurate sizing can be performed by running an SAP ABAP report.*

SAP HANA on Azure Storage Sizing

SAP HANA's storage requirements on Azure are covered in depth at:
http://scn.sap.com/docs/DOC-62595

Networking

With the deployment of resources into Microsoft Azure comes the need to provide connectivity, both between virtual machines deployed in the cloud, as well as back to on-premises services and users. The requirements for networking differ due to the requirements in a hybrid network scenario for a persistent connection to the customer's on-premises network.

In a Microsoft Azure-Only network scenario, normal operations involve deploying from a single

virtual machine (up to perhaps many hundred), solely within Microsoft Azure with no direct connectivity back to an enterprise network. In this scenario, these virtual machines would likely be connected to one or more Microsoft Azure Virtual Networks to allow cross-system communication (e.g. HTTP, RFC) and all interaction with those systems would happen through the opened TCP endpoints on the Azure virtual machine across the Internet.

In the Hybrid-IT scenario, Azure VPN or ExpressRoute services are designed to enable customers to maintain a persistent, private connection with their Microsoft Azure Virtual Networks and the resources deployed there. When the customer deploys one of these scenarios, they are able to communicate directly with their resources (virtual machines, Web roles, worker roles, Azure Storage etc.) and manage them as if they were an extension of their current network, much as many customers do today with remote data centers or colocation facilities.

When planning a Microsoft Azure Virtual Network, there are several aspects you should consider. For example, are you deploying in Microsoft Azure only a set of virtual machines, to be accessed using Remote Desktop or using SAPGUI across the Internet. Or are you going to connect this back to your network with VPN or ExpressRoute? How many Virtual Networks do you want or need, to segregate your virtual machines? What IP addressing range, subnets and gateways will you use? DHCP or static IP addressing? Beyond IP addresses, also to be considered are the needs of other network services, like DNS and Active Directory.

Domain Name Servers

Customers have the choice of providing their own DNS servers, and if they do not, the virtual machines will default to using Microsoft Azure DNS Servers.

Active Directory

In a Hybrid-IT scenario, Active Directory from on-premises can be extended to serve as the authentication mechanism through an Azure deployed domain controller (as well as potentially using the integrated DNS).

It is important to distinguish between traditional Active Directory Servers and Microsoft Azure Active Directory that provides only a subset of the traditional on-premises AD offering. This subset includes Identity and Access Management, but does not have the full AD schema or services that many 3rd party application take advantage of. While Azure Active Directory IS a requirement to establish authentication for the Azure virtual machines in use, and it can synchronize users with customers' on-premises AD, the two are explicitly different and customers will likely continue to require full Active Directory servers deployed in Microsoft Azure.

For an enterprise application such as SAP HANA, there are a few use cases where it can function in an isolated network environment (perhaps a training landscape or sandbox), however the reality is that an enterprise application such as SAP HANA requires network connectivity back to the enterprise for many reasons:

- SAPGUI / RFC traffic
- SAP transports

- SAP Solution Manager / OSS connectivity
- SAP System Interfaces
- Connectivity to other SAP solutions
- Printing
- Backups
- Monitoring

When connecting Microsoft Azure Virtual Networks back to the enterprise, there are two high-level options: VPN and ExpressRoute. Within those, the network topology utilized can vary, based on the customer requirements.

Common Network Scenarios

To illustrate key SAP HANA on Azure network considerations in perspective, the three most common network configurations are listed below.

- **Hybrid w/dev and QA**: Development and QA environments on Microsoft Azure with a Site-2-Site (S2S) VPN or ExpressRoute back to on-premises where a production environment is running has been implemented.
- **Hybrid w/dev, QA, and DR**: Development and QA environments on Microsoft Azure with S2S/ExpressRoute connecting to on-premises where a production environment is running. As well, production replicating to Microsoft Azure using Azure Site Recovery (ASR) for Disaster Recovery has been implemented.
- **Full cloud landscape**: Development, QA, production, and disaster recovery (DR) are all running on Microsoft Azure. S2S/ExpressRoute connects back to the internal corporate network to provide user connectivity.

Point to Site VPN

Point-to-site VPN means that client machines will connect through VPN directly into Azure, rather than routed through an Enterprise WAN. Regarding the deployment of an ERP Application like SAP, a connection by a single user into the system is an impractical solution and so, while there may be exception cases, it is recommended to disregard this option.

Site-2-Site VPN

This mechanism employs a Virtual Private Network (VPN) connection from either a single, or perhaps multiple (i.e. Multi-Site) locations in the customers' WAN into the Microsoft Azure virtual networks to allow the simplified routing of network traffic both from on-premises into Microsoft Azure and in return from Microsoft Azure, back to the resources on the customers' enterprise network.

VPNs are the simplest and most cost-effective connection type that can be employed to connect to Windows Azure, although they do have the drawbacks that their traffic travels directly over the Internet and therefore can suffer from some of the drawbacks of that shared Internet path, such as no guarantees on throughput, latency or the consistency of either.

There are many VPN devices which can be employed to connect these VPN endpoints, including a software-based connection utilizing Windows Server 2012 R2's Routing and Remote Access Service (RRAS).

VNet-2-VNet

It's also possible to create a VNet-2-VNet VPN. This kind of VPN can join two different Microsoft Azure Virtual Networks even if they are in different data centers. For example, this might be required when creating a Windows Cluster that spans from Singapore to Hong Kong regions.

For more information on VNet-2-VNet considerations, see:

- http://azure.microsoft.com/blog/2014/06/17/vnet-to-vnet-connecting-virtual-networks-in-azure-across-different-regions/
- https://msdn.microsoft.com/en-us/library/azure/dn690122.aspx.

ExpressRoute

Microsoft Azure ExpressRoute is the managed service which allows the private, dedicated connection between Azure data centers and the customer's Enterprise network and resources. ExpressRoute can establish private dedicated connections to Azure that do not travel across the open Internet in one of two ways, through an ExpressRoute partner or connect to Azure directly from the customer's network using technologies like MPLS (packet switched) network through a Network Service Provider.

This level of higher security and throughput while mitigating the issue of network latency sometimes seen with VPN circuits is often a requirement with Enterprise implementations to ensure operational availability and consistency.

Azure Speed Test

The performance of a Site-2-Site VPN can be approximated using these tools:

- http://azurespeedtest.azurewebsites.net
- http://www.azurespeed.com.

Testing considerations:

- Perform this test from a wired connection close to the core switch.
- Do not run this test on a wireless link to on a laptop on a user VLAN.
- Most customers achieve 5-60ms latency from their site to the nearest Microsoft Azure Region and if values are significantly higher than this, routing issues at the ISP may be present.
- Try testing using another ISP, even a residential fibre or DSL link can provide a useful comparison.

High Availability

The Microsoft Azure IaaS platform SLA is documented in detail and available for download. In summary, Microsoft commits to having virtual machines, configured in an availability set, available to external sources 99.95% of the time. Occasionally, the Microsoft Azure platform itself will be updated in order to provide more features and functionalities. There are ways to minimize the impact of planned Microsoft Azure restarts through various techniques including application-level high availability (HA).

System Replication

SAP HANA provides different methods of high availability with a range of recovery capabilities. There are three basic scenarios of HA:

- **Host auto-failover**: One (or more) standby nodes are added to an SAP HANA system and are in a 'standby mode' configuration (within a "scale-out" context).
- **Storage replication**: Actual storage replicates all of its data to another place (provided by hardware partners). Disks are mirrored without intervention from SAP HANA.
- **System replication**: SAP HANA itself replicates all of its data to a secondary SAP HANA instance which is an out-of-the-box, standard feature. Data is pre-loaded on the secondary system at all times to minimize downtime.

In the context of this whitepaper, Microsoft recommends SAP HANA System Replication for HA purposes.

For more system replication for SAP information and setup guidance, see:

http://scn.sap.com/docs/DOC-47702).

SAP Application Server Availability

The availability of SAP application servers is improved by configuring Autostart. In a scenario where a Microsoft Azure component fails and the platform self-heals and moves a virtual machine to another node, the impact of this restart is much less if the SAP application servers restart automatically. Autostart is configured by adding 'Autostart = 1' to the SAP default profile (default.pfl).

To achieve HA on Azure, there are additional steps required. The steps listed below are required as a minimum:

- **SAP ASCS and ERS availability**: Have two virtual machines as an SAP ASCS and ERS in their own Azure Availability Set.
- **SAP application availability**: Have two virtual machines running SAP application instances in their own Azure Availability Set.

Monitoring

For Windows Server virtual machines, the Microsoft Azure Enhanced Monitoring Extension for SAP is a Microsoft Azure extension that collects data from SAP systems through integration with SAP Kernel Components (SAPOSCOL/ SAP HostControl) installed on virtual machines running in Microsoft Azure, as well as host information about the virtual machine and makes it available inside SAP transactions as well as externally for support.

The Azure Enhanced Monitoring Extension for SAP is required for support from Microsoft and SAP, and thus must be installed in virtual machines running SAP for support purposes. The extension may be installed in virtual machines that are created in Microsoft Azure as well as virtual machines that are migrated from on premise environments to Microsoft Azure. The Microsoft Azure Enhanced Monitoring Extension for SAP makes use of the Microsoft Azure Extension Framework, and thus requires the Microsoft Azure virtual machine agent be installed as well. The virtual machine agent can be installed either during deployment of a virtual machine or after the virtual machine is deployed. Once the virtual machine agent is installed, the Microsoft Azure Enhanced Monitoring

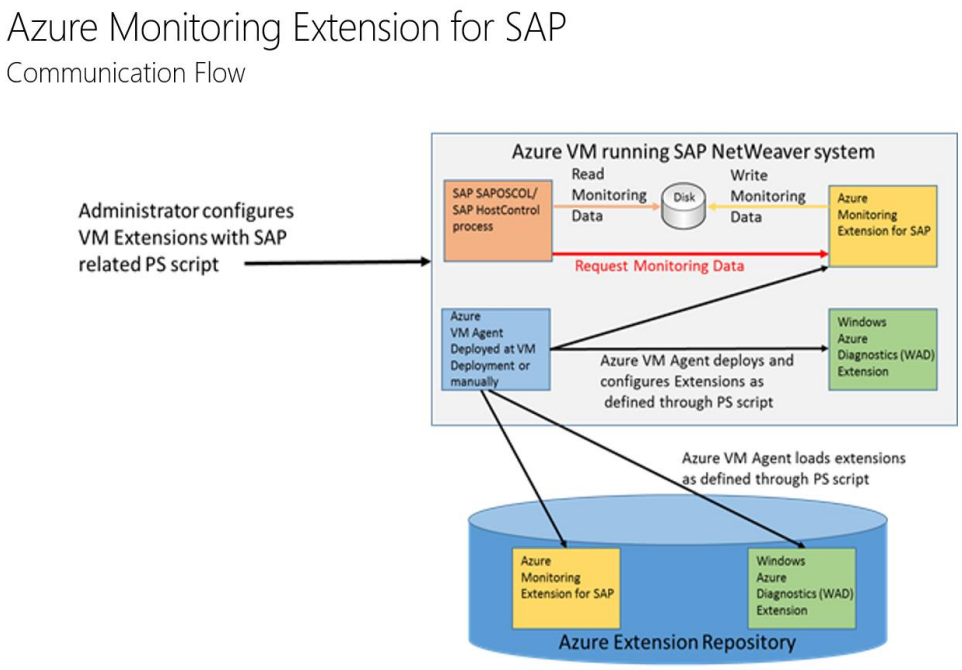

Extension for SAP can be installed through PowerShell.

Figure 12: Azure Monitoring Extension for SAP communication flow.

For additional information about the Azure Enhanced Monitoring Extension for SAP, please see: https://azure.microsoft.com/en-us/blog/azure-enhanced-monitoring-for-sap/.

For monitoring the Linux operating system, you can use the SUSE Manager combined with Microsoft System Center.

Migration of Existing SAP HANA into Azure

In order to migrate an existing SAP HANA system into Azure, a SAP homogeneous system copy is performed. A new virtual machine with SUSE Linux as the OS is deployed followed by installation of the specific SAP application being implemented. As part of the SAP Software Provisioning Manager (SWPM) installation process, an import is performed from the previous export of the SAP HANA database.

For more information regarding a homogeneous system copy, please see:

SAP Note #1844468 ("Homogeneous System Copy on SAP HANA").

Third-Party Tools

In order to access the SUSE Linux server, it's necessary to use a terminal program from a Windows client desktop. You can download 'PuTTY" for this purpose.

For more information on how to Log on to a virtual machine Running Linux, please see:

https://azure.microsoft.com/en-us/documentation/articles/virtual-machines-linux-how-to-log-on/.

Support

To ensure full support for the deployment of SAP systems on the Microsoft Azure Platform, SAP provides SAP Note #2015553 ("SAP on Microsoft Azure: Support Prerequisites").

Important points respective to ensuring full support:

- **Microsoft Premier Support**: Customer must have a Microsoft Premier Support contract in place.
- **Microsoft Azure Enhanced Monitoring Extension for SAP**: Must be implemented and reporting data into the SAP system.
- **SAP access**: Azure site-to-site VPN or Azure ExpressRoute is being used to access the SAP systems.
- **No 3-tier hybrid support**: No 3-tier SAP system should have its systems reside partly in Azure and partly on-premises in a way that splits the application and database servers due to inevitable latency issues across the networks.
- **Premium Storage for Database VHDs**: Virtual machines should go into Standard Storage and database server VHDs should go into Premium Storage.
- **Same storage account**: All VHDs for a specific SAP system should reside in the same Microsoft Azure Storage Account – multiple virtual machine VHDs can coexist in a single storage account.

- **Standard-Tier virtual machines**: Only Standard-Tier Microsoft Azure virtual machines are supported for SAP systems.
- **SAP HANA support**: Both production and non-production SAP HANA systems are supported.

Troubleshooting

Common troubleshooting links are listed below for your convenience:

- Storage account monitoring: http://azure.microsoft.com/en-us/documentation/articles/storage-monitor-storage-account
- Setup monitoring for storage account throttling on Azure platform: http://blogs.msdn.com/b/mast/archive/2014/08/02/how-to-monitor-for-storage-account-throttling.aspx
- Setup analytics logging: http://blogs.msdn.com/b/cie/archive/2013/10/10/storage-analytics-logging-how-to-enable-and-where-to-find-the-logs.aspx
- Troubleshooting Microsoft Azure Storage with Message Analyzer: http://azure.microsoft.com/blog/2015/01/27/troubleshooting-microsoft-azure-storage-with-message-analyzer
- Monitoring, diagnosing and troubleshooting Azure storage: http://azure.microsoft.com/blog/2014/09/25/monitoring-diagnosing-and-troubleshooting-azure-storage
- End-to-end troubleshooting: http://azure.microsoft.com/en-us/documentation/articles/storage-e2e-troubleshooting/#mma-analyze-data
- Monitor, diagnose and troubleshoot: http://azure.microsoft.com/en-us/documentation/articles/storage-monitoring-diagnosing-troubleshooting
- Troubleshooting using Azure Log Collector: http://azure.microsoft.com/blog/2015/03/09/simplifying-virtual-machine-troubleshooting-using-azure-log-collector
- Troubleshooting virtual machines: http://azure.microsoft.com/blog/2015/03/25/diagnostics-troubleshoot-remote-desktop-connections
- Get the reason for a virtual machine restart by yourself: http://azure.microsoft.com/blog/2015/04/01/viewing-vm-reboot-logs

- Azure gateway diagnostic: http://blogs.technet.com/b/keithmayer/archive/2014/12/18/diagnose-azure-virtual-network-vpn-connectivity-issues-with-powershell.aspx and https://gallery.technet.microsoft.com/scriptcenter/Azure-Virtual-Network-2b4d0793

Appendix

Of all you've read from the this 40-page whitepaper, the following point are the most important takeaway points to know when installing SAP HANA on Azure.

- From the Azure Gallery, use only SUSE Linux Server (SLES) 11 SP4 or SLES 12 SP1. (A special SUSE image is listed as being 'for SAP', but that is to be used only when using with the SAP Cloud Appliance Library.)
- All SAP HANA on Linux deployments must utilize the Azure Resource Manager deployment method.
- Never mount Azure data disks to a Linux server by using Device ID. Instead, use the Universally Unique Identifier (UUID). (Also be careful when using GUI tools when configuring disks as you must pay attention to this setting.) Verify the entries in the /etc/fstab file.
- Do not use LVM (Logical Volume Manager) as it isn't yet fully validated on Azure. That said, use MDADM (setting up RAID for large logical volumes needed for SAP HANA).
- Use the SAP recommended file system of XFS for SAP HANA. (See SAP Support Note #405827 – Linux: Recommended file systems.)
- For production systems, always use Azure premium storage.
- Make sure the Running SAP HANA on SLES SAP support note is followed (with all parameter changes performed, patterns/add-ons installed, etc.).
- Licensing: when SAP HANA performs a hardware check during installation (for example, during a homogeneous system copy installation via SWPM), it will install a temporary license (good for 90 days). A permanent license will need to be requested from the SAP Support Portal and applied to the new SAP HANA system deployed in Azure.
- Gnome desktop: if a GUI is desired for installing or managing SAP HANA on Linux in Azure, run **zypper in -t pattern gnome** (for SLES 11) or **zypper in -t pattern gnome-basic** (for SLES 12).

(Further guidance for running SAP HANA in Azure can be found at https://azure.microsoft.com/en-us/documentation/articles/virtual-machines-linux-sap-on-suse-quickstart/.