

Tales from Building a SQL Server Data Warehouse in Azure

Experiences & lessons learned from a migration to Azure

August 21, 2017

Lead Analytics Architect, SentryOne
Microsoft Data Platform MVP



Blog: www.sqlchick.com

Twitter: [@sqlchick](https://twitter.com/sqlchick)

Melissa Coates

Goals for This Session

- 1 Share **decisions and lessons learned** from a recent Azure implementation
- 2 Introduce **key concepts** for a deployment to Azure
- 3 Discuss items involved with **building a DW environment in Azure**

Azure services & features change frequently.

The information in the presentation is accurate as of mid-August 2017.

Agenda

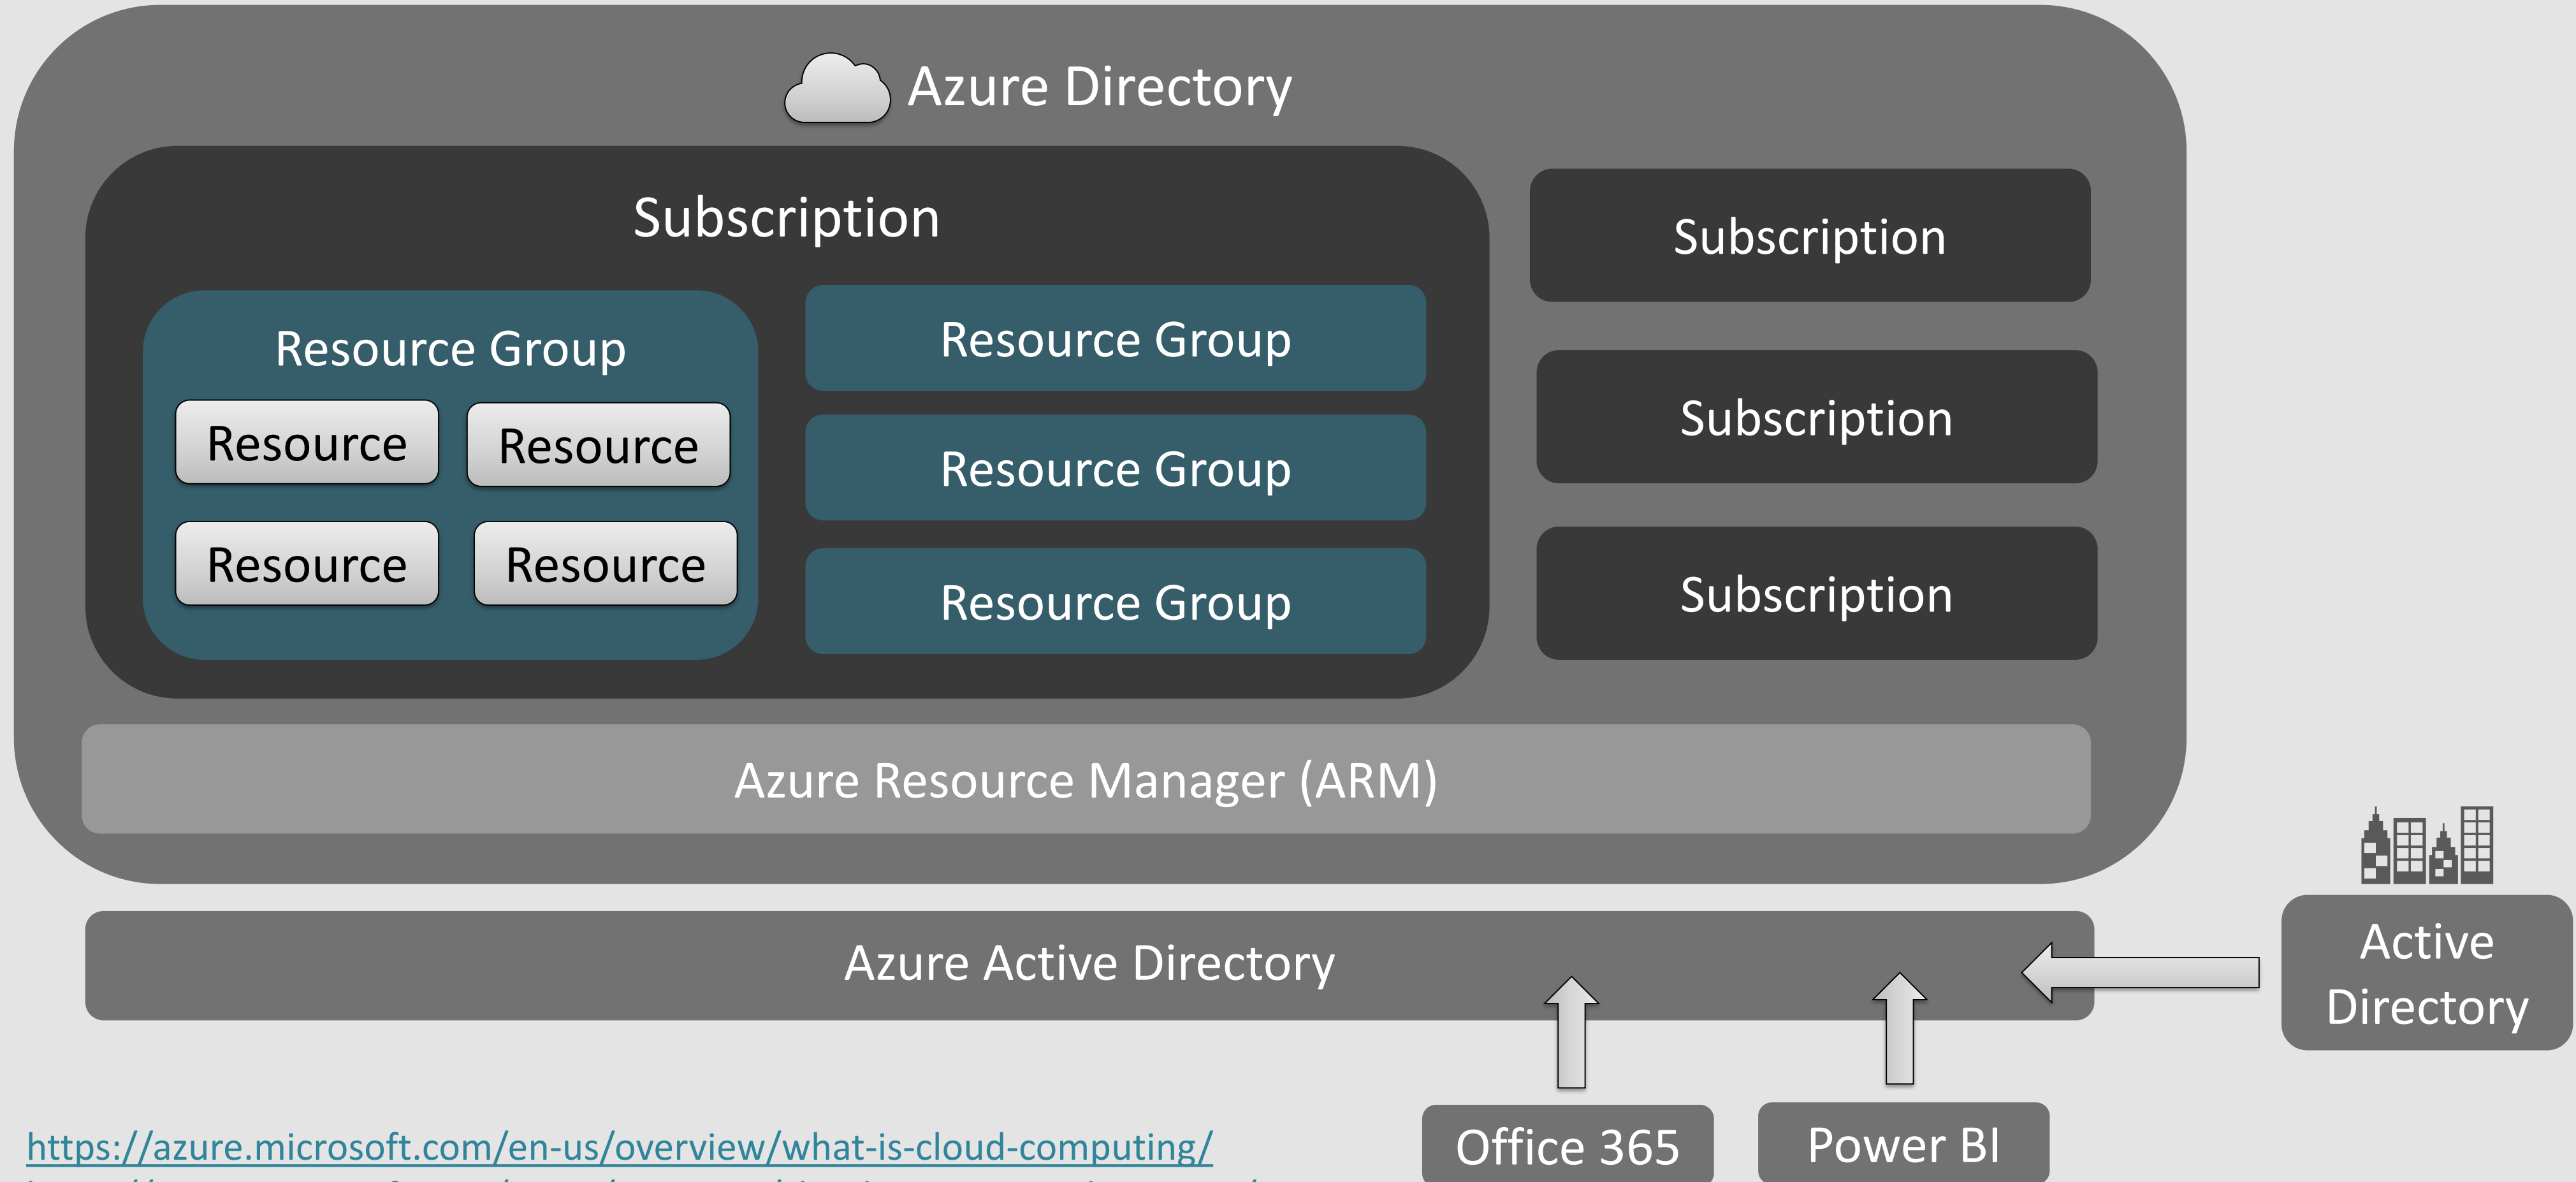
- ❑ Key Azure Concepts
- ❑ Deciding on Azure VM vs. Azure SQLDB vs. Azure SQLDW
- ❑ Additional Planning Considerations
- ❑ Building the Azure Environment
 - ❑ Structuring Dev, Test, & Prod Environments
 - ❑ Naming Conventions & Tags
 - ❑ Configuration Decisions
 - ❑ Automation & Scheduling
 - ❑ Monitoring the Solution
- ❑ Key Takeaways & Open Q&A

Out of scope:

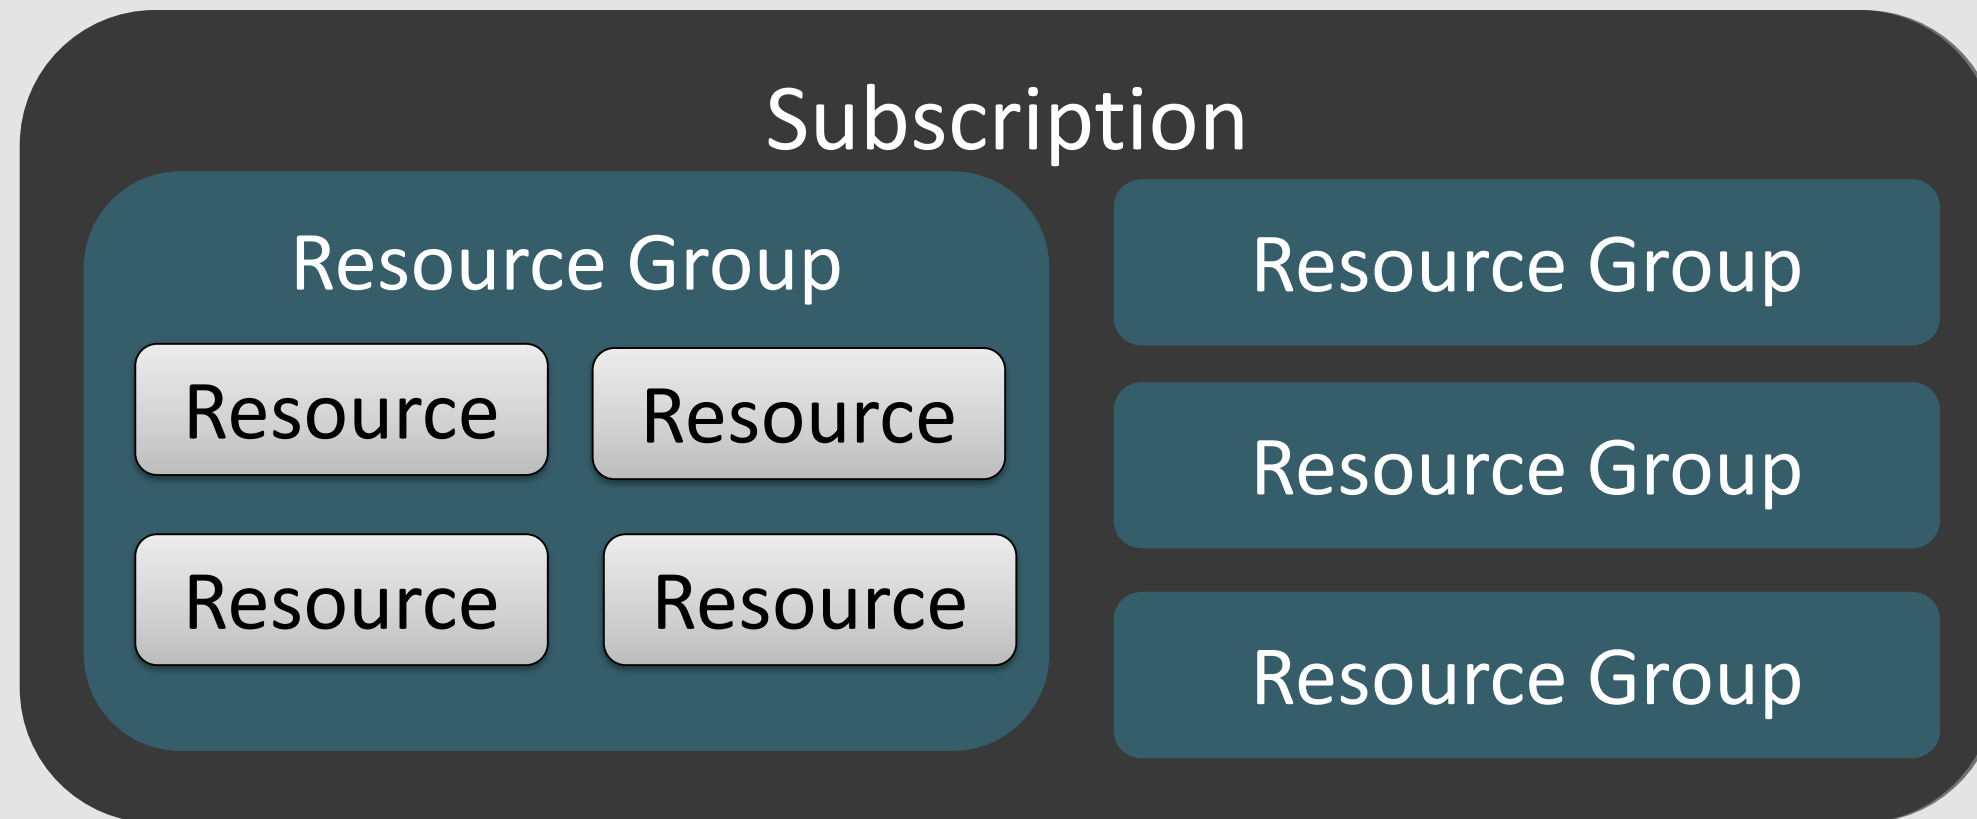
- Linux deployments
- Azure Stack & private cloud deployments
- On-premises physical & virtualized deployments (including Fast Track DW & APS/PDW)
- Security in depth (though we touch on a few points)
- Networking & connectivity
- Troubleshooting, performance tuning & growth
- Details on how to install and configure SQL Server

Key Azure Concepts

The Azure Lingo



Resource Groups



We have learned:

- ✓ Keep resource groups more narrow than broad
- ✓ Select the region (location) carefully

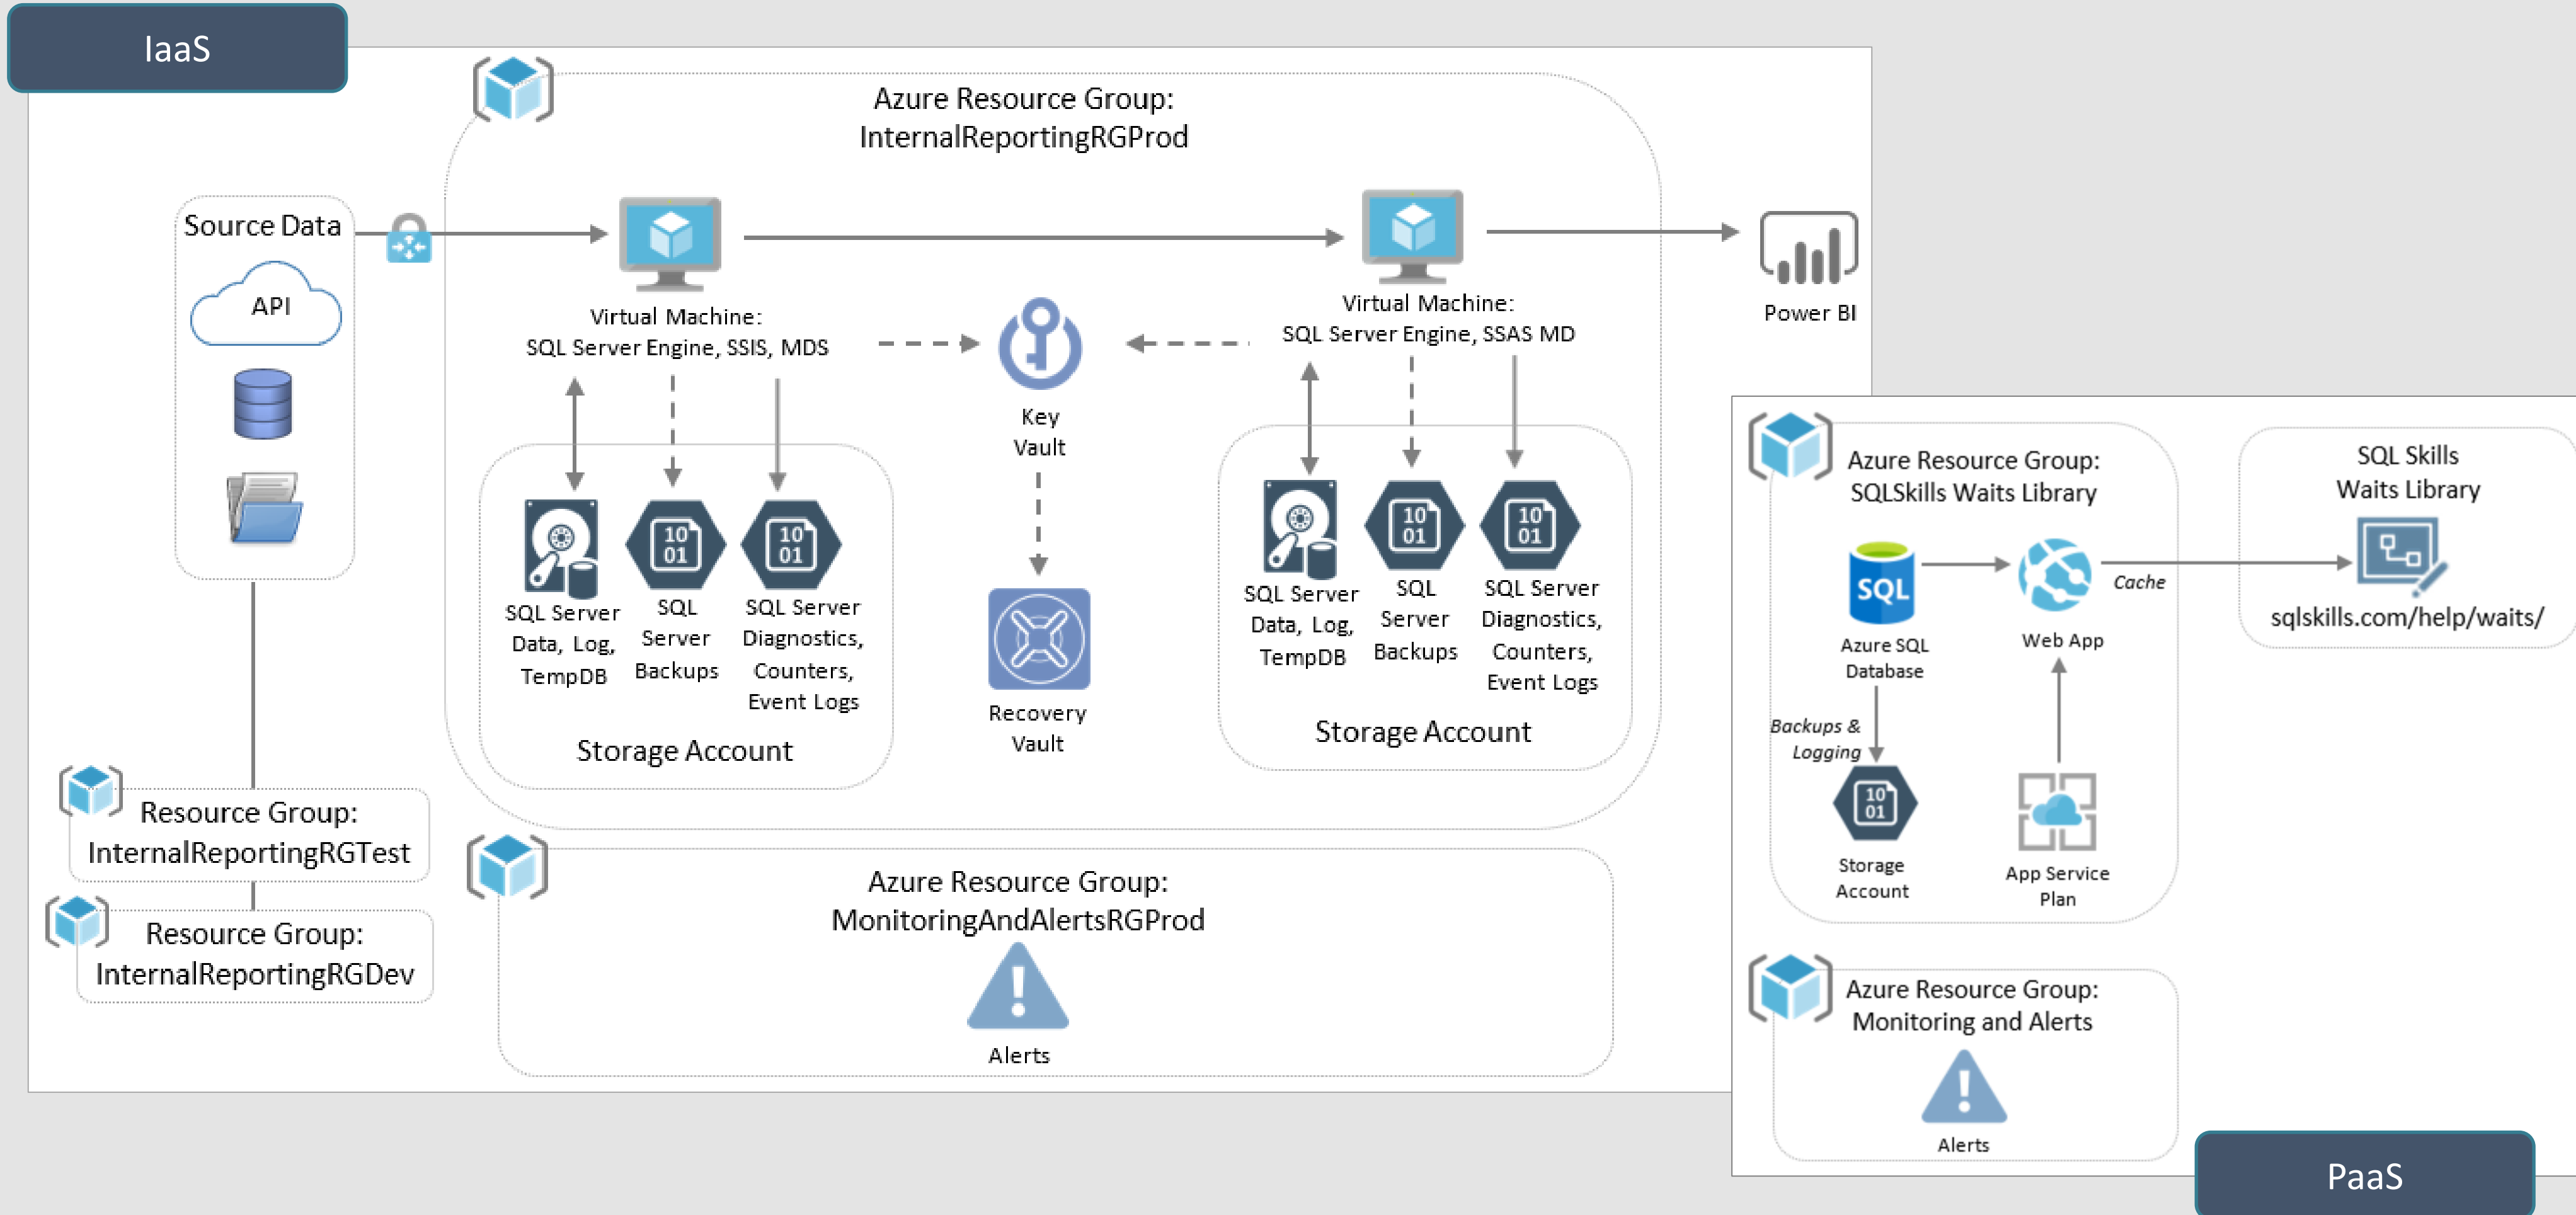
Planning for resource groups is critical

Focus on:

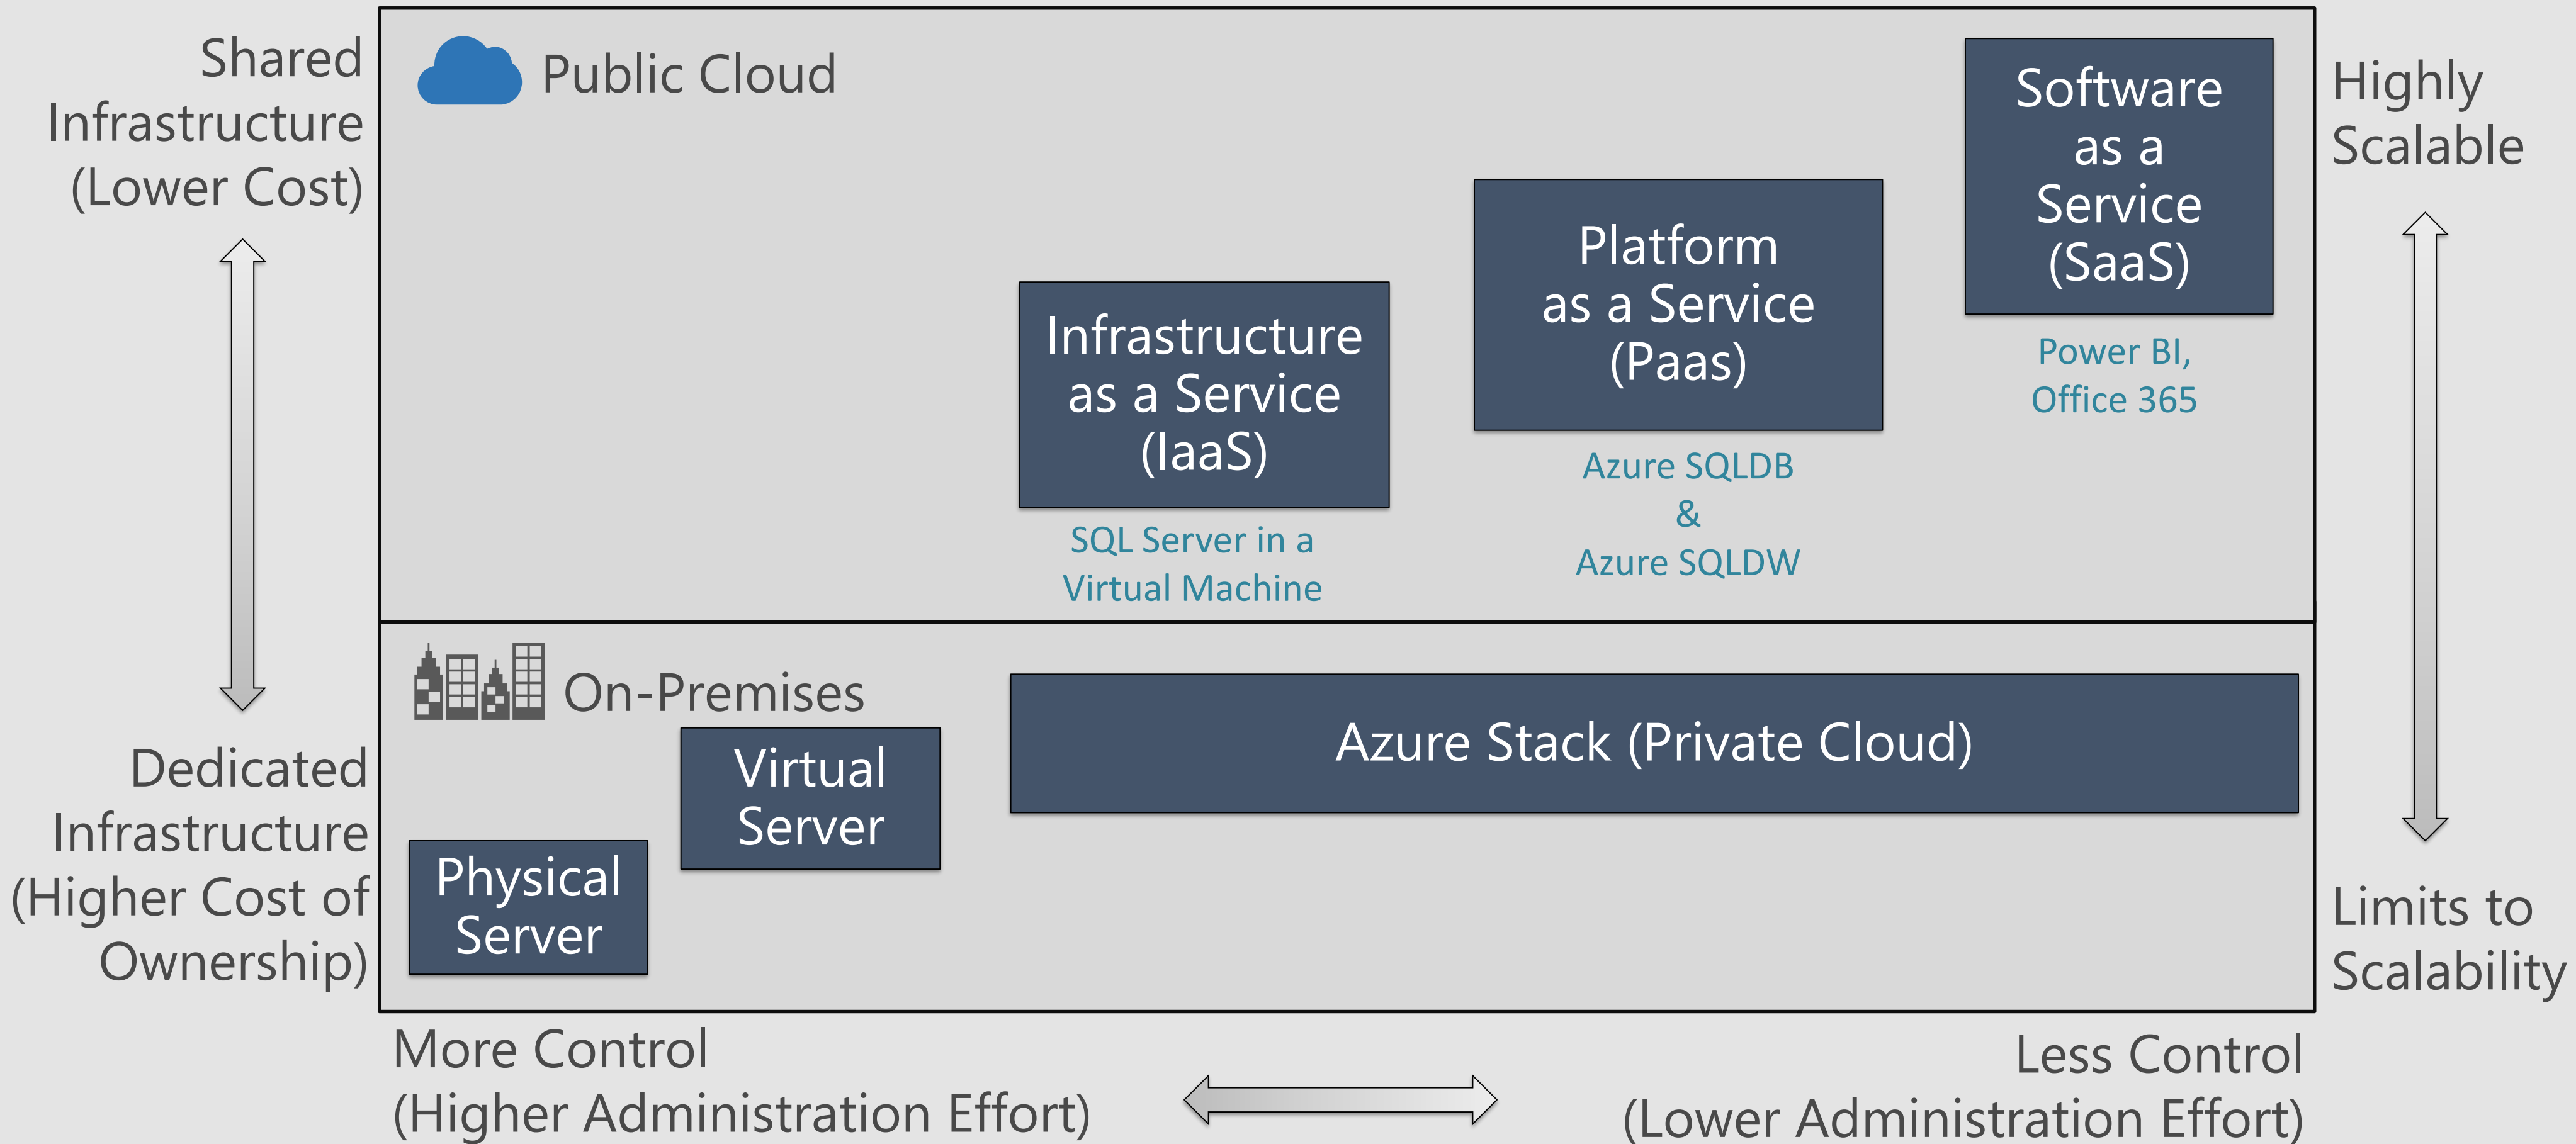
- Logical organization
- Permissions
- Policies

Scope of ARM automation scripts
(exception: Resource Explorer)

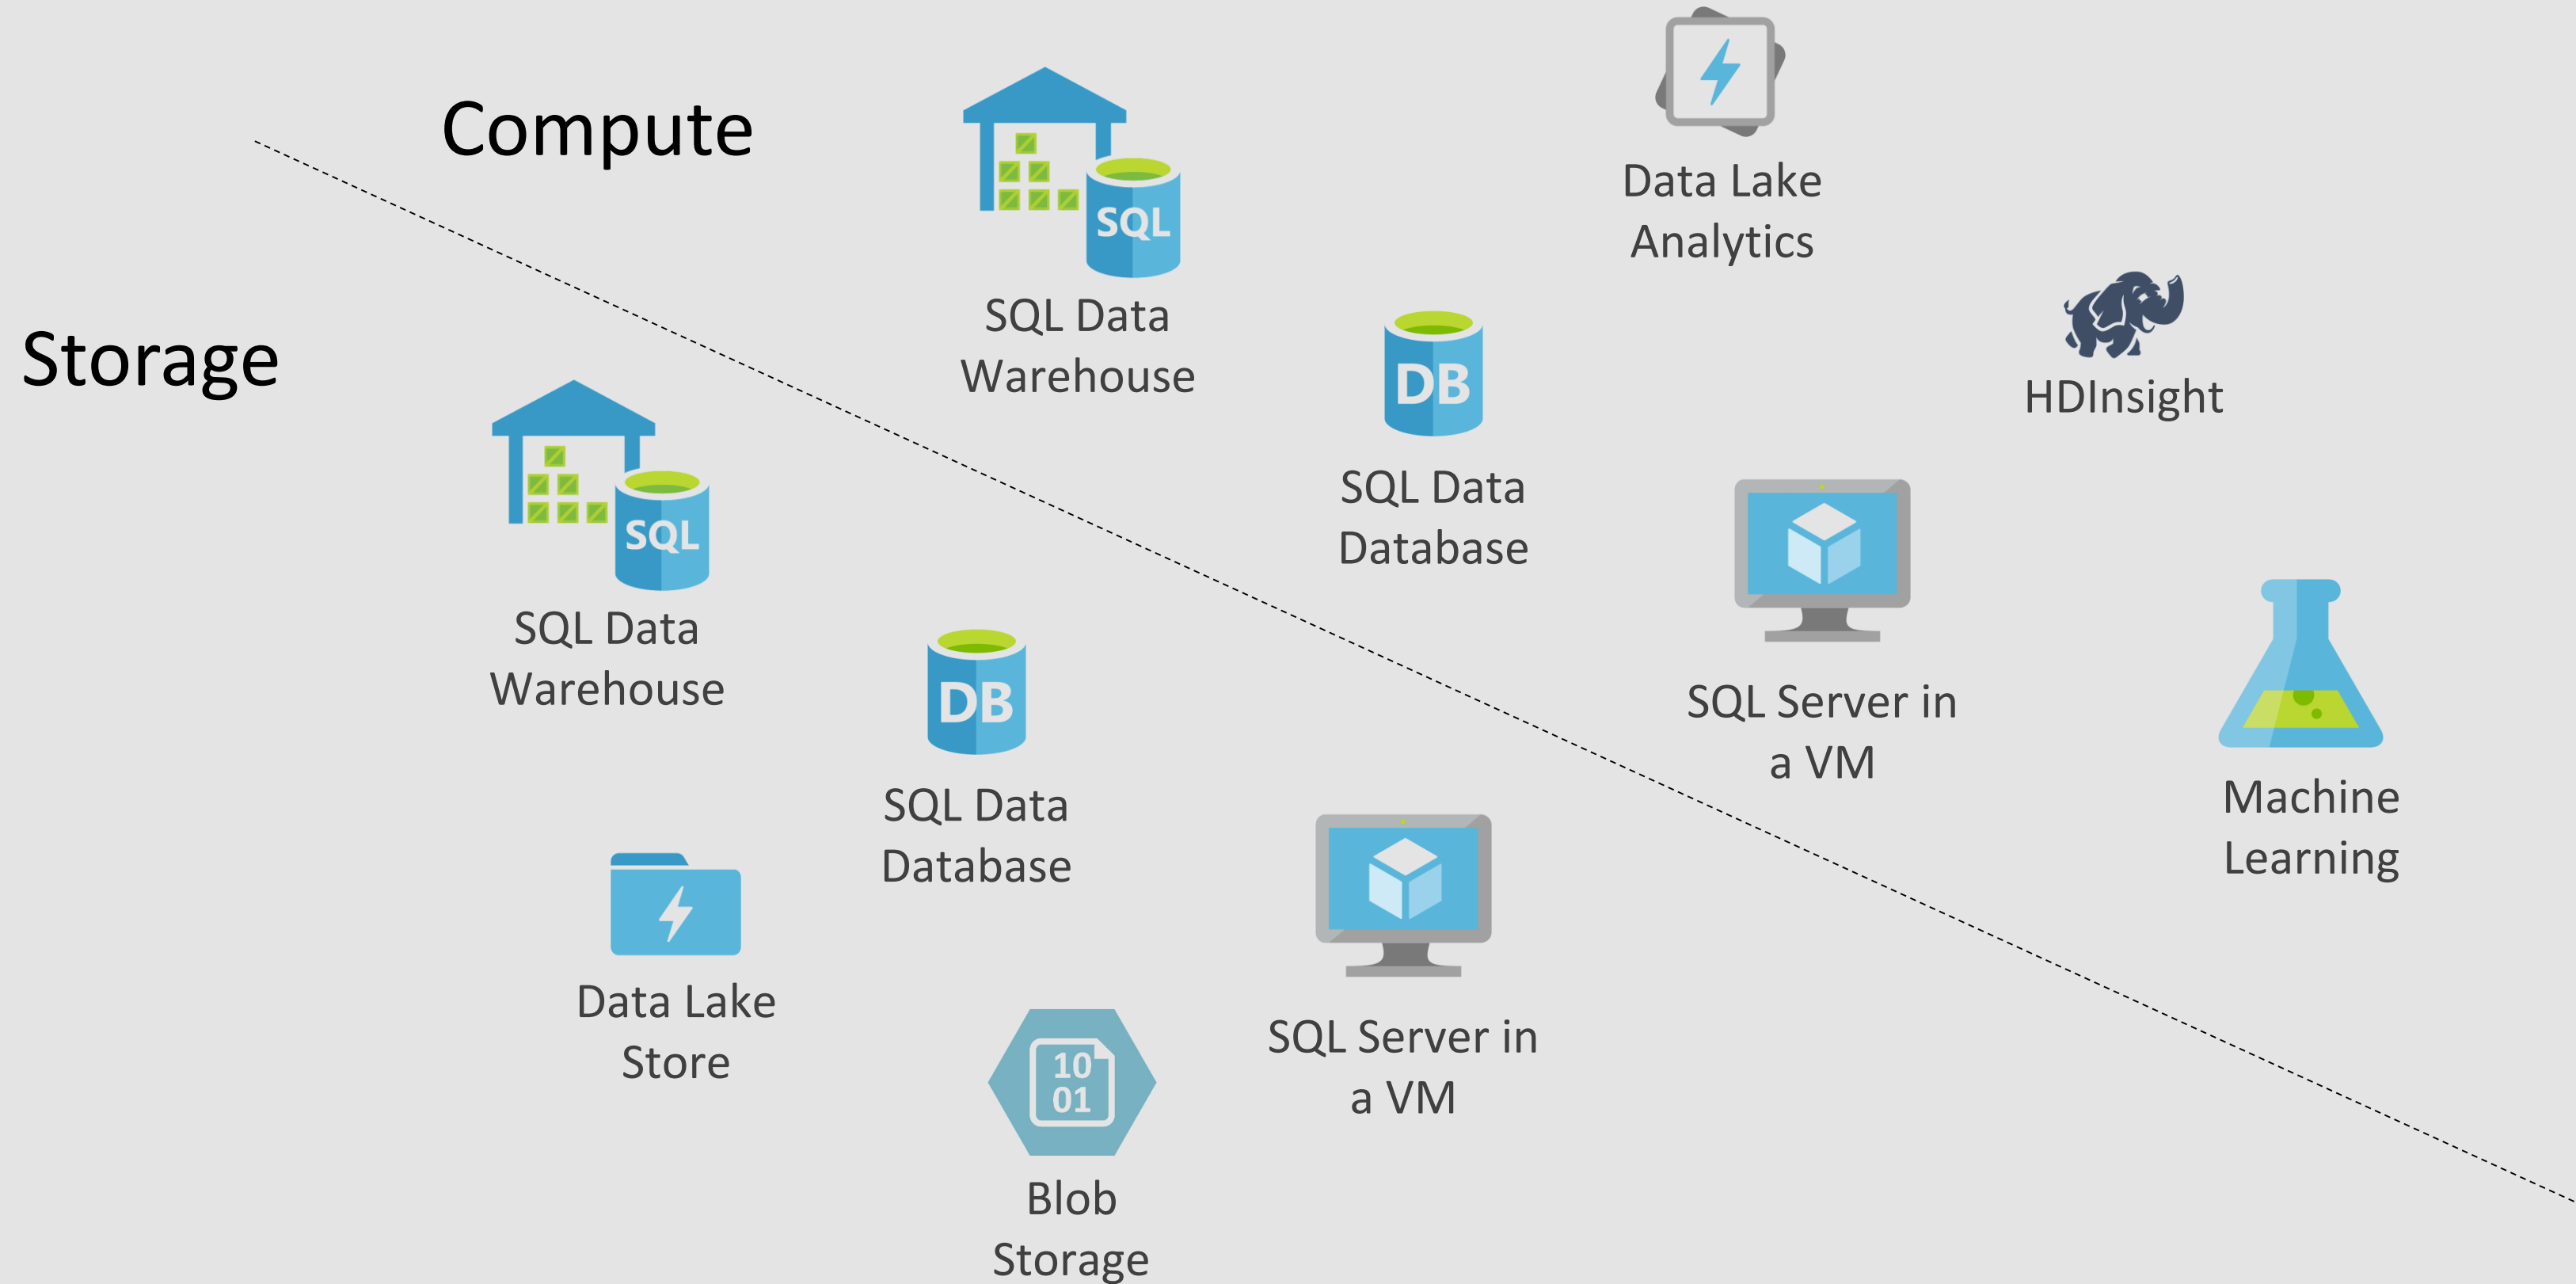
Examples – Resources & Resource Groups



IaaS vs. PaaS vs. SaaS



Compute vs. Storage



Some resources scale, or even pause, compute separately from storage.

Deciding on *Azure VM*
vs. Azure SQLDB
vs. Azure SQLDW

Comparing the SQL Offerings in Azure

(1/2)



SQL Server in a Virtual Machine
(IaaS)

Run full workload within an Azure virtual machine, including SQL Server, SSIS, SSAS, SSRS, etc



Azure SQL Database
(PaaS)

A relational database-as-a-service (DBaaS)

Non-Managed

A traditional Azure SQLDB deployment (isolated DB)

Managed Instance

Newer - closer feature parity to SQL Server (instance level features)



Azure SQL Data Warehouse
(PaaS)

An data warehouse-as-a-service (DWaaS) optimized for performance and large-scale distributed workloads

MPP architecture (massively parallel processing)

Comparing the SQL Offerings in Azure

(2/2)



SQL Server in a Virtual Machine (IaaS)

Best for:

- ✓ Migrating existing solutions
- ✓ Running any software and/or all SQL Server features
- ✓ Administering all aspects
- ✓ Bring your own license (Software Assurance)
- ✓ Isolated dev/test environments
- ✓ SLA: for the VM



Azure SQL Database (PaaS)

Best for:

- ✓ New database solutions
(If a DW it should be < 4TB data size--sharding across DBs is not suitable for DW workloads)
- ✓ OLTP with scaling & pooling needs (unpredictable workloads)
- ✓ Reduced administration of DB, OS, HA, and DR
- ✓ SLA: for the database



Azure SQL Data Warehouse (PaaS)

Best for:

- ✓ DW with larger data volumes (bare min. of 1-4TB)
- ✓ Ability to scale compute up/down, or pause (elasticity)
- ✓ Data Lake Store integration (relational + nonrelational data)
- ✓ Reduced administration
- ✓ SLA: for the database

Key Differences with Azure SQL Database

Many features go first to Azure SQLDB (“cloud first”). However, there are some key features not available in SQLDB (PaaS):

- PolyBase
- R Services
- Change data capture
- CLR
- DB snapshots
- Some T-SQL syntax
- Profiler
- Non-primary filegroups

Also, some features rely on Premium edition:

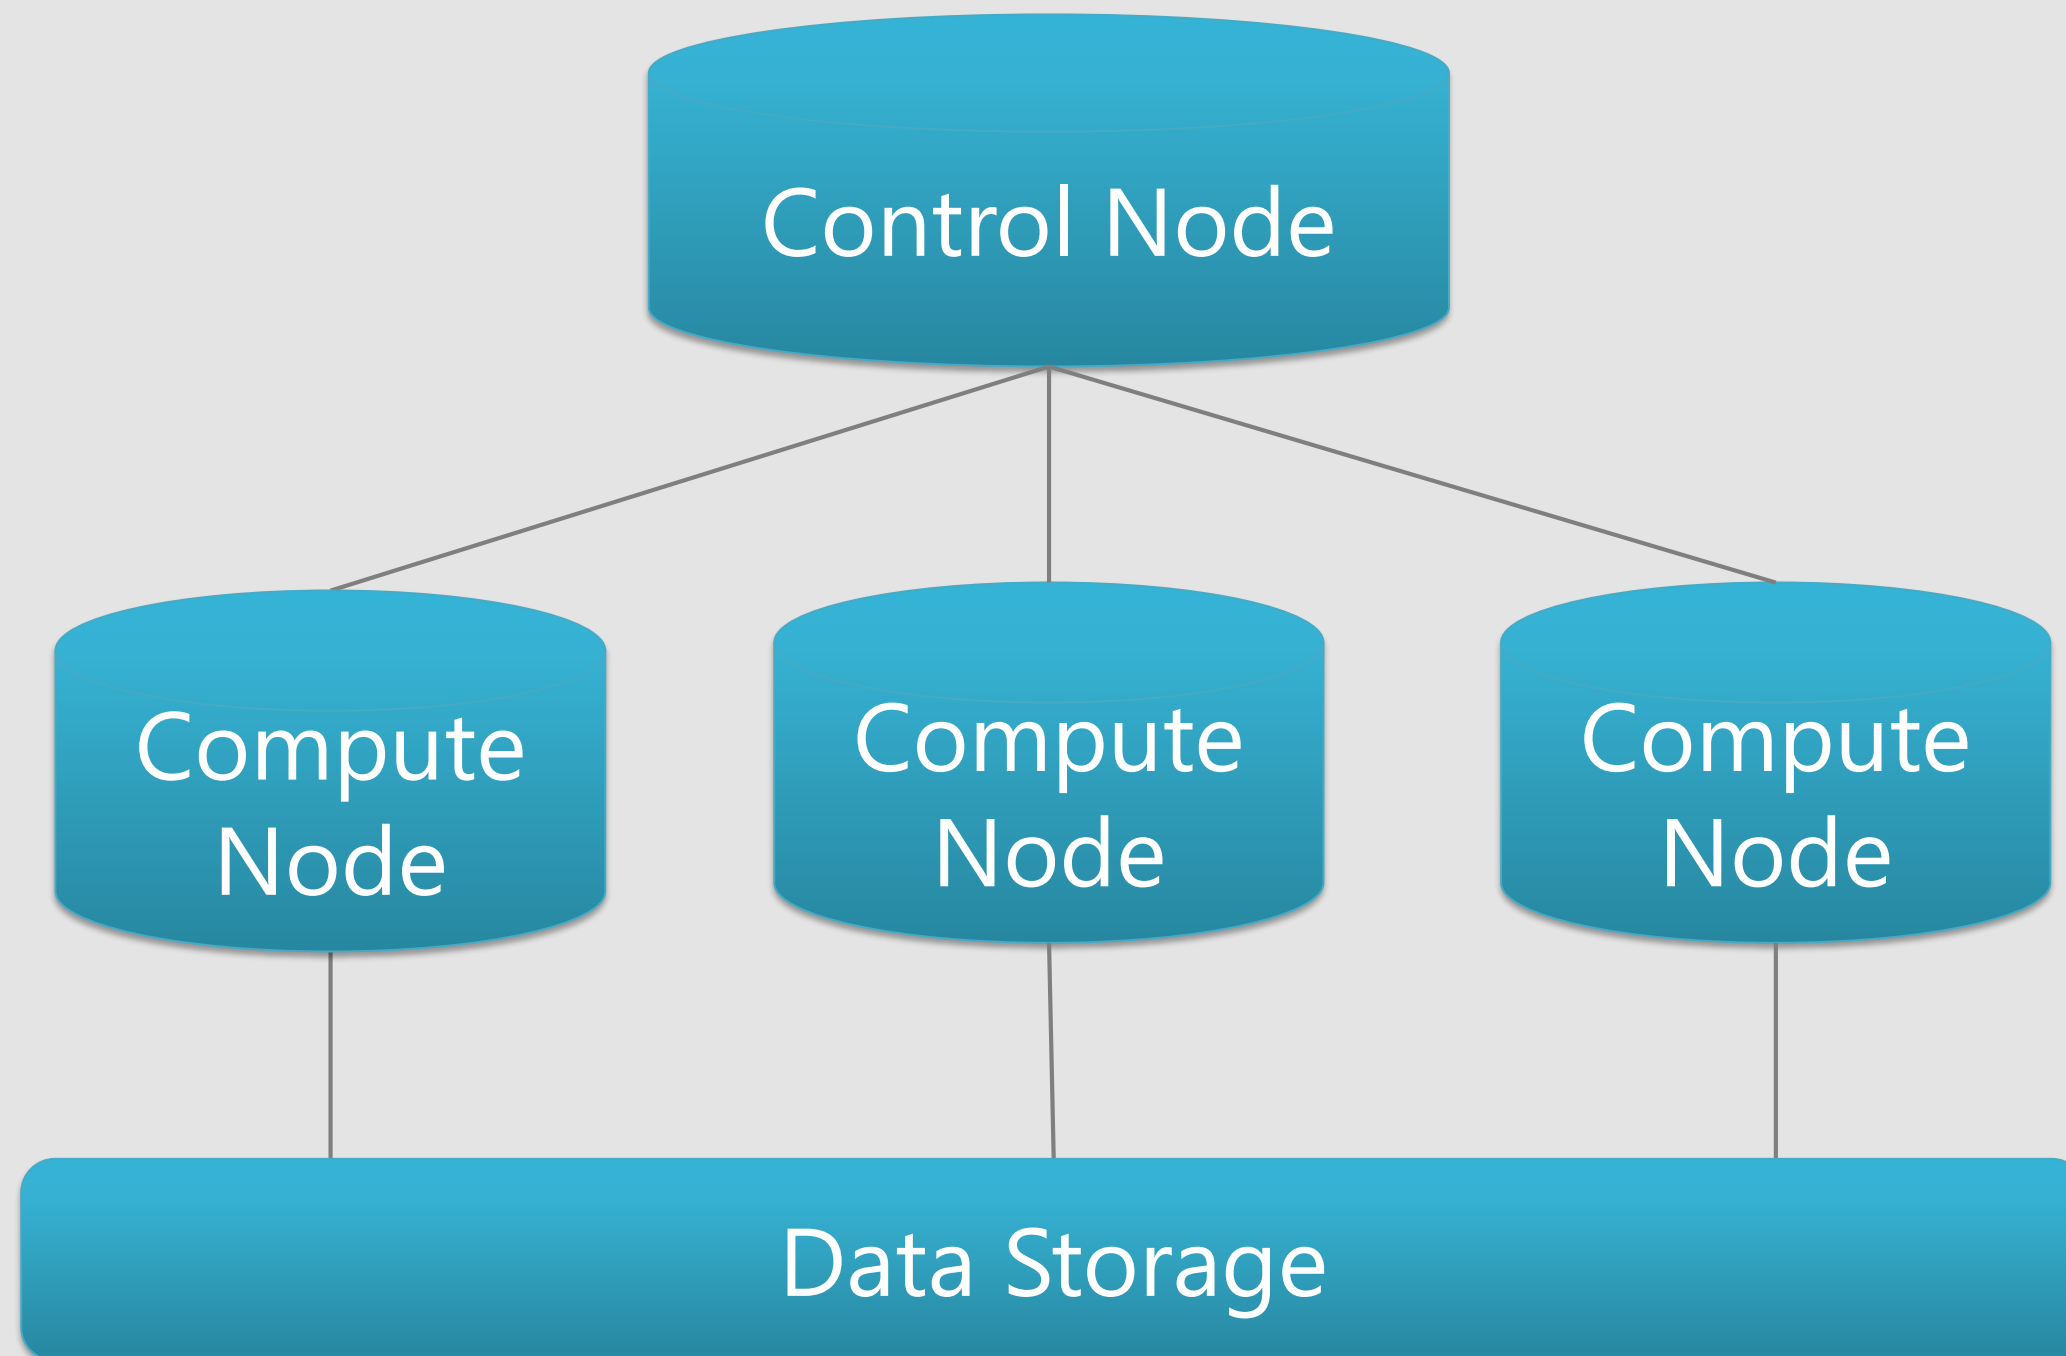
- Columnstore indexes

Feature	SQL Server	Azure SQL Database
Active geo-replication	Not supported - see Always On Availability Groups	Supported
Always Encrypted	Supported	Supported - see Cert store and Key vault
AlwaysOn Availability Groups	Supported	Not supported - See active geo-replication
Attach a database	Supported	Not supported
Application roles	Supported	Supported
Auto scale	Not supported	Supported - see Service tiers
Azure Active Directory	Not supported	Supported
Azure Data Factory	Supported	Supported
Auditing	Supported	Supported

Full list:

<https://docs.microsoft.com/en-us/azure/sql-database/sql-database-features>

Key Differences with Azure SQL Data Warehouse (1/2)



MPP Architecture

Shared-Nothing Architecture

Decoupled Storage & Compute

Scale Up, Down, Pause

PolyBase

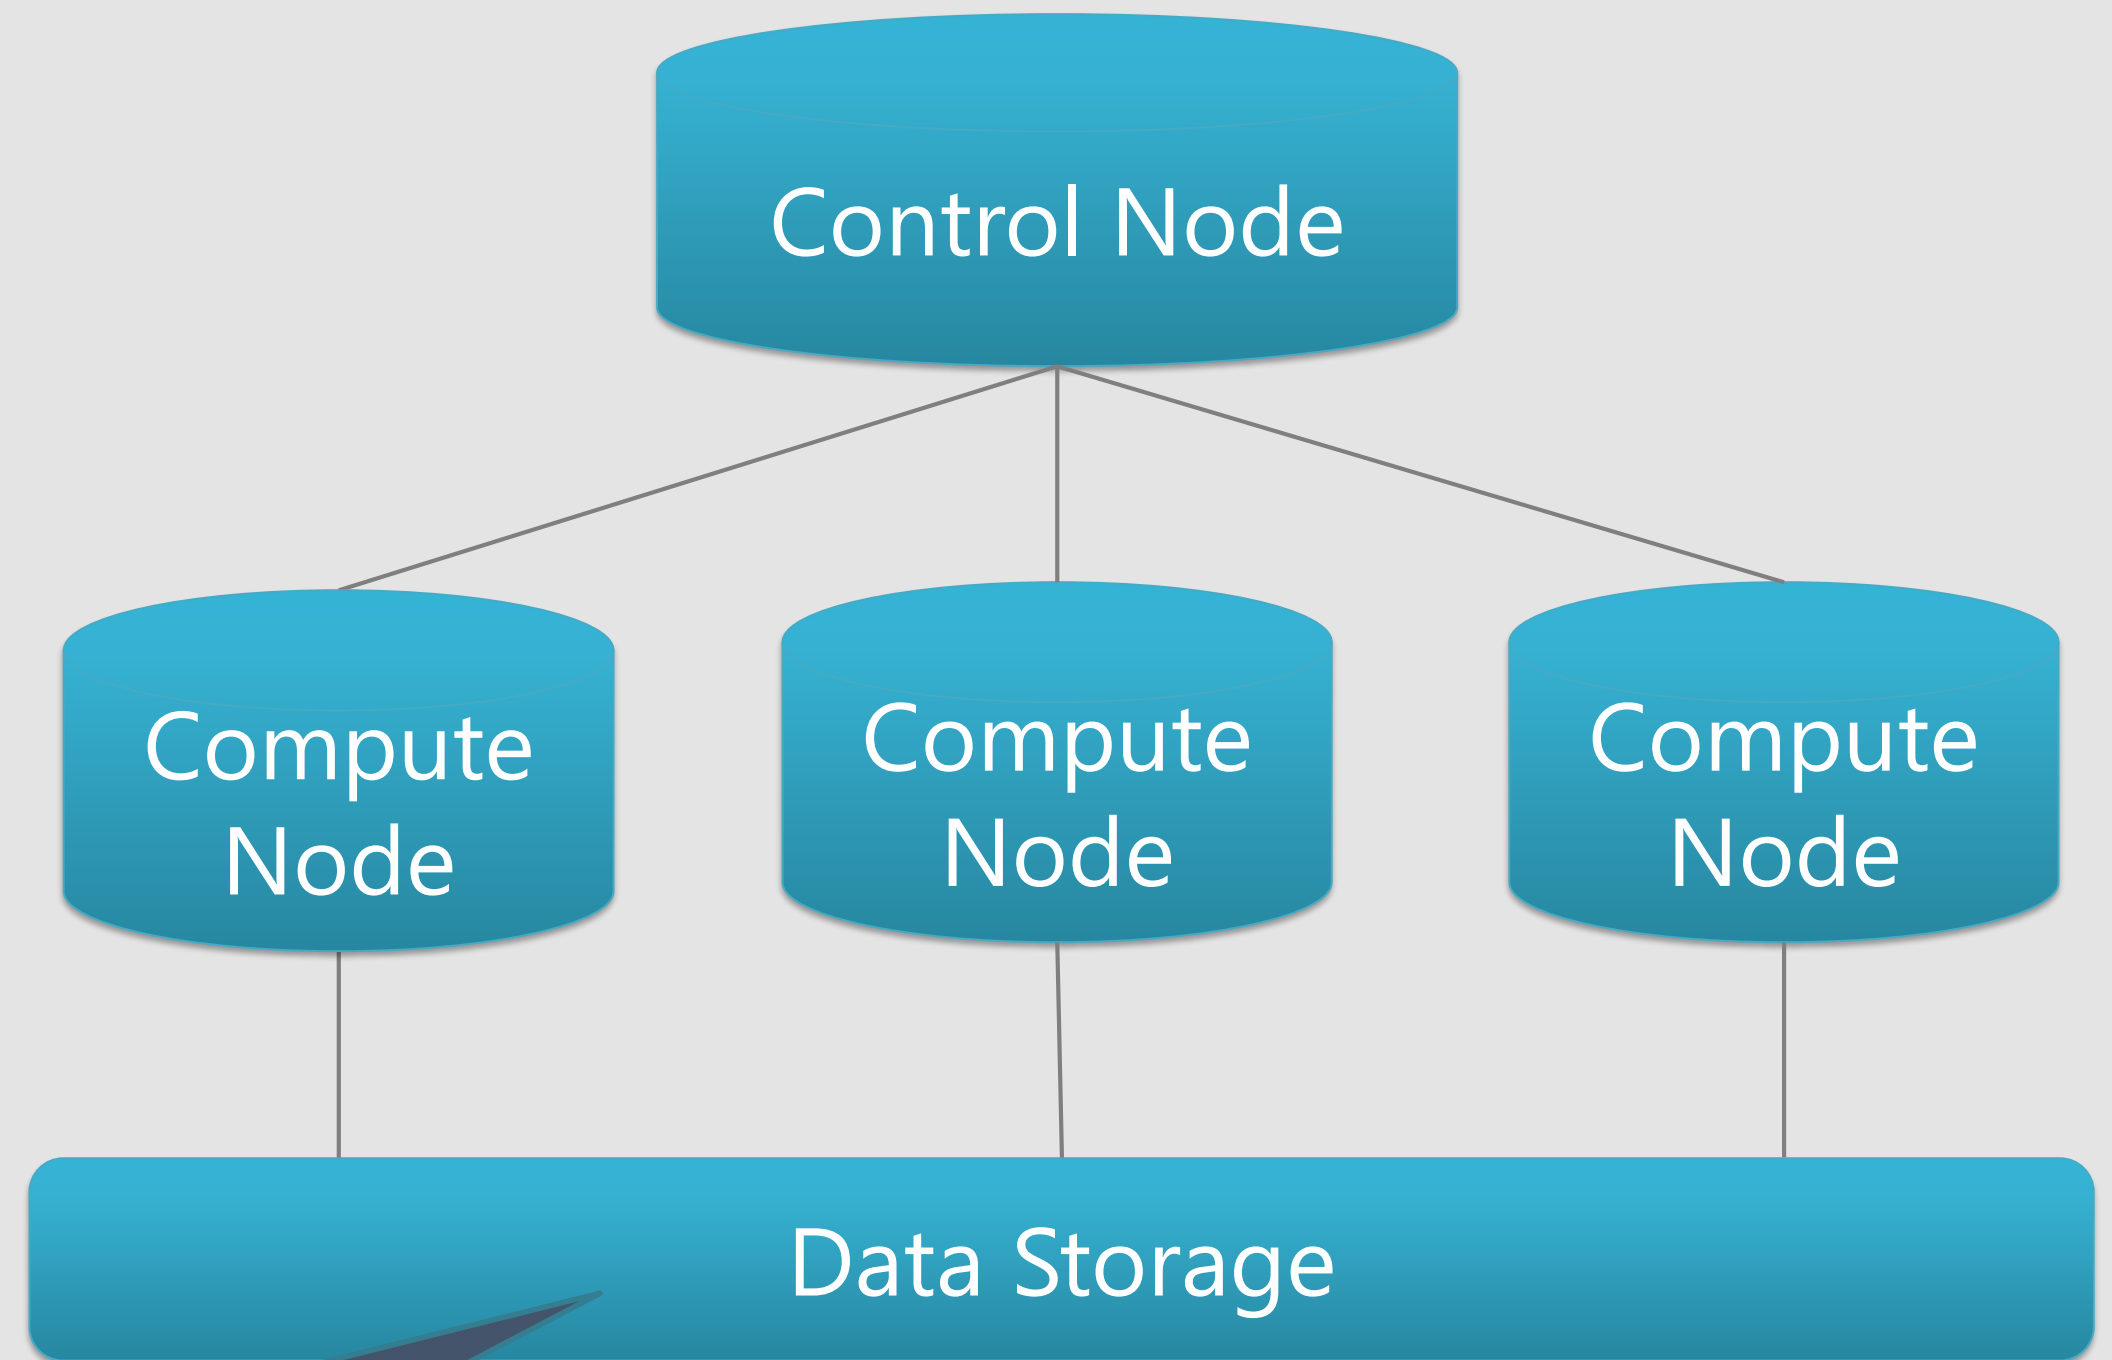
Key Differences with Azure SQL Data Warehouse (2/2)

Not All Features Supported

Different Data Loading Patterns

Distribution Keys are Critical

Denormalized Data Model is Best



Take time to educate yourself on the key differences with the MPP architecture—it will affect the design & the data load processes

Our Decisions on What to Use



SQL Server in a Virtual Machine (IaaS)

We are using a VM for:

- ✓ SQL Server DW
- ✓ Integration Services
- ✓ Analysis Services (MD)
- ✓ Master Data Services
- ✓ R Services

We need VMs for SSIS and SSAS anyway, so we couldn't justify migrating the relational DW to a PaaS solution at this time



Azure SQL Database (PaaS)

We are using SQLDB for:

- ✓ A specific use case: public reporting solution via SQLSkills Waits Library
- ✓ This SQLDB is loaded from the DW (in SQL Server)




Azure SQL Data Warehouse (PaaS)

We put SQLDW on roadmap:

- ✓ Future data growth
- ✓ Future PolyBase integration with multi-structured data in Azure Data Lake Store

Requires some refactoring of the data load processes so we are planning the move to SQLDW strategically

Our Current State



Additional Planning Considerations

Our Goals & Requirements for the Move to Azure

Expand infrastructure to support future growth

Ensure ability for Analytics Team to manage environment independently

Support existing solutions with little to no redesign or refactoring

Secure connectivity via VPN

Minimize cost where practical

Acceptable performance of hourly ETL jobs

Think about trade-offs you're willing to make for cost, performance, security, regulatory compliance, DR/backups/redundancy, and simplicity

Initial Planning Before Provisioning Any Resources (1/3)

First Steps

- Licensing and Editions
- Full Cloud Implementation vs. Hybrid (Partially On-Premises)
- IaaS vs. PaaS decisions; feature comparison

Capacity Planning & Cost Estimates

- Big 3: Storage, memory, CPU
- Networking
- Scalability needs
- Cost estimates: <https://azure.microsoft.com/en-us/pricing/calculator/>

High Availability & Disaster Recovery

- Down time sensitivity (RPO, RTO)
- SLAs from Azure

Compliance & Security

- Compliance: <https://www.microsoft.com/en-us/trustcenter/compliance/default.aspx>

Initial Planning Before Provisioning Any Resources (2/3)

Domain service
accounts &
credentials

- How administrative & owner permissions will work if decentralized
- Read and/or write permissions for source or related systems
- One domain service account per service, per environment
- Sync to Azure Active Directory for domain users & groups
- Service principals for certain resources in Azure are very important

Azure region:
primary location
&
data redundancy

- Geographic location of data
- Proximity to business users
- Co-location of related resources
- Minimizing latency
- Minimizing data egress charges (very inexpensive though)
- Not all resources/services are available in every region

Initial Planning Before Provisioning Any Resources (3/3)

Migration Method

Backup/
Restore

Upload
a VHD

Fail Over from
AlwaysOn

Replication

Ship a
Hard Drive

Database
Migration Service

We went with a backup/restore approach for the SQL Server files for the DW.

The exception to this was SSISDB custom auditing & logging objects. This DDL was deployed from SSDT after the SSISDB catalog was configured.

Info on migration techniques:

<https://docs.microsoft.com/en-us/azure/virtual-machines/windows/sql/virtual-machines-windows-migrate-sql>

Azure SQL Data Warehouse migration utility:

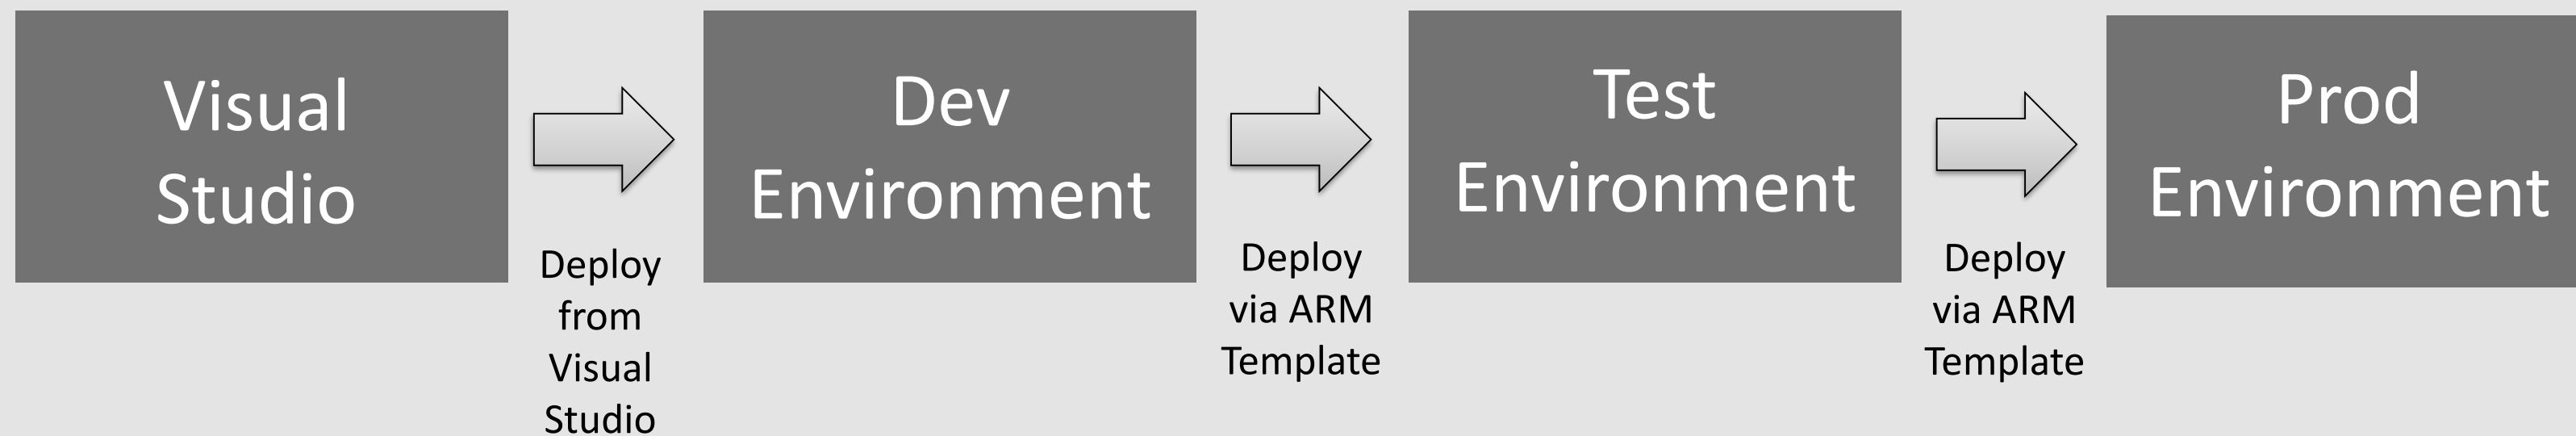
<https://docs.microsoft.com/en-us/azure/sql-data-warehouse/sql-data-warehouse-migrate-migration-utility>

Structuring Dev, Test & Prod Environments

Separation of Dev, Test, Prod Environments

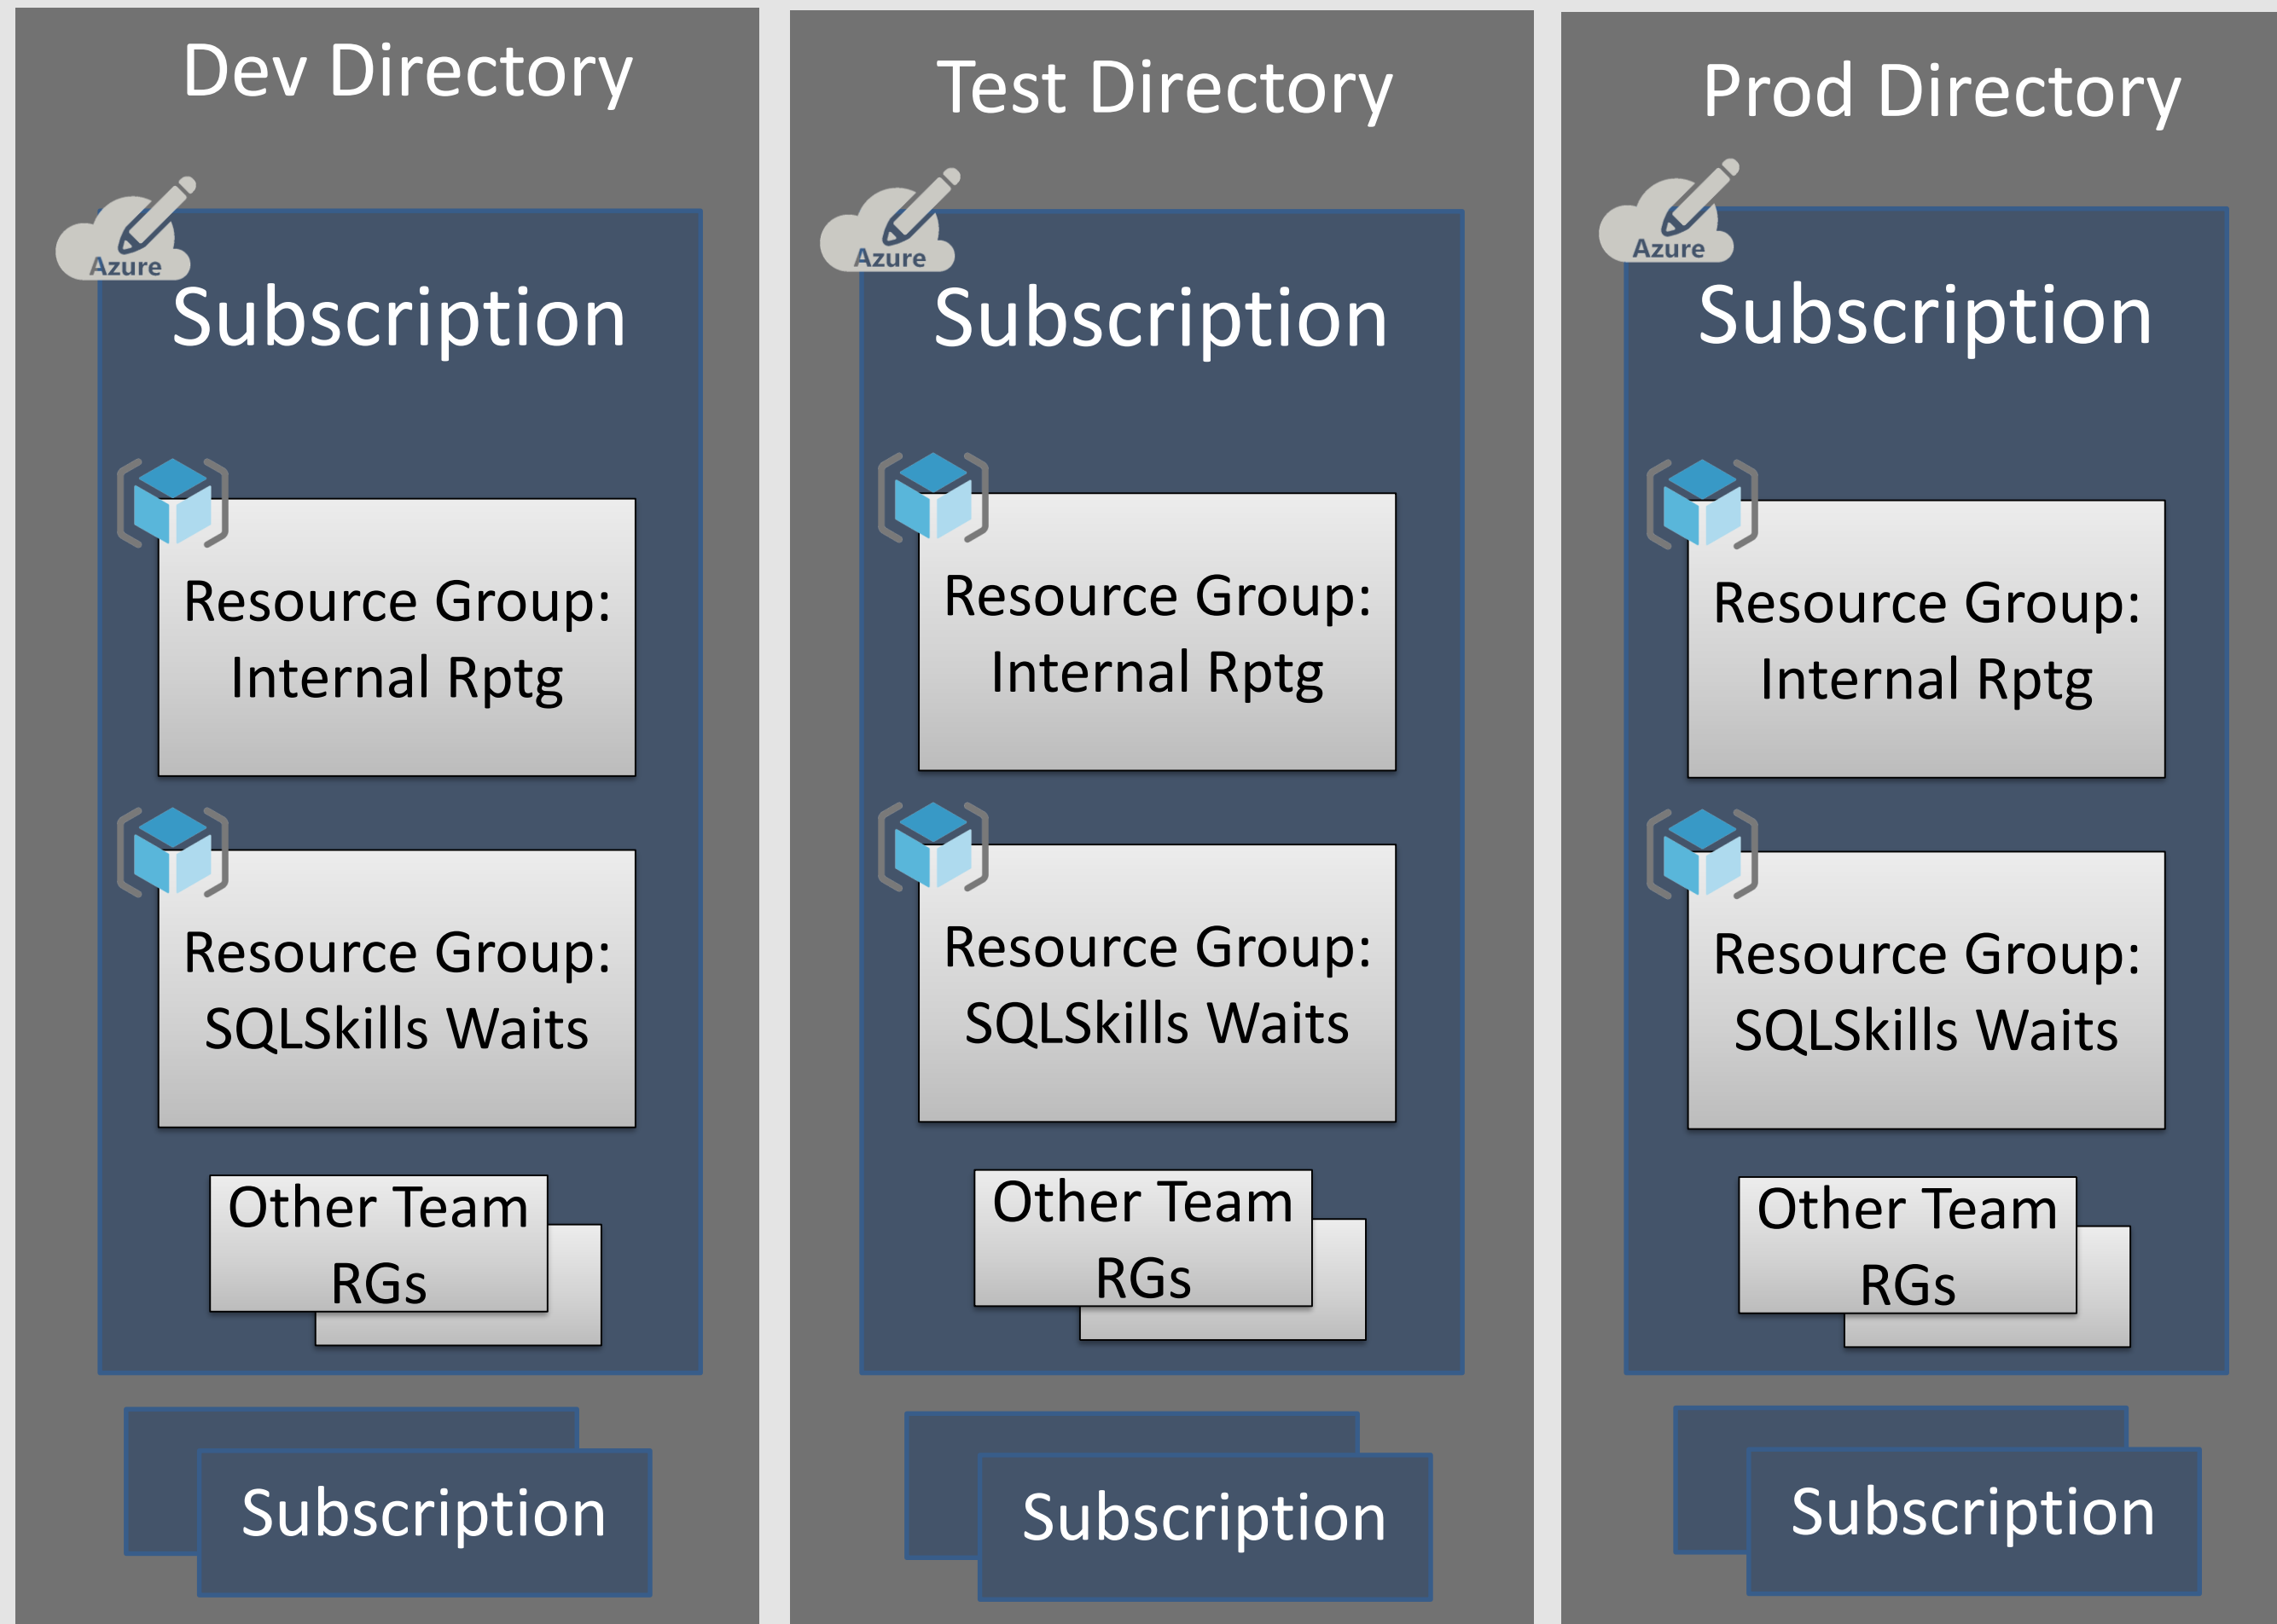
Most commonly environments are segregated by:

1. Resource Groups,
2. Subscription, or
3. Directory, or
4. A combination of 1 and 3, or 2 and 3



For large or multi-tenant implementations: be aware of Azure limits before deciding.

Option: Separate By Directory



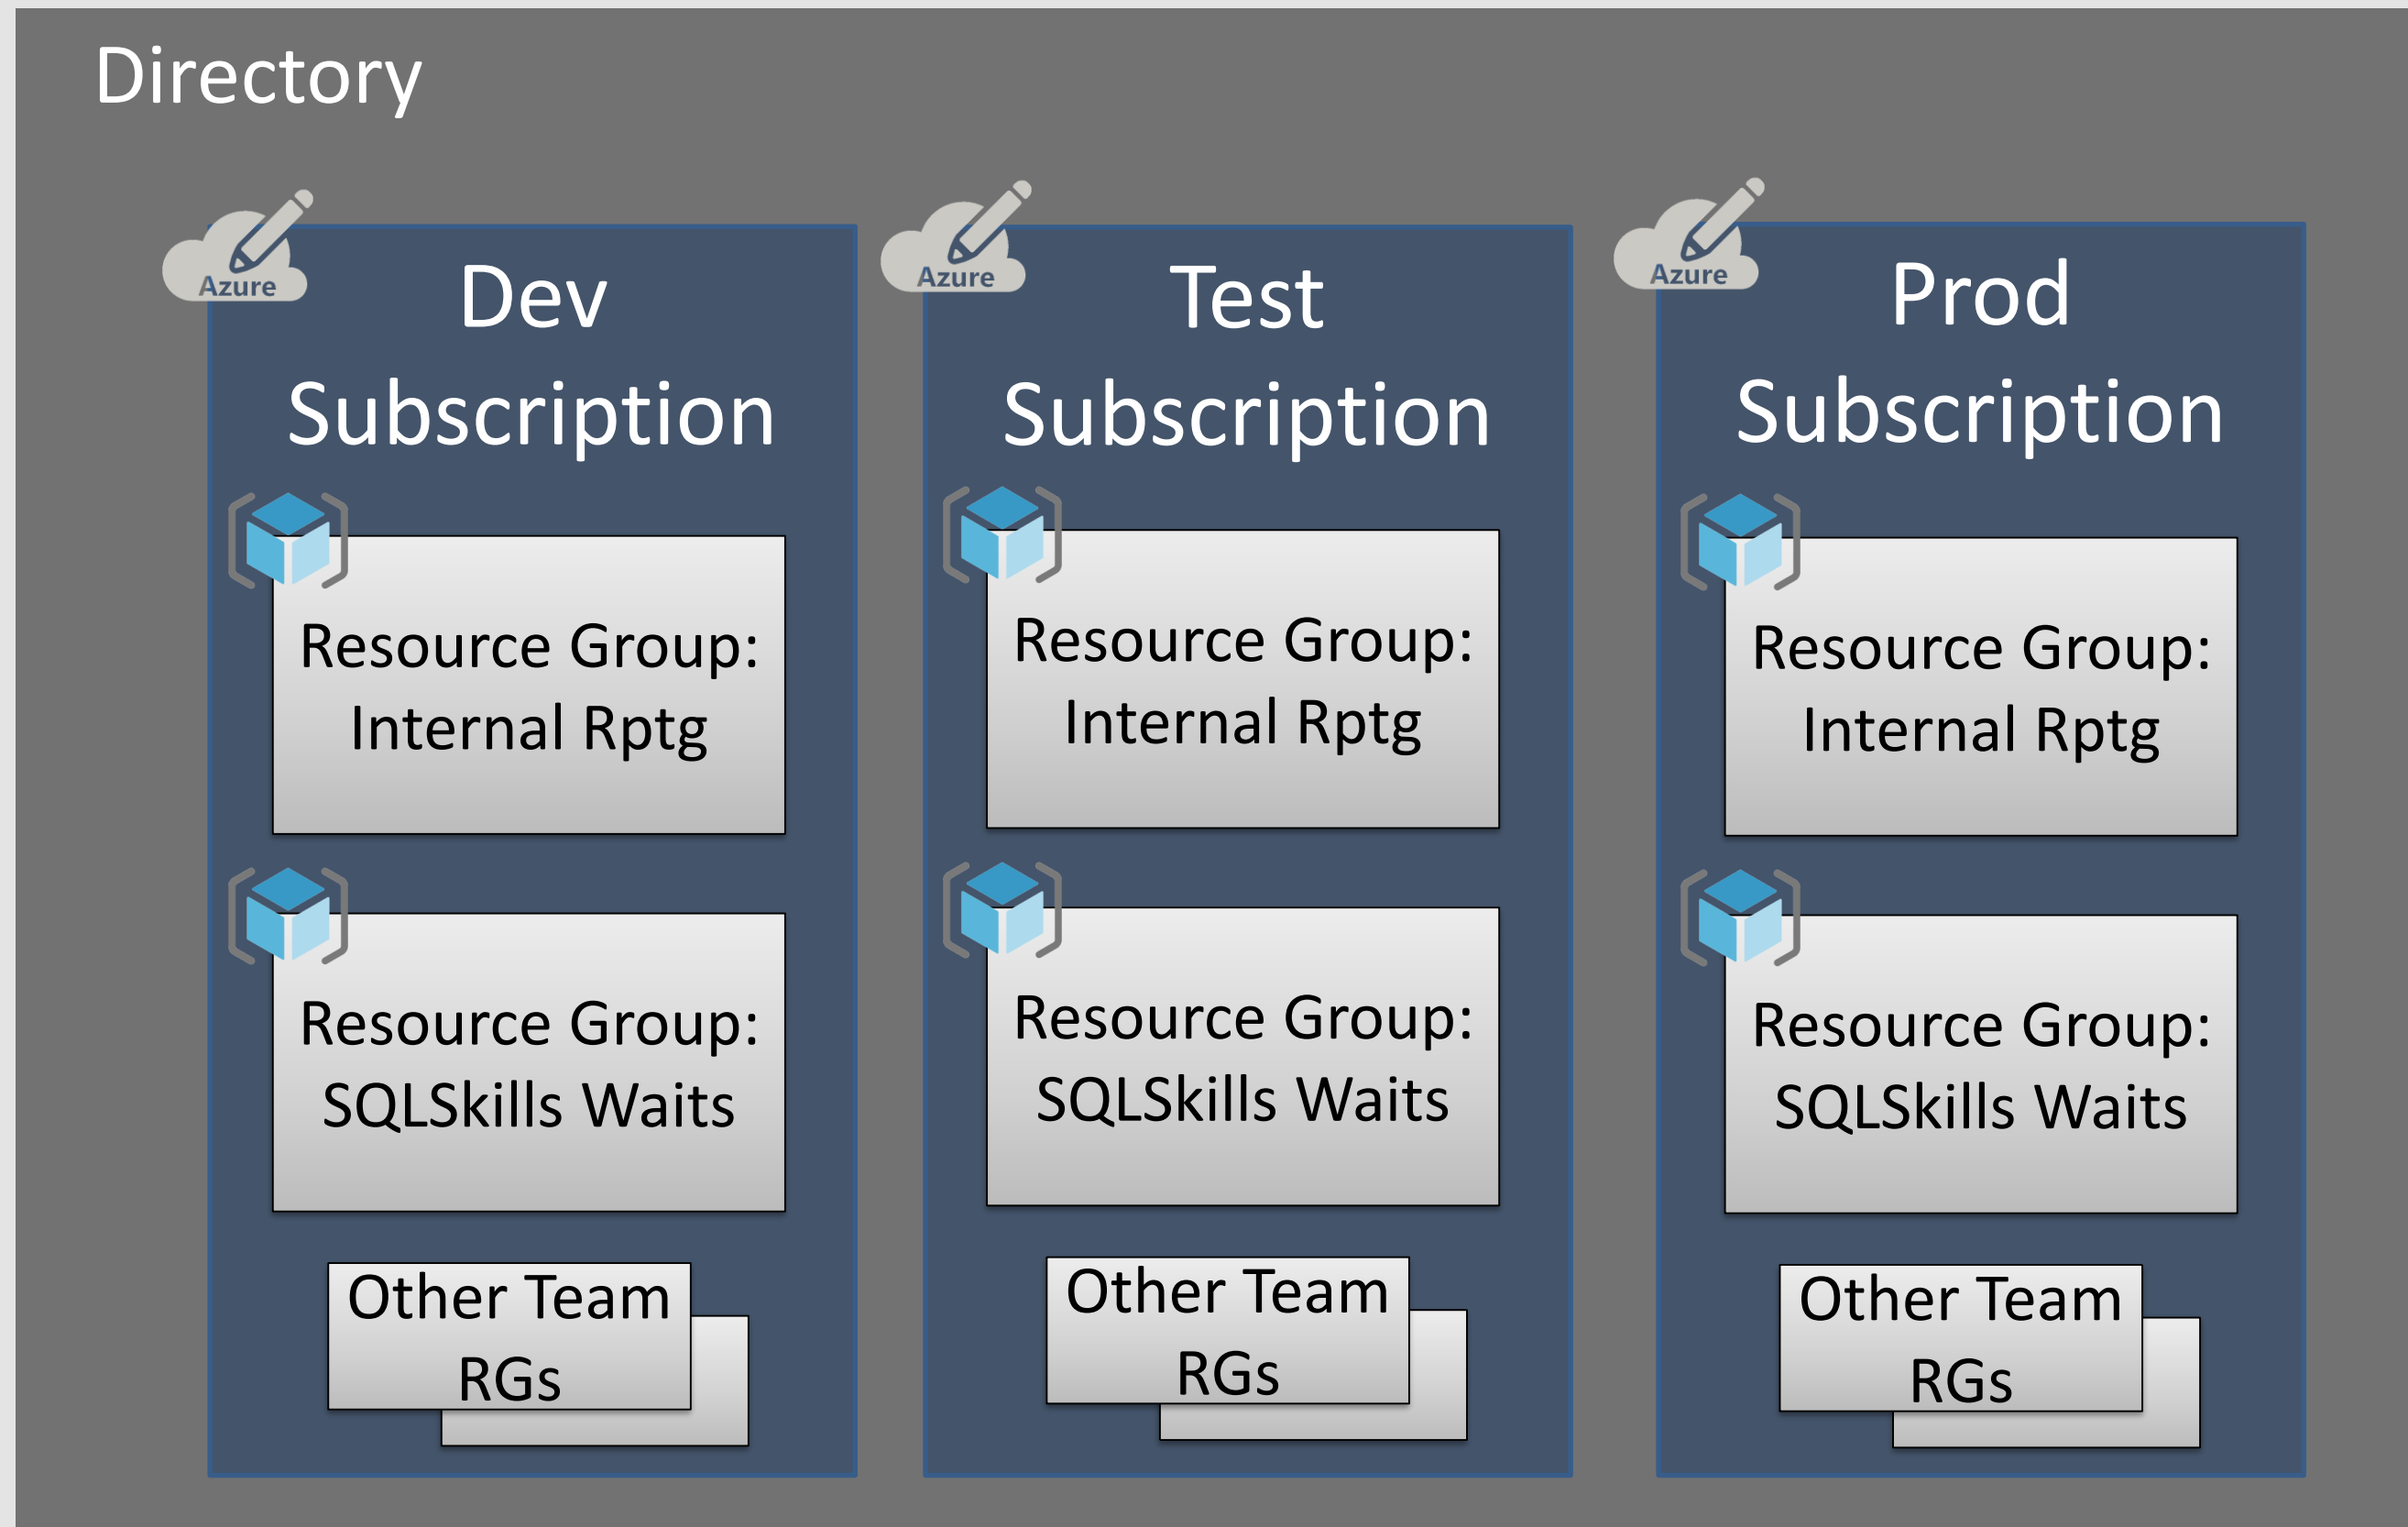
Pros:

- ✓ Clear boundary
- ✓ Offers the most scalability

Cons:

- ✓ More infrastructure to manage
- ✓ A lot of objects intermixed in a subscription - need clear resource group names and/or tags to tell what belongs to which team

Option: Separate By Subscription



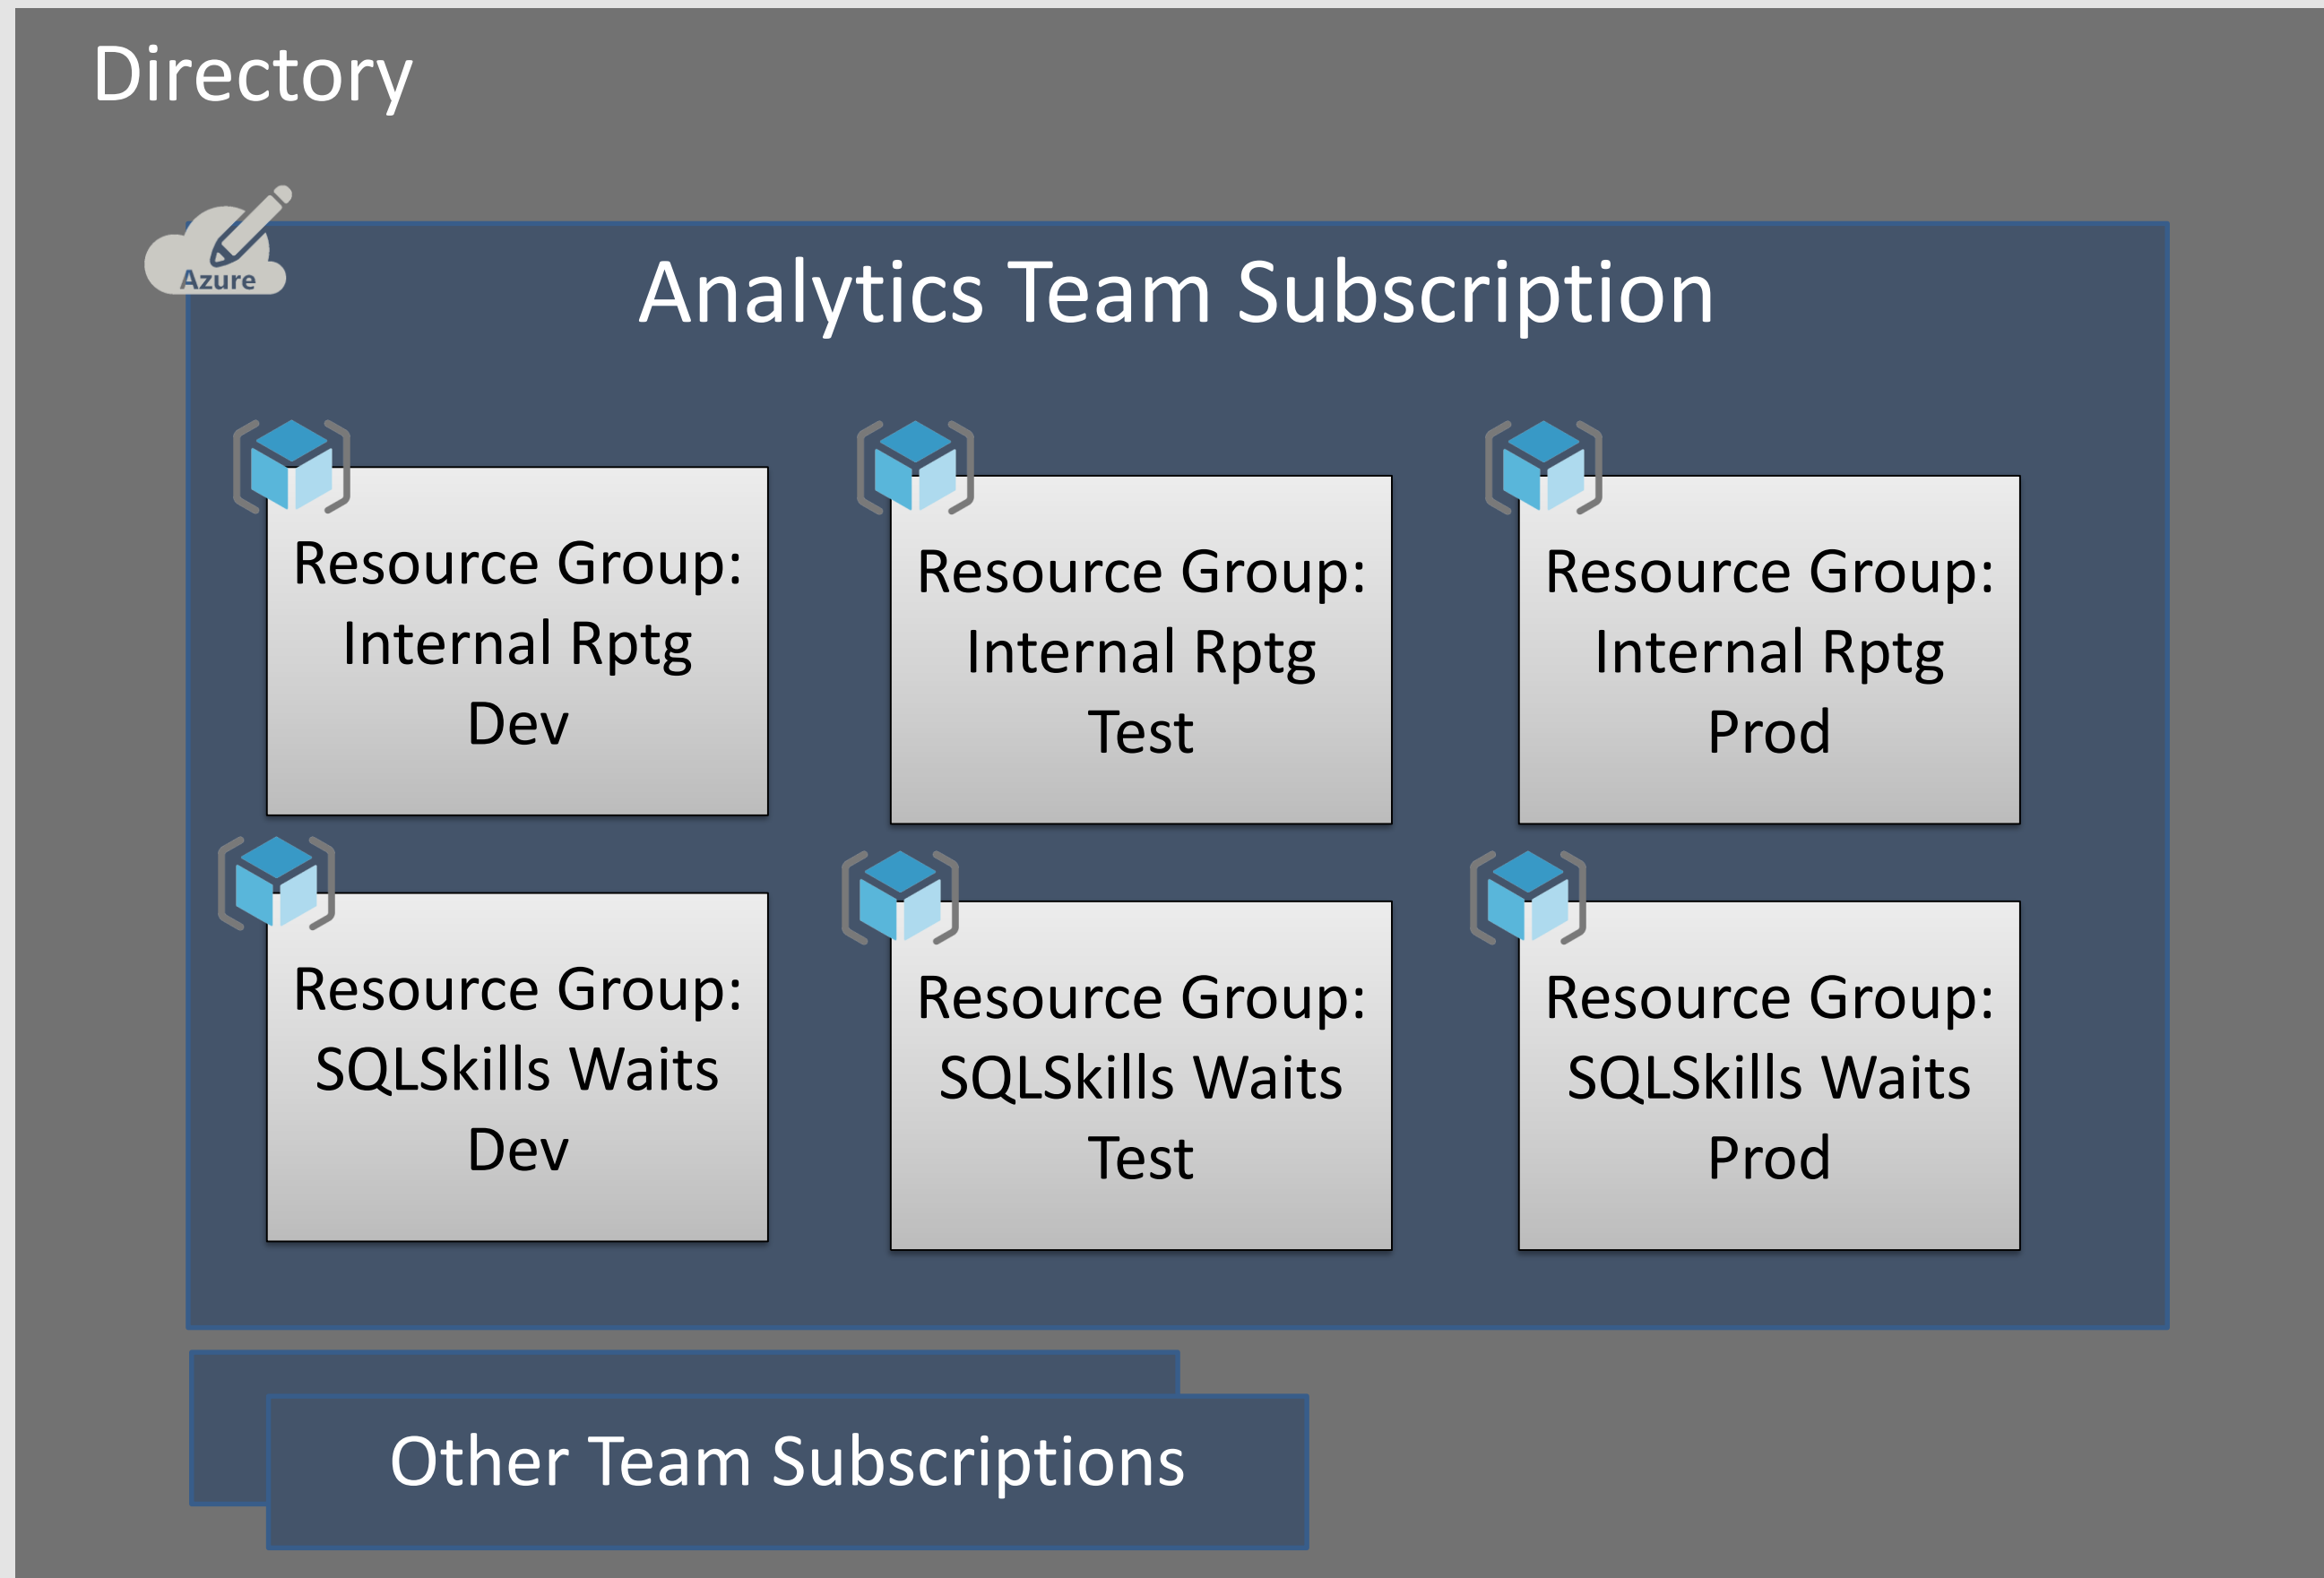
Pros:

- ✓ Clear boundary

Cons:

- ✓ Cannot always provision new resources if you're just an owner at the resource group level
- ✓ Separate virtual networks & VPN set up
- ✓ A lot of objects intermixed in a subscription - need clear resource group names and/or tags to tell what belongs to which team

Option: Separate by Resource Group



Pros:

- ✓ Isolation of subscriptions by the team who owns/manages
- ✓ Set up of virtual network & VPN just once
- ✓ Co-admin privileges easier to delegate at the subscription level for each team
- ✓ Billing segregation
- ✓ Less likely to hit Azure limits

Cons:

- ✓ More risk because Prod is mixed in with Dev & Test (mitigate w/ diff IDs to access Prod resources)

We are currently using this option.

Naming Conventions & Tags

Naming Conventions

Purpose → Type of Service → Environment

- ✓ **Type of service** in the name helps with logging/metrics in monitoring scenarios
- ✓ **Environment as the suffix** makes any concatenations easy within scripts
- ✓ **Prod is enumerated** because we work within one subscription
- ✓ **No dashes** since all services don't allow them
- ✓ **Camel case** if the service allows it; otherwise lower case

Resource Groups

InternalReporting**RG**Dev
InternalReporting**RG**Test
InternalReporting**RG**Prod

Virtual Machines

BISQL**VM**1Dev
BISQL**VM**1Test
BISQL**VM**1Prod

Storage Accounts

BISQLVM1**DataStd****Strg**Dev
BISQLVM1**BckStd****Strg**Dev
BISQLVM1**DiagStd****Strg**Dev

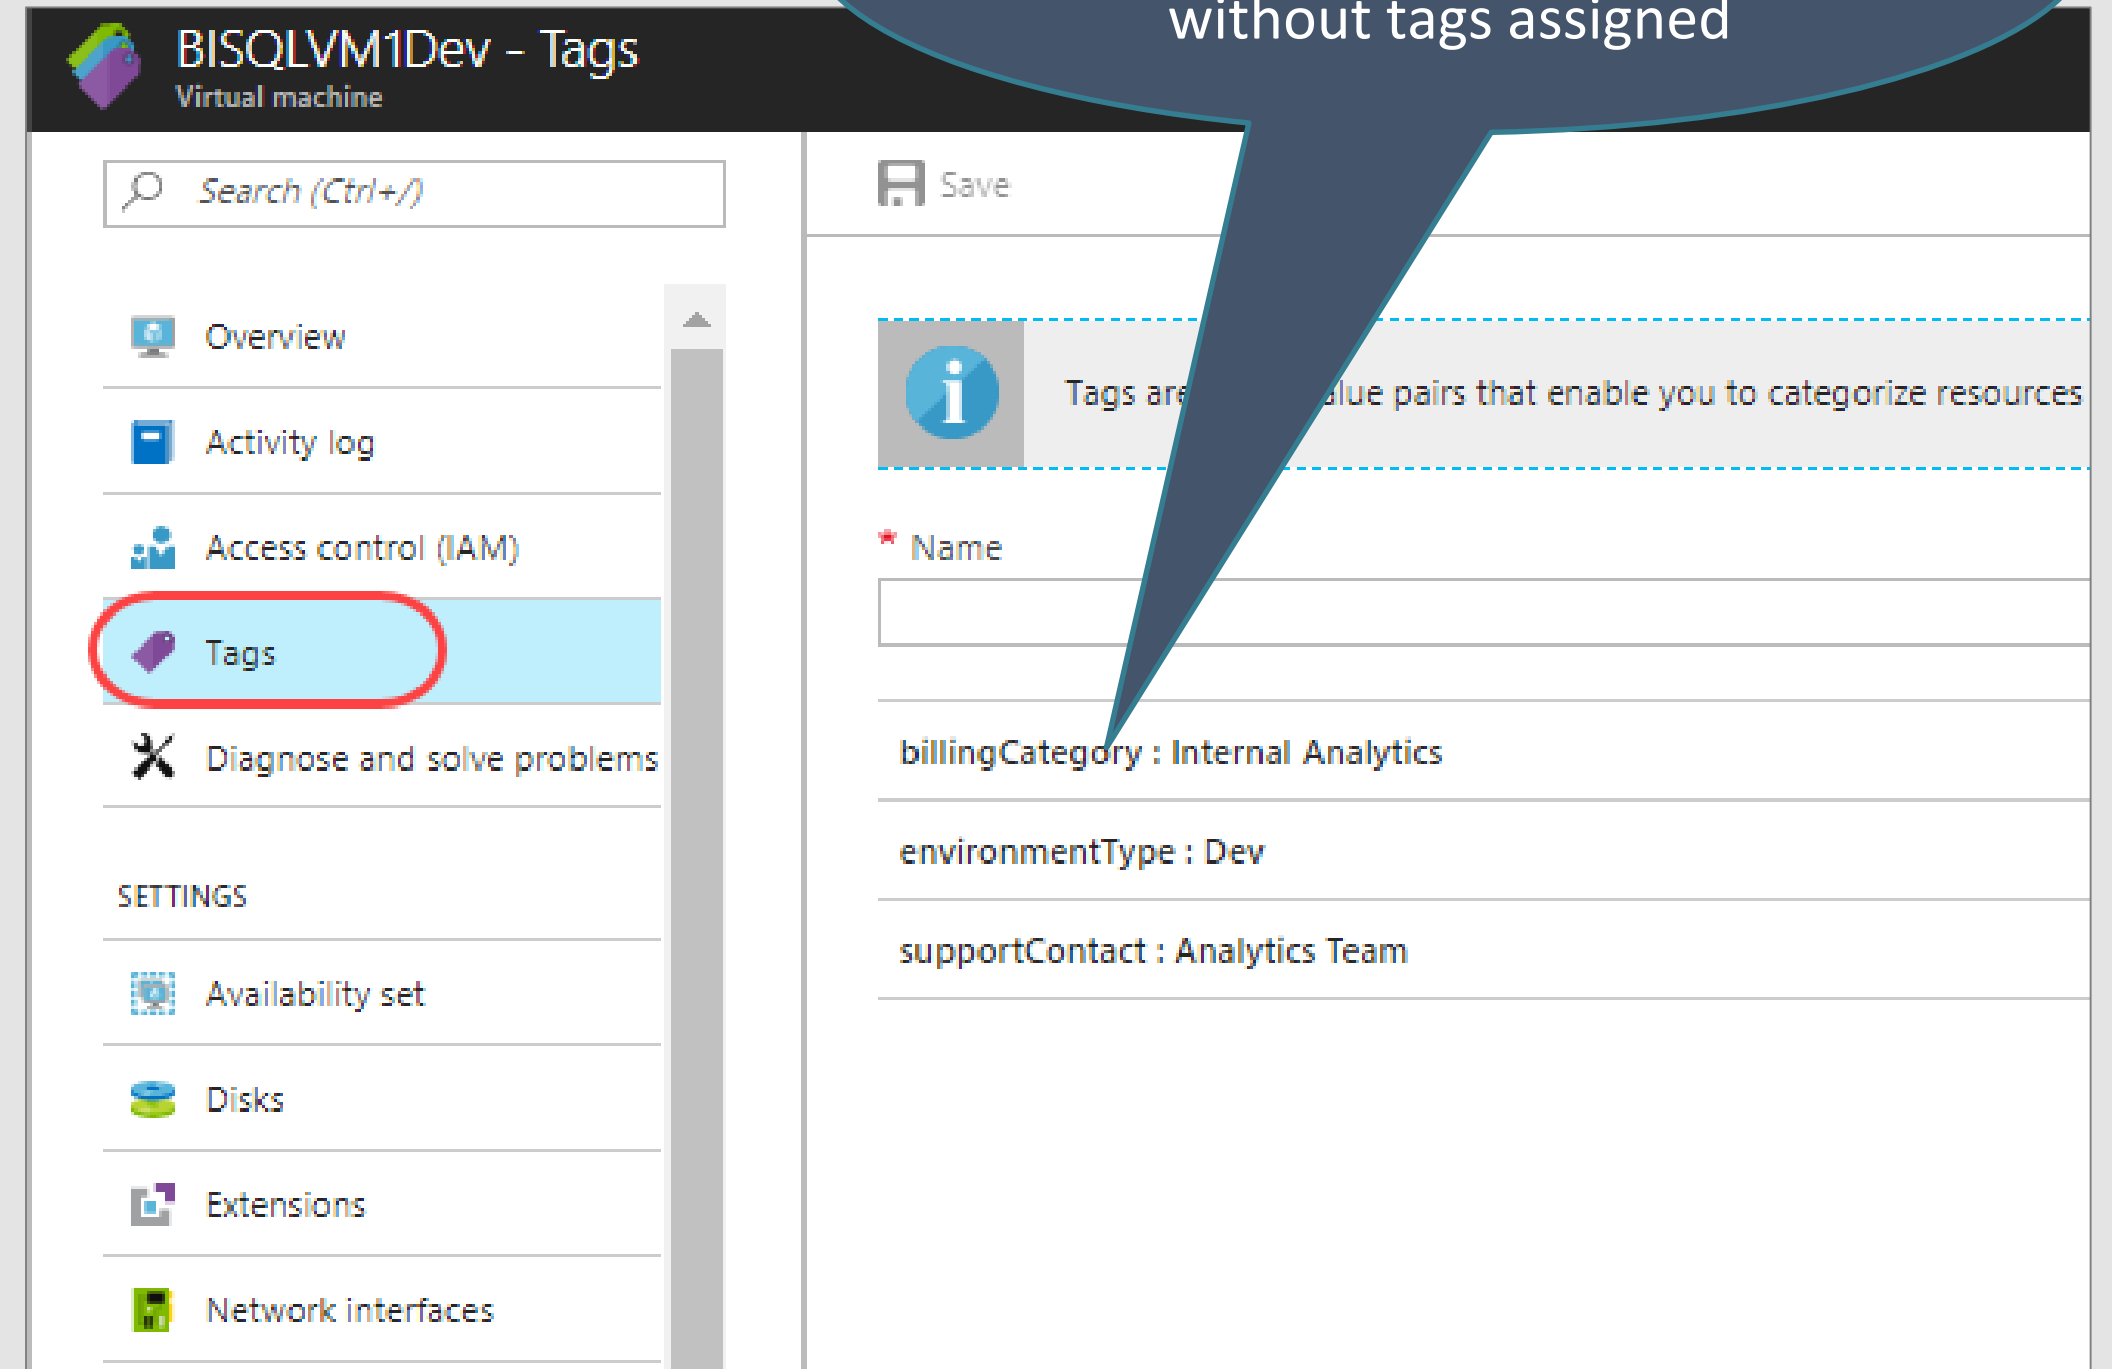
Not our "real" object names but close enough

Standard Tags

Tags are very helpful for:

- ✓ Billing or cost center categories
- ✓ Environment names
- ✓ Project or system
- ✓ Purpose or application
- ✓ Team, department, business unit
- ✓ Who owns or supports a resource
- ✓ Release or version #s (ex: testing infrastructure)
- ✓ Archival date (ex: if needed only temporarily)
- ✓ Patching or maintenance window or SLA
- ✓ Which customer it applies to (ex: if an ISV)
- ✓ etc...

Try to assign tags right away so you don't have partial billing without tags assigned



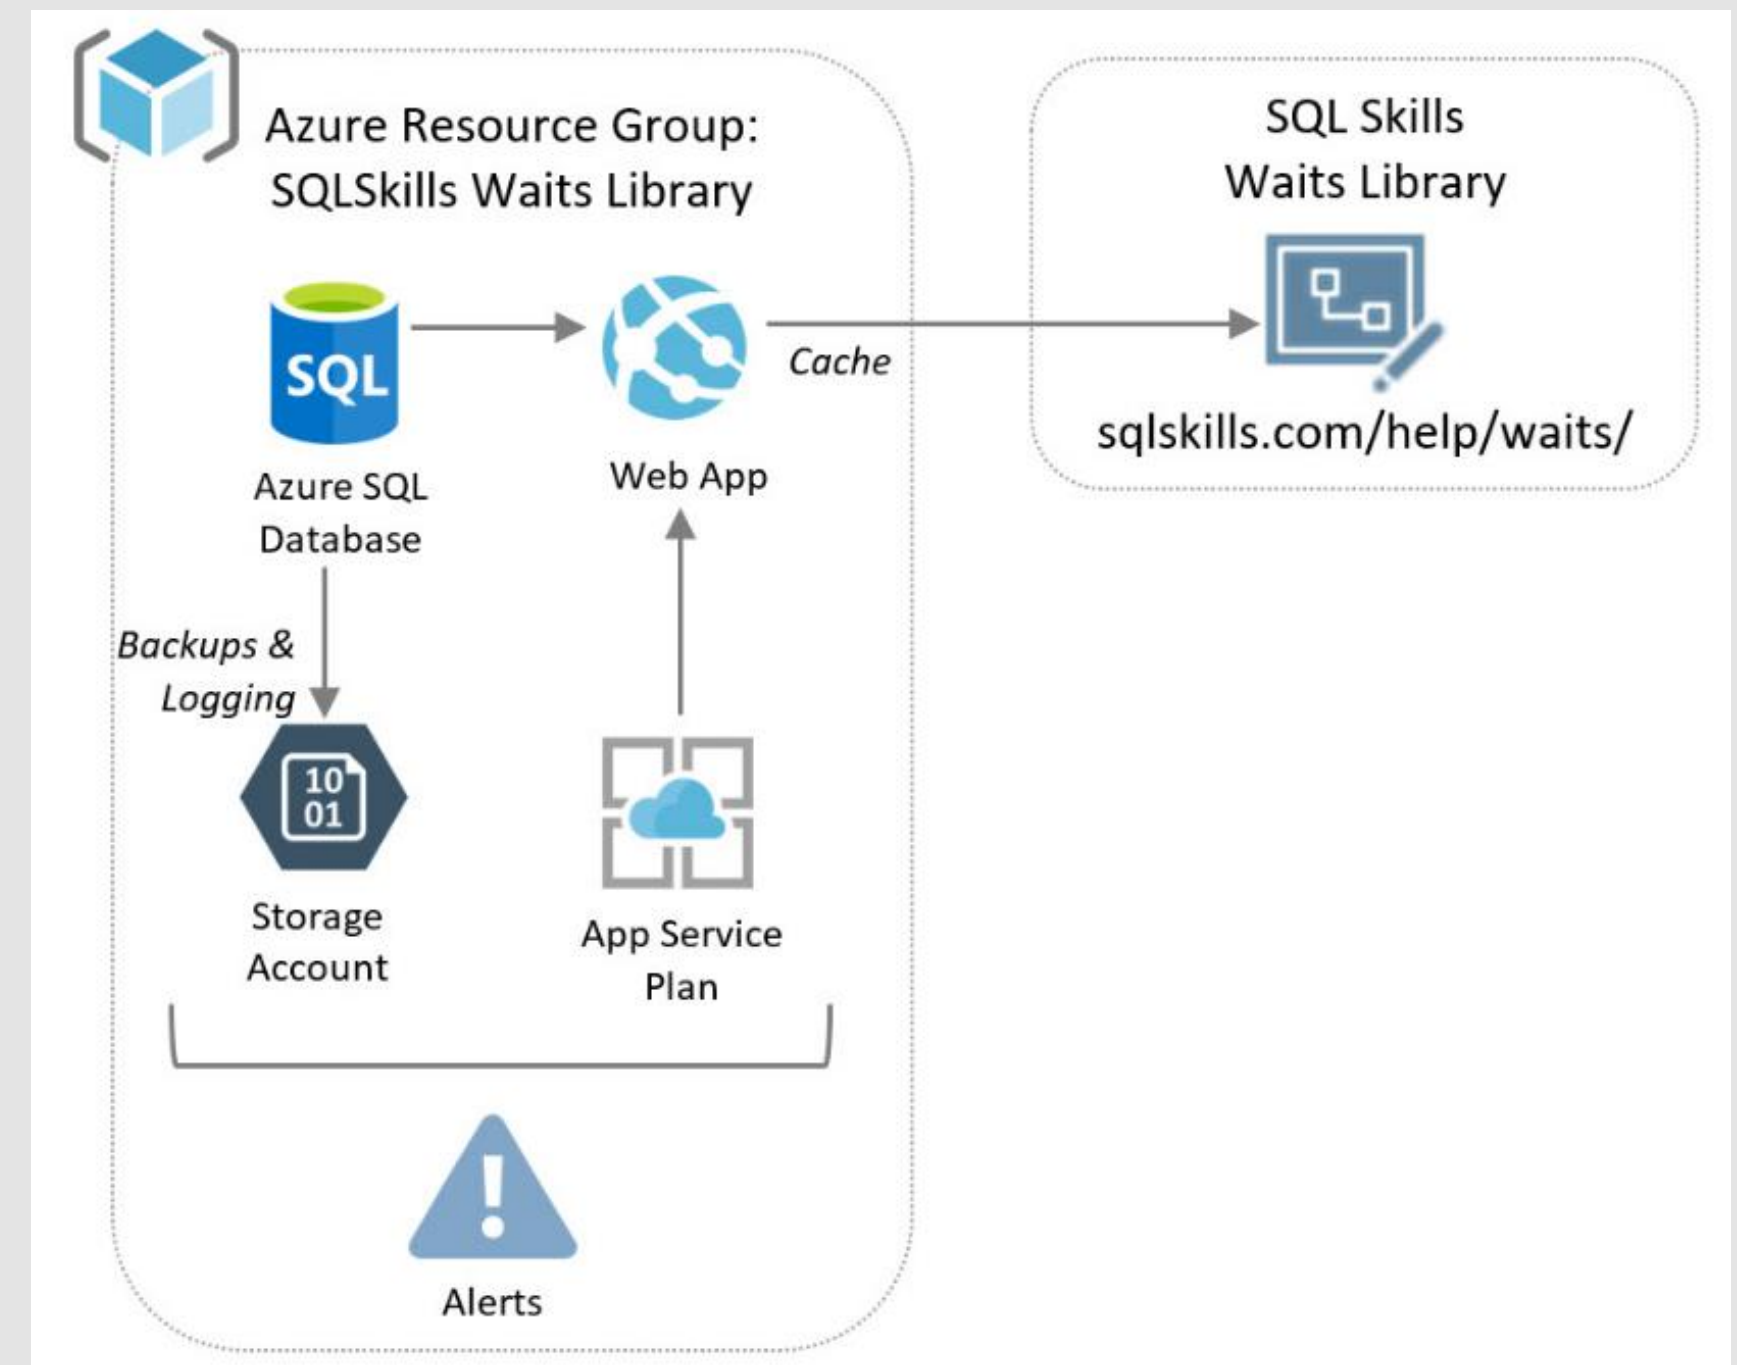
Meter Name	Consumed Service	Resource Group	Tags
"Standard IO - Block Blob Read Operation Units (in 10,000s)"	"Microsoft.Storage"	"internalreportingrgdev"	"{"billingCategory":"Internal Analytics","environmentType":"Dev","supportContact":"Analytics Team"}"
"Standard IO - Table Write Operation Units (in 10,000s)"	"Microsoft.Storage"	"internalreportingrgdev"	"{"billingCategory":"Internal Analytics","environmentType":"Dev","supportContact":"Analytics Team"}"
"Standard IO - Page Blob/Disk (GB)"	"Microsoft.Storage"	"internalreportingrgdev"	"{"billingCategory":"Internal Analytics","environmentType":"Dev","supportContact":"Analytics Team"}"
"Standard IO - Table (GB)"	"Microsoft.Storage"	"internalreportingrgdev"	"{"billingCategory":"Internal Analytics","environmentType":"Dev","supportContact":"Analytics Team"}"

Configuration Decisions: Azure SQL Database

Use of Azure SQLDB

We're using Azure SQL Database for a new solution just created recently.

NAME	STATUS	REPLICATION ROLE	SERVER	PRICING TIER
sqlwaitssqldatabase	Online	None	sqlwaitsserverdev	Standard: S0
sqlwaitssqldatabase	Online	None	sqlwaitsserverprod	Standard: S2



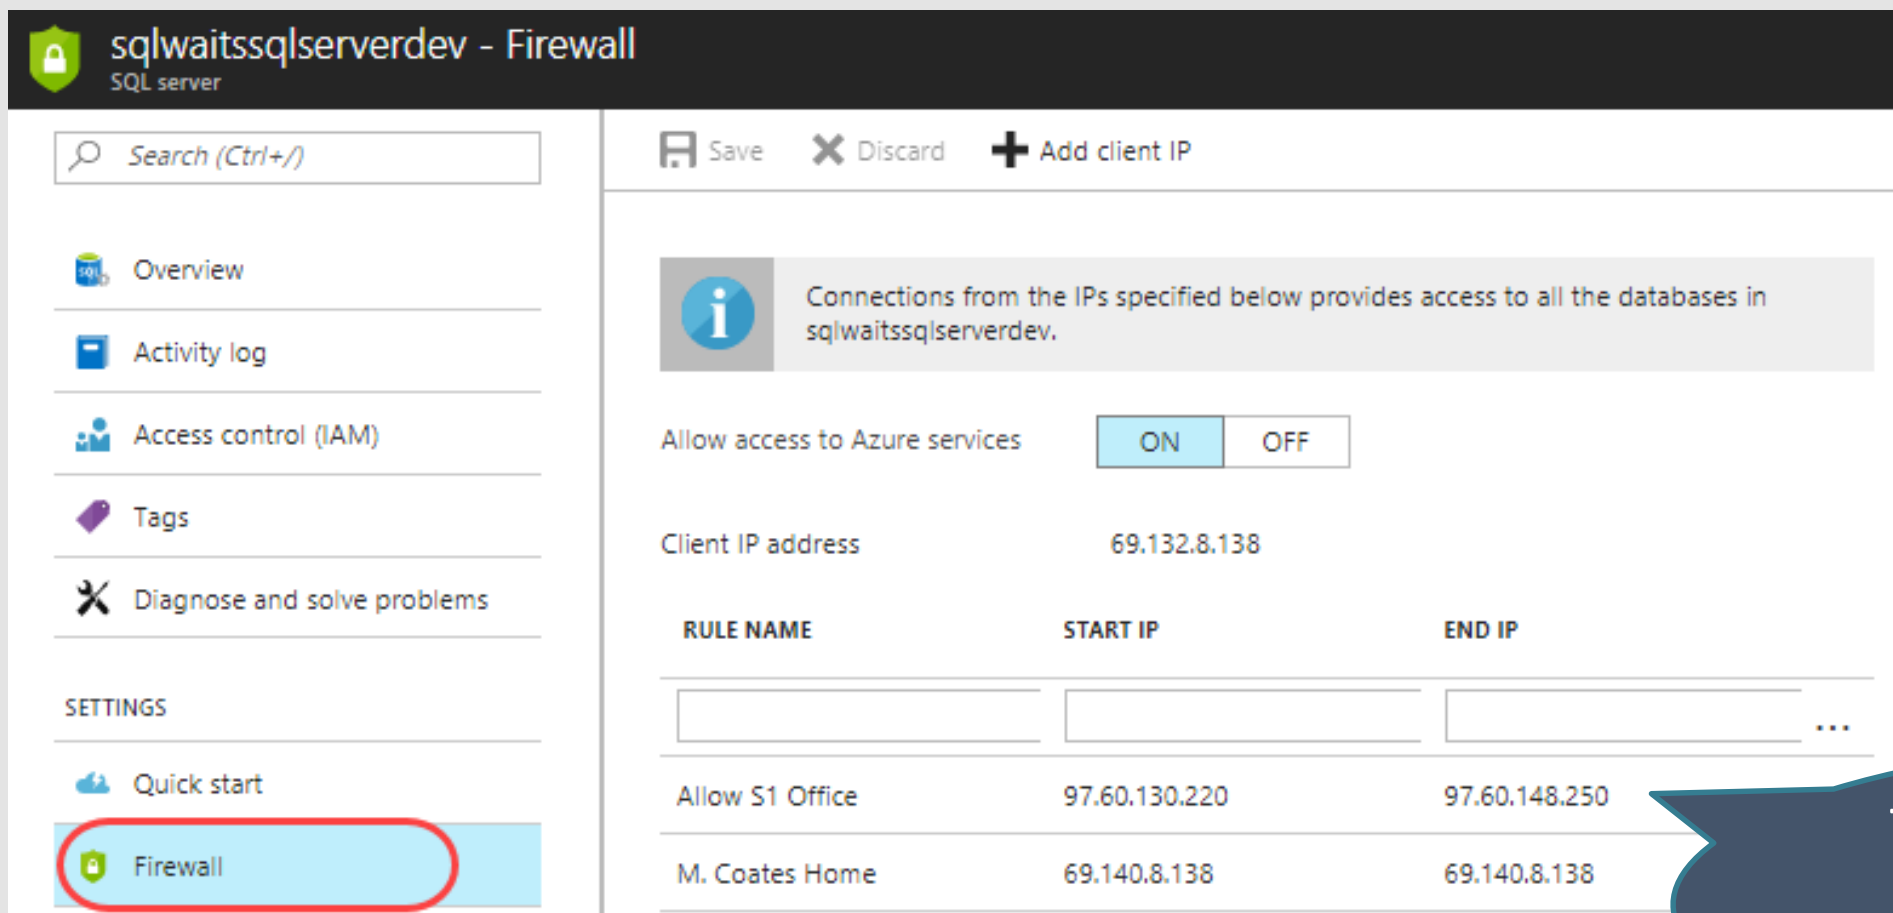
Our solution contains aggregate, anonymous data.

There is no Virtual Network/VPN functionality for SQLDB or SQLDW yet (coming in "Managed Instances" of Azure SQLDB).

Firewall for Azure SQLDB

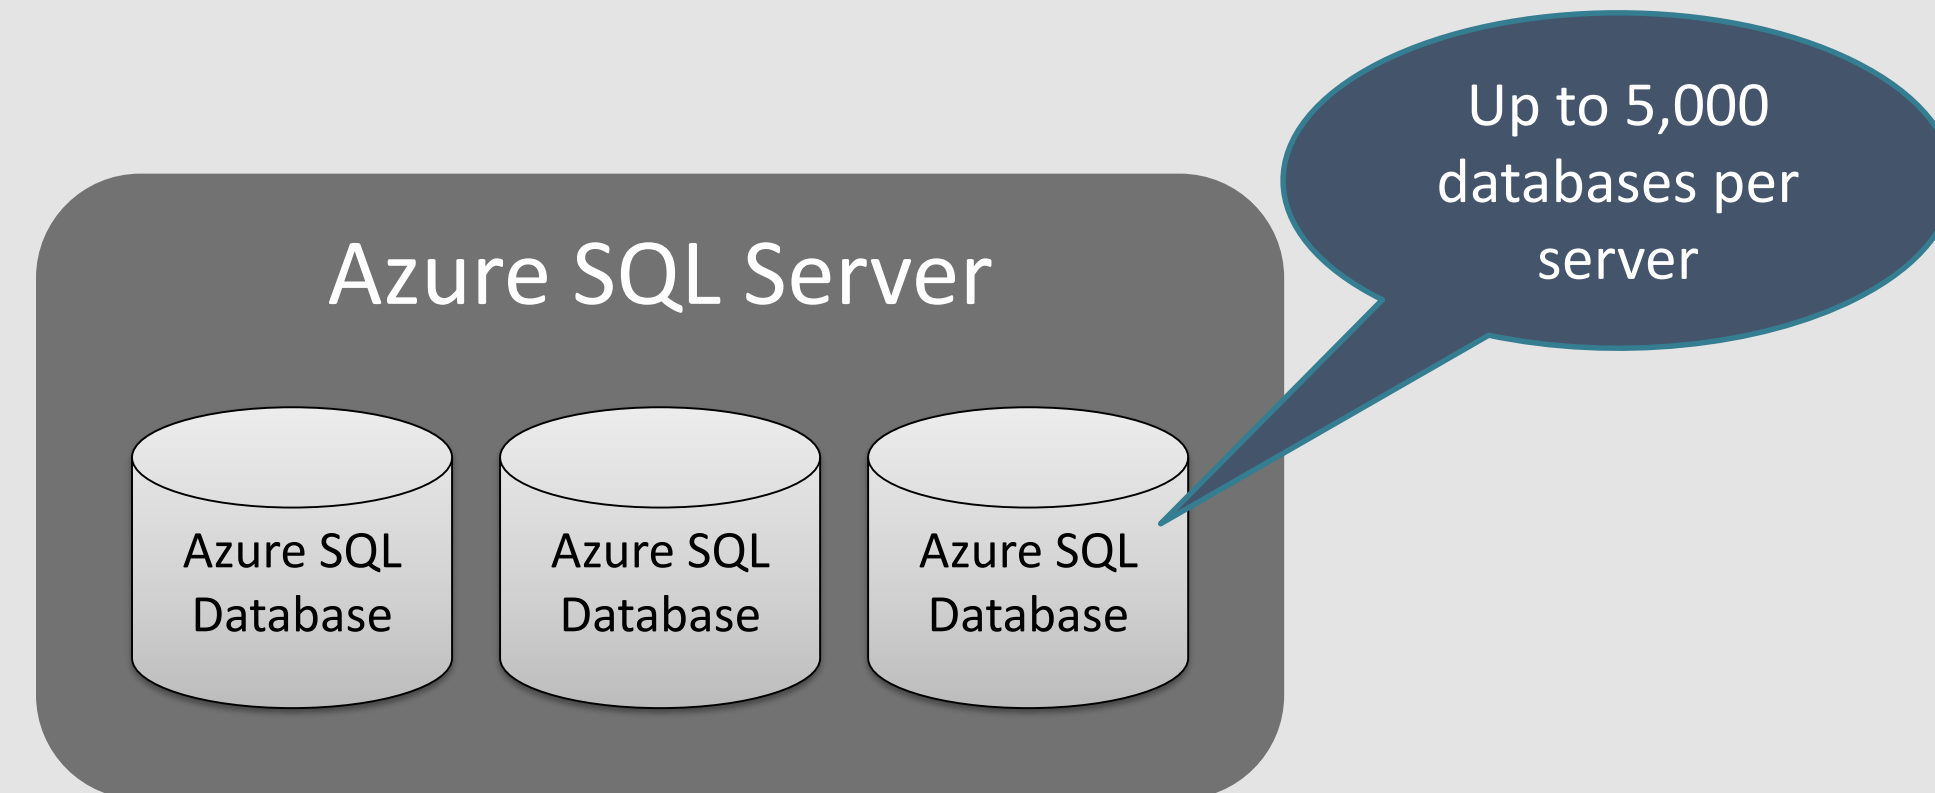
A firewall can be set at the server level, or at the database level.

Server level firewall rules can be set in the portal:



The screenshot shows the Azure portal interface for the Firewall settings of an SQL server named 'sqlwaitssqlserverdev'. The left sidebar contains navigation options: Overview, Activity log, Access control (IAM), Tags, Diagnose and solve problems, and SETTINGS. Under SETTINGS, 'Quick start' and 'Firewall' are listed, with 'Firewall' highlighted in blue and circled in red. The main content area shows the Firewall configuration for the server. It includes a search bar, 'Save', 'Discard', and 'Add client IP' buttons. An information icon indicates that connections from the specified IPs provide access to all databases in the server. The 'Allow access to Azure services' toggle is set to 'ON'. The 'Client IP address' is set to '69.132.8.138'. Below this is a table of firewall rules:

RULE NAME	START IP	END IP
		...
Allow S1 Office	97.60.130.220	97.60.148.250
M. Coates Home	69.140.8.138	69.140.8.138



Database level firewall rules must be set with T-SQL. Although that makes them less visible, it is a better practice.

Try to assign a name to each entry. The generic default name isn't helpful.

Configuration Decisions: SQL Server in Azure Virtual Machine

Starting Point: VM Image

Virtual machines Directory

Virtual machines and Virtual machines (classic) can now be managed together in the combined list below.

Subscriptions: 1 of 2 selected – Don't see a subscription? Switch directories

Filter by name...

2 items

NAME	
BISQLVM1Dev	...
BISQLVM2Dev	...

Compute

Filter

sql server 2016 SP1

Results

NAME	PUBLISHER
SQL Server 2016 SP1 Standard on Windows Server 2016	Microsoft
Free License: SQL Server 2016 SP1 Developer on Windows Server 2016	Microsoft
SQL Server 2016 SP1 Enterprise on Windows Server 2016	Microsoft
Free License: SQL Server 2016 SP1 Express on Windows Server 2016	Microsoft
SQL Server 2016 SP1 Web on Windows Server 2016	Microsoft
{BYOL} SQL Server 2016 SP1 Enterprise on Windows Server 2016	Microsoft
{BYOL} SQL Server 2016 SP1 Standard on Windows Server 2016	Microsoft

There's not currently a VM image which is Fast Track certified

SQL Server settings

SQL connectivity: Private (within Virtual Network)

Port: 1433

SQL Authentication: Disable | Enable

Storage configuration: **Data warehousing**

DS12_V2 Standard

- 4 Cores
- 28 GB ← RAM
- 8 Data disks
- 12800 Max IOPS
- 56 GB Local SSD
- Load balancing
- Premium disk support

485.09 USD/MONTH (ESTIMATED)

Storage configuration

Select your desired performance, storage size, and workload to optimize the storage on your virtual machine.

Storage size (TB): 1

1 data disks will be added to the virtual machine. This value was computed based on the value of IOPS, throughput, and storage size.

Storage optimization: **Data warehousing**

This sets storage disk stripe size to 256KB and trace flags 610 & 1117

If the VM scale you want isn't available: <http://www.sqlchick.com/entries/2017/3/30/why-some-azure-vm-sizes-are-unavailable-when-resizing-in-the-portal>


Best Practices Guidance

Azure / Virtual Machines / Windows

Filter

- > Overview
- > Get Started
- ▼ How to
 - > Configure High Availability
 - > Manage
 - > Backup and restore
 - Performance best practices**
 - Configure storage
 - Application patterns
- > Reference

Performance best practices for SQL Server in Azure Virtual Machines

2017-4-28 • 11 min to read • Contributors 

Overview

This topic provides best practices for optimizing SQL Server performance in Microsoft Azure Virtual Machine. While running SQL Server in Azure Virtual Machines, we recommend that you continue using the same database performance tuning options that are applicable to SQL Server in on-premises server environment. However, the performance of a relational database in a public cloud depends on many factors such as the size of a virtual machine, and the configuration of the data disks.

Area	Optimizations
VM size	DS3 or higher for SQL Enterprise edition. DS2 or higher for SQL Standard and Web editions.
Storage	Use Premium Storage . Standard storage is only recommended for dev/test. Keep the storage account and SQL Server VM in the same region. Disable Azure geo-redundant storage (geo-replication) on the storage account.
Disks	Use a minimum of 2 P30 disks (1 for log files; 1 for data files and TempDB). Avoid using operating system or temporary disks for database storage or logging. Enable read caching on the disk(s) hosting the data files and TempDB. Do not enable caching on disk(s) hosting the log file. Important: Stop the SQL Server service when changing the cache settings for an Azure VM disk. Stripe multiple Azure data disks to get increased IO throughput. Format with documented allocation sizes.
I/O	Enable database page compression. Enable instant file initialization for data files. Limit or disable autogrow on the database. Disable autoshrink on the database. Move all databases to data disks, including system databases. Move SQL Server error log and trace file directories to data disks. Setup default backup and database file locations. Enable locked pages. Apply SQL Server performance fixes.
Feature specific	Back up directly to blob storage.

Most of the guidance in this Performance Best Practices article is implemented in the pre-built image.

<https://docs.microsoft.com/en-us/azure/virtual-machines/windows/sql/virtual-machines-windows-sql-performance>

Changes Made to Pre-Built VM Image

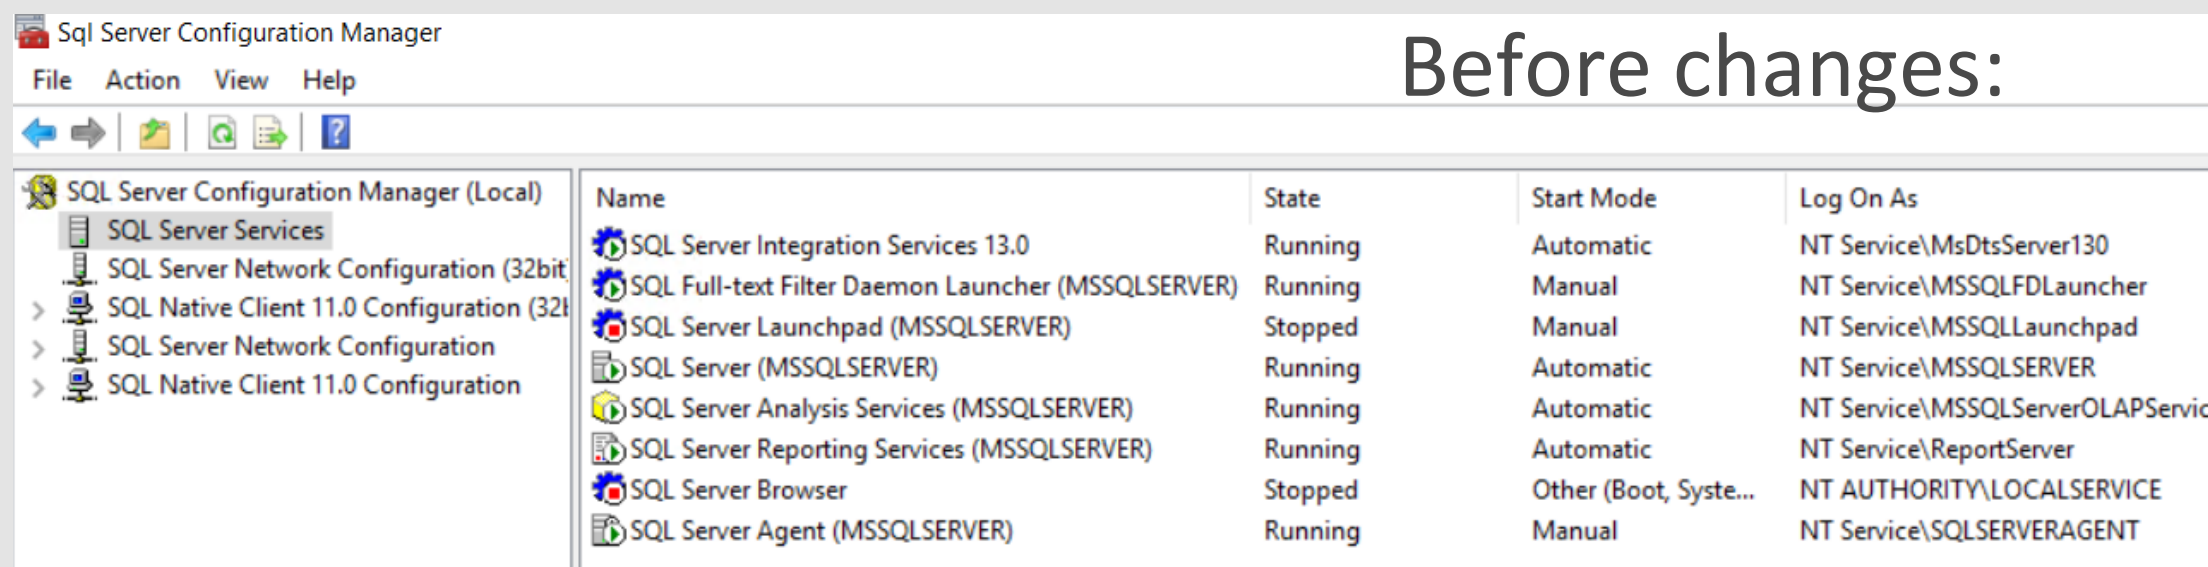
We started with a pre-built image, and made customizations after it was domain-joined, such as:

- Domain service accounts for each service
- Disable unused services
- Disk structure + permissions
 - G: Data
 - L: Logs
 - T: TempDB
- Enable disk encryption
- Policies (volume maintenance, etc)

SQL Server changes for each new VM:

- Enable the remote DAC
- Default database file locations
- Fill factor settings at database level
- Optimize for ad hoc workloads
- Max degree of parallelism
- Cost threshold for parallelism
- Startup parameters & trace flags
- Expand # of TempDB files & relocate
- Time of SSIS maintenance job
- Default for backup checksum
- Limited account to act as DBOwner etc...

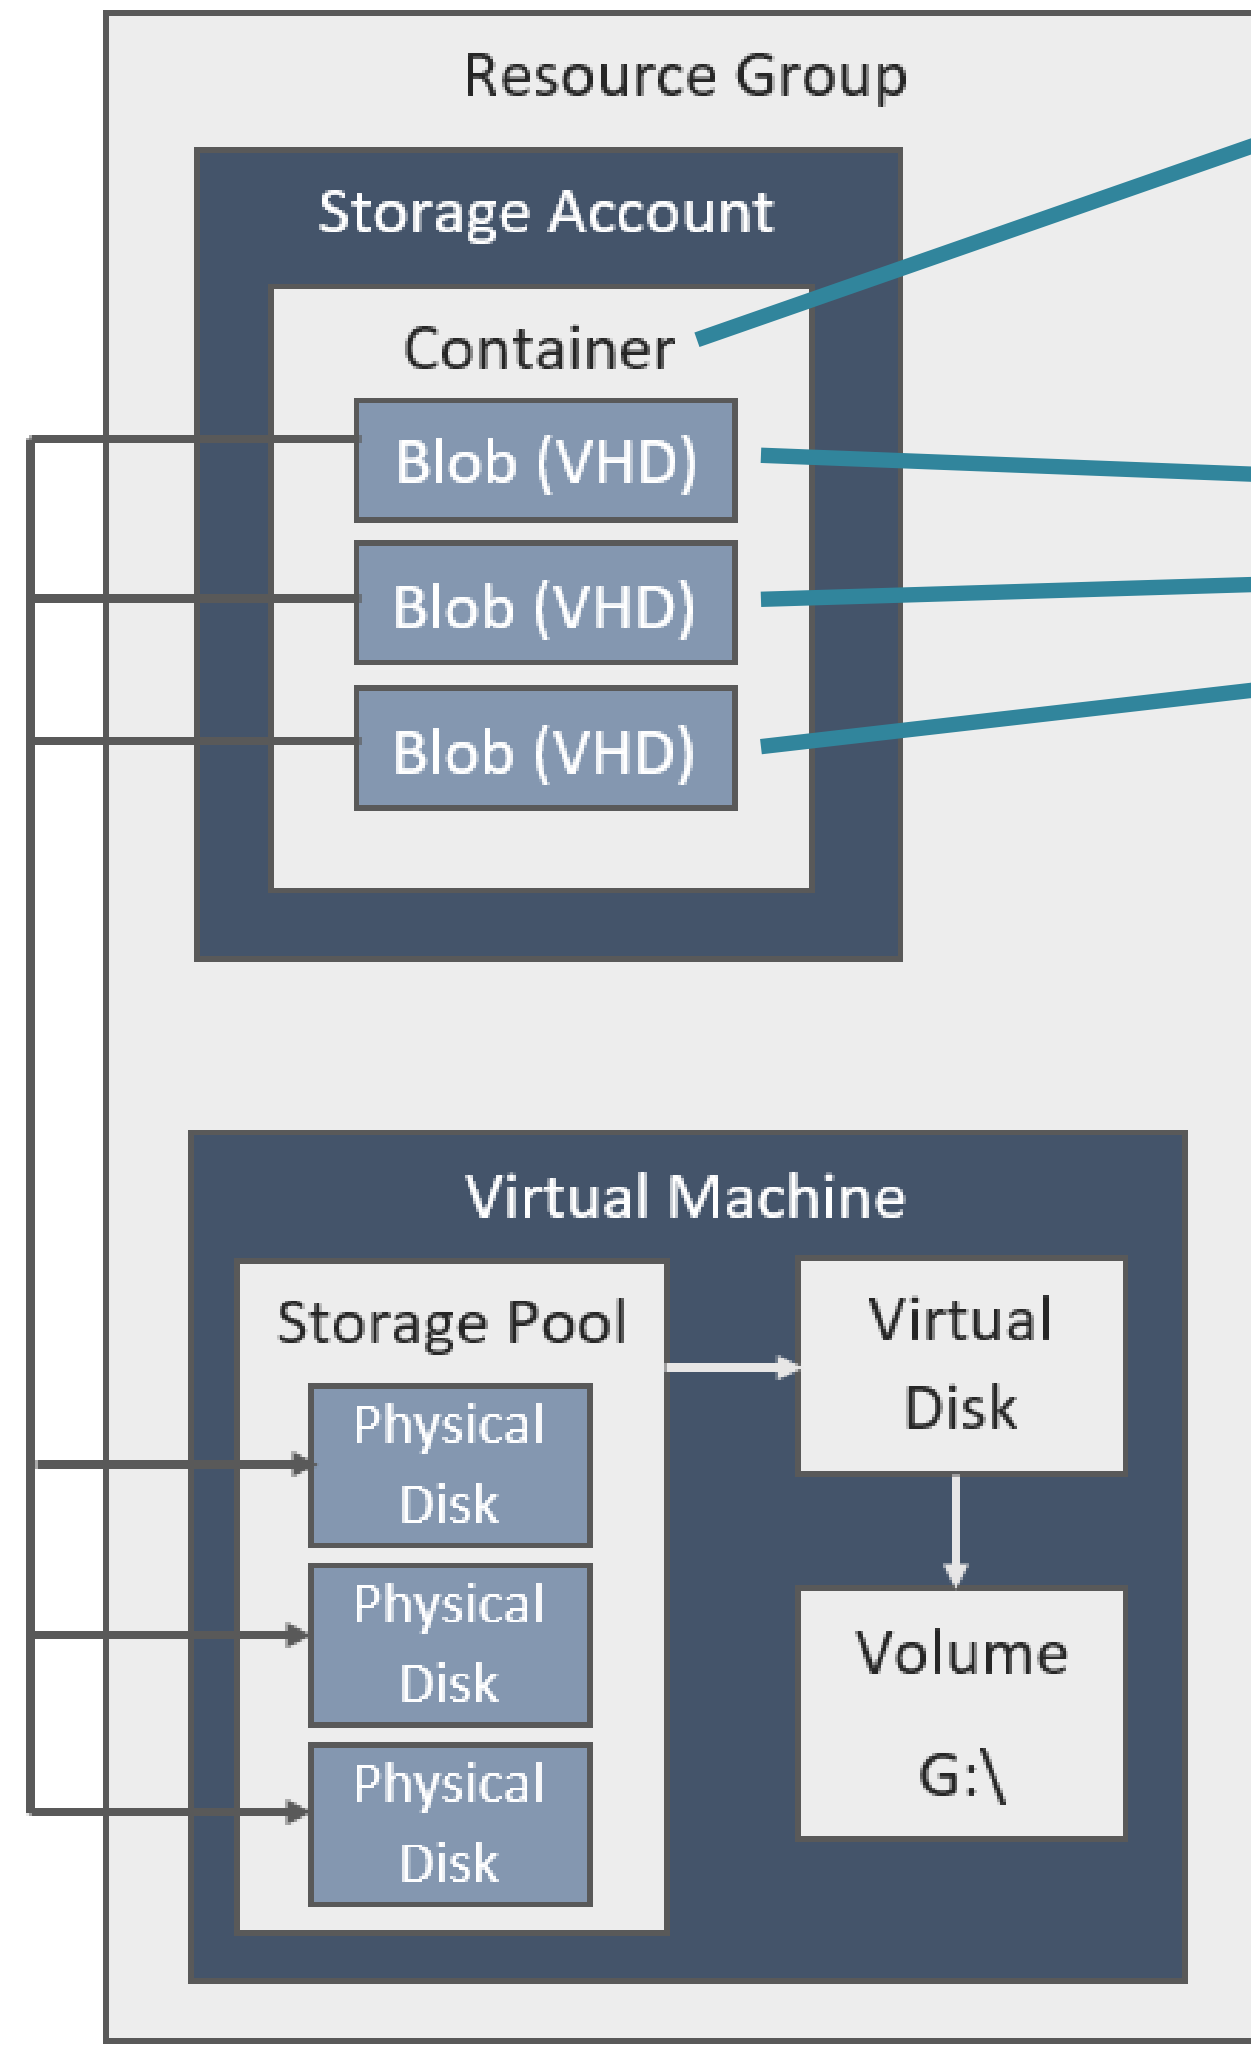
Before changes:



Name	State	Start Mode	Log On As
SQL Server Integration Services 13.0	Running	Automatic	NT Service\MsDtsServer130
SQL Full-text Filter Daemon Launcher (MSSQLSERVER)	Running	Manual	NT Service\MSSQLFDLauncher
SQL Server Launchpad (MSSQLSERVER)	Stopped	Manual	NT Service\MSSQLLaunchpad
SQL Server (MSSQLSERVER)	Running	Automatic	NT Service\MSSQLSERVER
SQL Server Analysis Services (MSSQLSERVER)	Running	Automatic	NT Service\MSSQLServerOLAPService
SQL Server Reporting Services (MSSQLSERVER)	Running	Automatic	NT Service\ReportServer
SQL Server Browser	Stopped	Other (Boot, System, or Start on Demand)	NT AUTHORITY\LOCALSERVICE
SQL Server Agent (MSSQLSERVER)	Running	Manual	NT Service\SQLSERVERAGENT

Storage

The VHDs (virtual hard disks) behind a VM are in Azure Storage.



The screenshot shows the Azure portal interface for a container named **vhds**. The container contains a list of blobs, including **bisqlvm1datadisk1dev.vhd**, **bisqlvm1datadisk2dev.vhd**, **bisqlvm1datadisk3dev.vhd**, and others. The **bisqlvm1datadisk1dev.vhd** blob is selected, and its properties are displayed on the right. A warning message states: "This blob has an active lease and cannot be deleted. Click here to learn more about leases." The properties for this blob are:

- NAME:** bisqlvm1datadisk1dev.vhd
- URL:** https://bisqlvm1datastdstrgdev.blob.co...
- LAST MODIFIED:** 7/11/2017, 7:29:48 AM
- TYPE:** Page blob
- SIZE:** 550 GiB
- ETAG:** 0x8D4C8502348D378
- CONTENT-MD5:** -
- LEASE STATUS:** Locked
- LEASE STATE:** Leased
- LEASE DURATION:** Infinite

The lease associates it with the VM

Storage

Backups (“bck”), Data, and Diagnostic logging (“diag”) are separated

All data for a single VM is in the same storage account – it is a unit of recovery in case of failure

Associated with Azure VMs (IaaS)

Associated with Azure SQLDB (PaaS)

NAME	RESOURCE GROUP	LOCATION	SKU	BILLINGCATEGORY (TAG)	ENVIRONMENTTYPE (TAG)	SUPPORTCONTACT (TAG)
INTERNALREPORTINGRGDEV						
bisqlvm1 bck stdstrgdev	InternalReportingRGDev	East US	Standard_RAGRS	Internal Analytics	Dev	BI & Analytics Team
bisqlvm1 data stdstrgdev	InternalReportingRGDev	East US	Standard_LRS	Internal Analytics	Dev	BI & Analytics Team
bisqlvm1 diag stdstrgdev	InternalReportingRGDev	East US	Standard_LRS	Internal Analytics	Dev	BI & Analytics Team
SQLSKILLSWAITSLIBRARYRGDEV						
sqlwaits bck stdstrgdev	SQLSkillsWaitsLibraryRGDev	East US	Standard_RAGRS	SQLSkills Waits Library	Dev	BI & Analytics Team
sqlwaits diag stdstrgdev	SQLSkillsWaitsLibraryRGDev	East US	Standard_LRS	SQLSkills Waits Lib	Dev	BI & Analytics Team

Geo-redundancy (GRS) for backup files

Managed/Unmanaged

We chose to use unmanaged storage as shown above (managed was released to preview during our implementation). Managed VM storage doesn't display in the portal.

Premium/Standard

We are using standard storage for VMs in Dev and Test, and Premium in Production.

VM Disks

(1/8)

Default disks provisioned by the pre-built VM image.

We omitted the F: drive (after moving TempDB to its proper drive).

The default data disk size is 1TB

The screenshot shows the Windows Disk Management console. At the top, there is a table listing the volumes. Below this, three individual disk details are shown, each with a blue header bar and a white content area.

Volume	Layout	Type	File System	Status	Capacity	Free Space	% Free
(C:)	Simple	Basic	NTFS	Healthy (System, Boot, Active, Crash Dump, Primary Partition)	127.00 GB	99.11 GB	78 %
SQLVMDATA1 (F:)	Simple	Basic	NTFS	Healthy (Primary Partition)	1021.87 GB	1020.12 GB	100 %
Temporary Storage (D:)	Simple	Basic	NTFS	Healthy (Page File, Primary Partition)	14.00 GB	12.82 GB	92 %

Disk	Layout	Type	File System	Status	Capacity	Free Space	% Free
Disk 0	Basic	Basic	NTFS	Healthy (System, Boot, Active, Crash Dump, Primary Partition)	127.00 GB	99.11 GB	78 %
Disk 1	Basic	Basic	NTFS	Healthy (Page File, Primary Partition)	14.00 GB	12.82 GB	92 %
Disk 3	Basic	Basic	NTFS	Healthy (Primary Partition)	1021.87 GB	1020.12 GB	100 %

To release the lease on the storage blob:
1. Delete the volume, and
2. Detach disk from the VM

Don't forget to delete the file in Azure Storage as well.

VM Disks

(2/8)

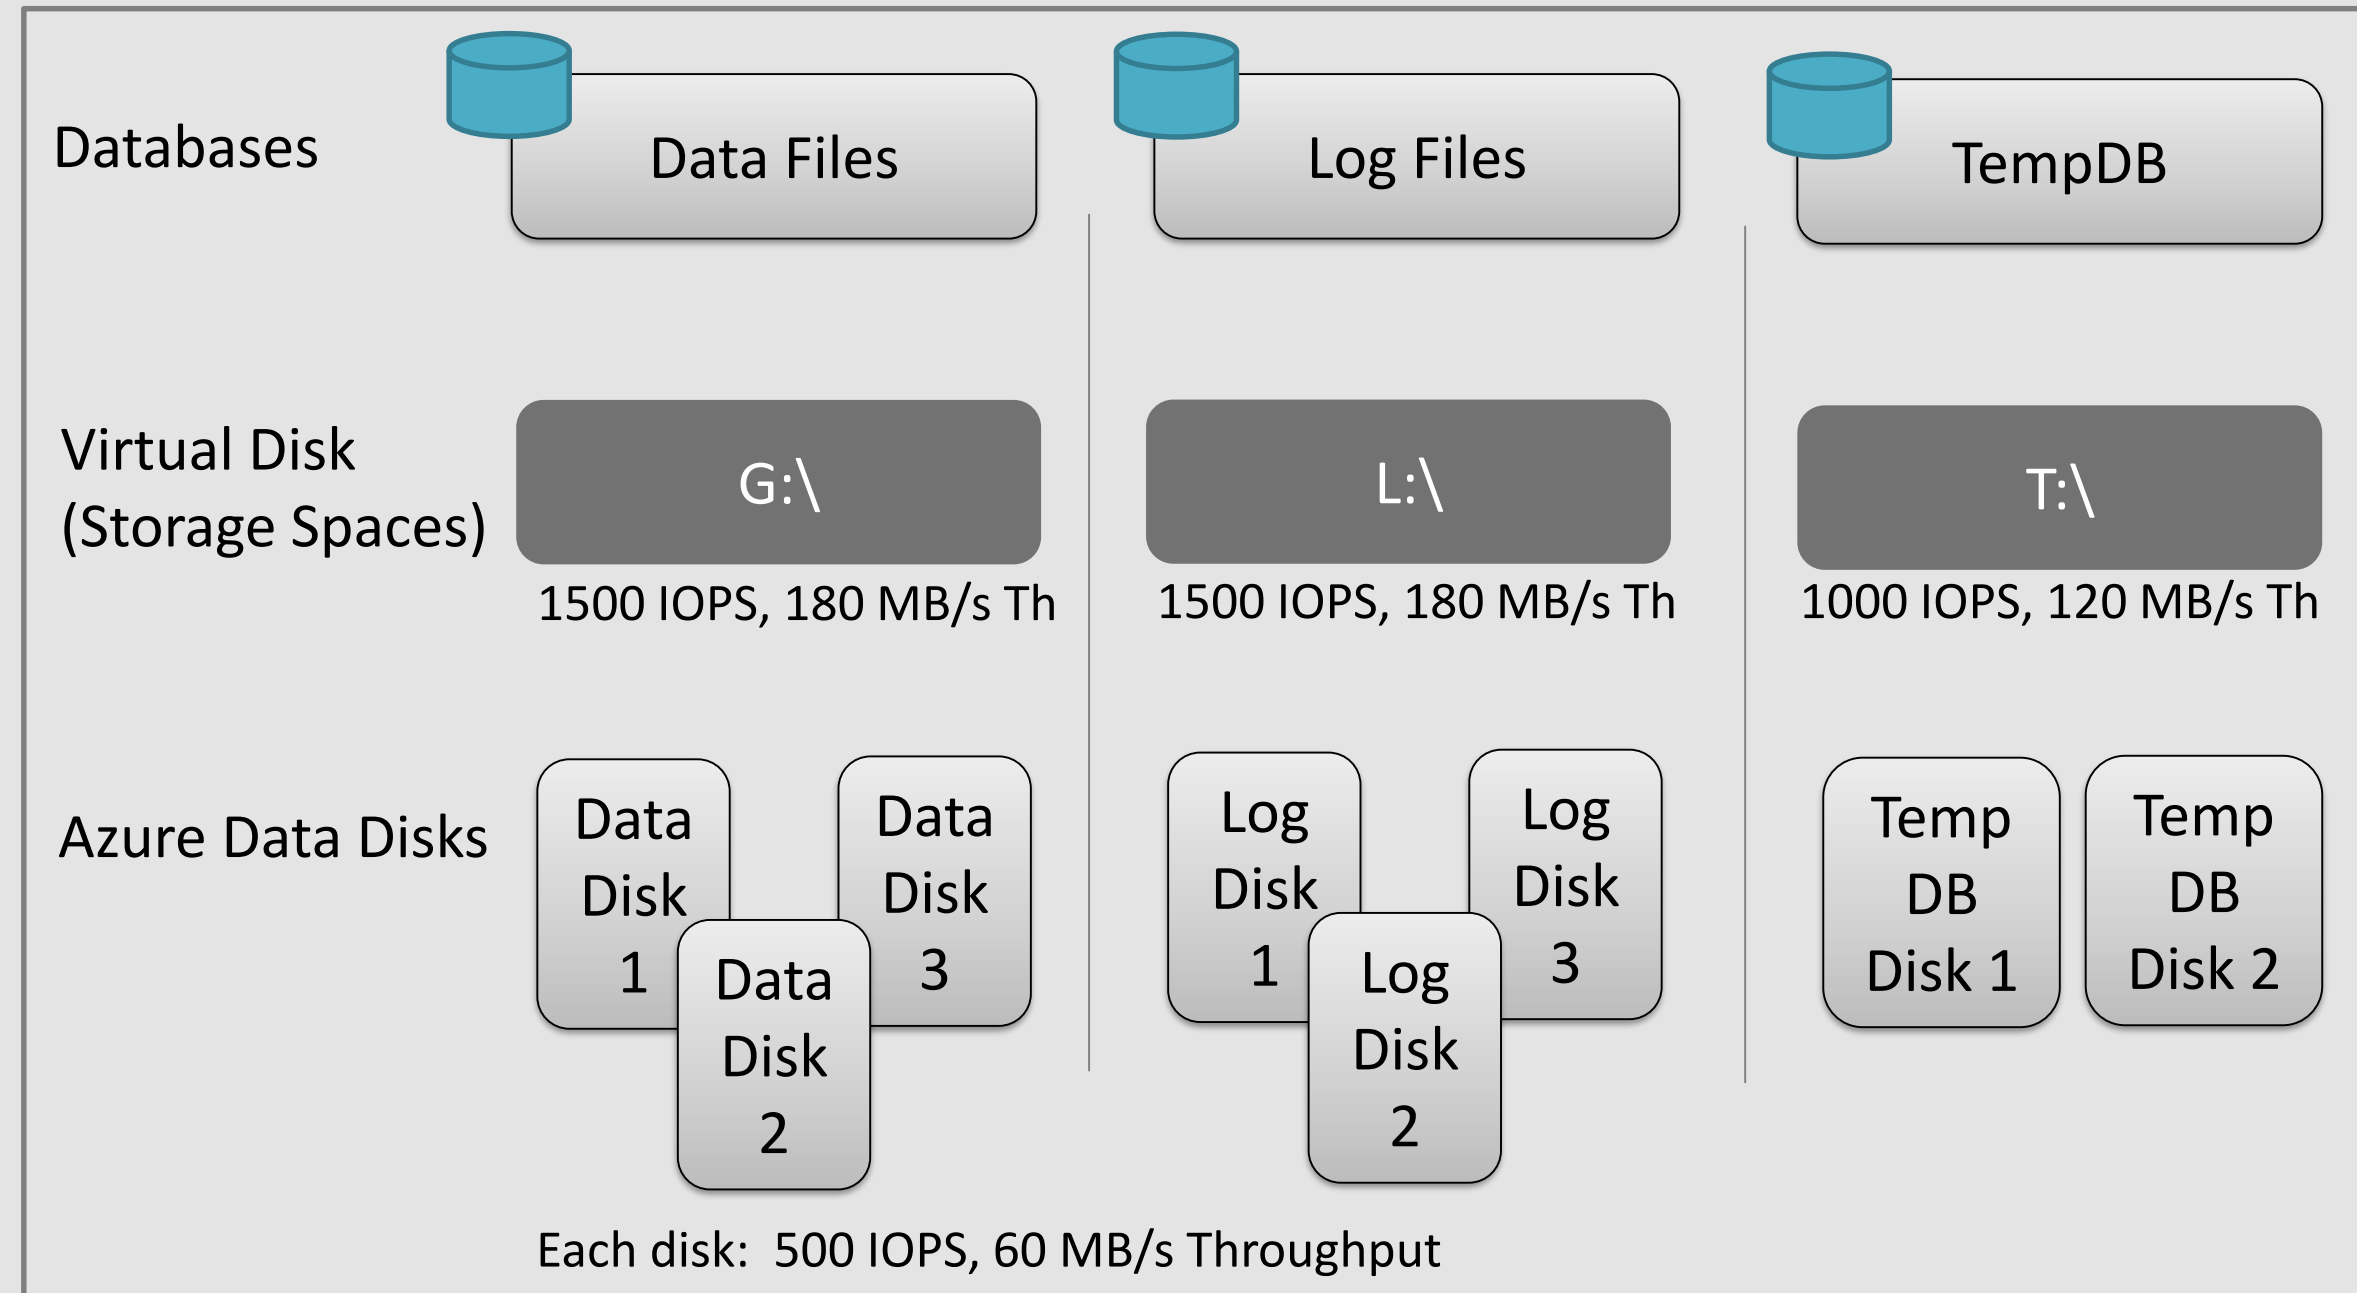
We are using **Storage Spaces (aka Virtual Disks)**: multiple Azure data disks in a storage pool to collectively share throughput and IOPS, limits.

This aggregation of disks is helpful due to lower Standard limits on IOPS & throughput.

(Another alternative to increase disk performance is to spread separate database files on different disks.)

IOPS = # of requests per second (OLTP)

Throughput (Bandwidth) = IOPS x I/O Size per specified interval (DW)



Adapted from:

<https://blogs.msdn.microsoft.com/sqlcat/2013/06/17/performance-guidance-for-sql-server-in-windows-azure-virtual-machines/>

VM Disks

(3/8)

Step 1: Check max # of disks which can be attached to the VM

Step 2: Add disks to VM (Azure portal or PowerShell)

Step 3: Add storage pool with disk striping
Stripe size (virtual disk interleave) for DW workload: 256K
Volume allocation unit size: 64K

Step 4: Create folders + reassign SQL Server pathing to your preference

Max # of disks which can be added

Step 1

Size	Cores	Memory (GB)	Data disks	Max IOPS	Local SSD (GB)	Estimated Cost (USD/MONTH)
DS12_V2 Standard	4	28	8	12800	56	485.09
DS13_V2 Standard	8	56	16	25600	112	872.71
DS14_V2 Standard	16	112	32	50000	224	1,570.58
DS15_V2 Standard	20	140	40	62500	280	1,963.42

VM Disks

Step 2: Add disks to VM
(Azure portal or PowerShell)

BISQLVM1Dev - Disks
Virtual machine

Search (Ctrl+*/*)

Overview
Activity log
Access control (IAM)
Tags
Diagnose and solve problems

SETTINGS

Availability set
Disks
Extensions
Network interfaces
Size
Backup
SQL Server configuration
Properties

OS disk

NAME	SIZE	STORAGE ACCOUNT TYPE	ENCRYPTION	HOST CACHING
BISQLVM1Dev		Standard_LRS	Not enabled	Read/write

Data disks

LUN	NAME	SIZE	STORAGE ACCOUNT TYPE	ENCRYPTION	HOST CACHING
1	bisqlvm1datadisk1dev	550 GIB	Standard_LRS	Not enabled	Read-only
2	bisqlvm1datadisk2dev	550 GIB	Standard_LRS	Not enabled	Read-only
3	bisqlvm1datadisk3dev	550 GIB	Standard_LRS	Not enabled	Read-only

+ Add data disk

bisqlvm1datadisk1dev (unmanaged)
BISQLVM1Dev

Save Discard

NAME
bisqlvm1datadisk1dev

LOGICAL UNIT NUMBER (LUN)
1

STORAGE ACCOUNT TYPE
Standard_LRS

* Size (GiB) ●
550

ESTIMATED PERFORMANCE ●

IOPS limit 500

Throughput limit (MB/s) 60

SOURCE IMAGE
MicrosoftSQLServer / SQL2016SP1-WS2016 / SQLDEV / latest

* Host caching ●
Read-only

VHD URI
<https://bisqlvm1datadisk1dev.blob.core.windows.net/vhds/bisqlvm1datadisk1dev.vhd>

Note the single disk performance before it's collectively shared

(4/8)

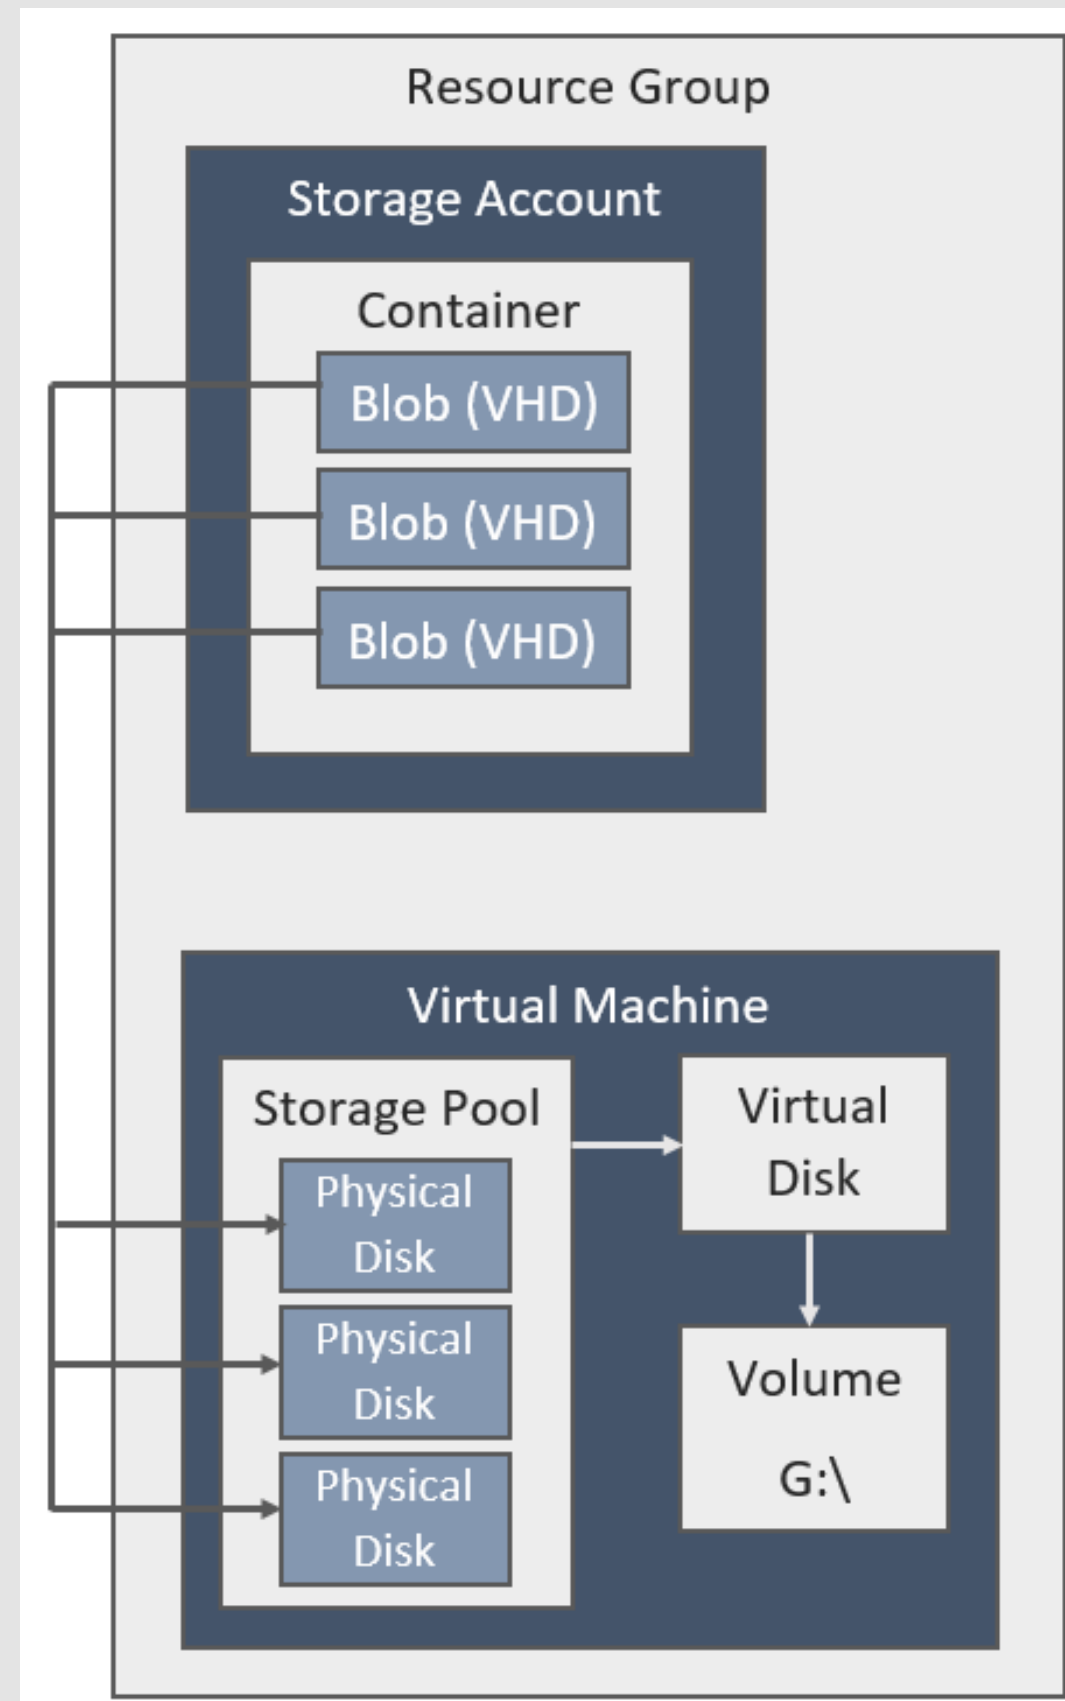
VM Disks

(5/8)

Step 3: Add storage pool + virtual disk + volume in the VM

Step 4: Create folders + reassign SQL Server pathing

- G:\MSSQL\Data
- G:\MSSQL\Backup
- G:\OLAP\Data
- G:\OLAP\Backup
- L:\MSSQL>ErrorLog
- L:\MSSQL\Log
- L:\MSSQL\Dump
- L:\OLAP\Log
- L:\OLAP\Dump
- T:\MSSQL\Data
- T:\OLAP\Temp



The screenshot shows the Windows Server Manager interface for File and Storage Services. The **Storage Pools** section is active, displaying a table of storage pools. The **Physical Disks** section is also visible, showing a table of physical disks.

Name	Type	Managed by	Available to	Read-Write Server	Capacity	Free Space	Percent Allocated
Primordial	Available Disks	BISQLVM1Dev	BISQLVM1Dev	BISQLVM1Dev			
SQLVMStoragePool1	Storage Pool	BISQLVM1Dev	BISQLVM1Dev	BISQLVM1Dev	1,022 GB	0.00 B	

Slot	Name	Status	Capacity	Bus	Usage	Characteristics
	Msf Virtual Disk (BISQLVM1Dev)		550 GB	SAS	Automatic	Integ
	Msf Virtual Disk (BISQLVM1Dev)		550 GB	SAS	Automatic	Integ
	Msf Virtual Disk (BISQLVM1Dev)		550 GB	SAS	Automatic	Integ

VM Disks

End Result:
8 data disks
attached to
the Azure
VM

We're going to
handle encryption
next

(6/8)

The screenshot shows the 'BISQLVM1Dev - Disks' page in the Azure portal. The left sidebar has the 'Disks' option highlighted with a red circle. The main area displays a table of disks:


OS disk				
NAME	SIZE	STORAGE ACCOUNT TYPE	ENCRYPTION	HOST CACHING
BISQLVM1Dev		Standard_LRS	Not enabled	Read/write

Data disks					
LUN	NAME	SIZE	STORAGE ACCOUNT TYPE	ENCRYPTION	HOST CACHING
1	bisqlvm1datadisk1dev	550 GIB	Standard_LRS	Not enabled	Read-only
2	bisqlvm1datadisk2dev	550 GIB	Standard_LRS	Not enabled	Read-only
3	bisqlvm1datadisk3dev	550 GIB	Standard_LRS	Not enabled	Read-only
4	bisqlvm1logdisk1dev	550 GIB	Standard_LRS	Not enabled	Read-only
5	bisqlvm1logdisk2dev	550 GIB	Standard_LRS	Not enabled	Read-only
6	bisqlvm1logdisk3dev	550 GIB	Standard_LRS	Not enabled	Read-only
7	bisqlvm1tempdbdisk1dev	550 GIB	Standard_LRS	Not enabled	Read-only
8	bisqlvm1tempdbdisk2dev	550 GIB	Standard_LRS	Not enabled	Read-only

At the bottom of the table, there is a '+ Add data disk' button.

The screenshot shows the 'vhds' container in the 'bisqlvm1datastdstrgdev' storage account. The container contains the following blobs:

- bisqlvm1datadisk1dev.vhd
- bisqlvm1datadisk2dev.vhd
- bisqlvm1datadisk3dev.vhd
- BISQLVM1Dev20170521150516.vhd
- bisqlvm1logdisk1dev.vhd
- bisqlvm1logdisk2dev.vhd
- bisqlvm1logdisk3dev.vhd
- bisqlvm1tempdbdisk1dev.vhd
- bisqlvm1tempdbdisk2dev.vhd

Disks attached
to the VM 

 Blob storage

VM Disks

(7/8)

End Result:
3 storage
pools

The screenshot shows the Server Manager interface for File and Storage Services. The 'Storage Pools' section displays three pools: DataDisksStoragePool (1.61 TB), LogDisksStoragePool (1.61 TB), and TempDBDisksStoragePool (1.07 TB). The 'Virtual Disks' section shows a DataVirtualDisk on the DataDisksStoragePool with a Simple layout and Fixed provisioning. The 'Physical Disks' section shows three Msft Virtual Disk (BISQLVM1Dev) instances, each 550 GB.

STORAGE POOLS
All storage pools | 3 total

Name	Type	Managed by	Available to	Read-Write Server	Capacity	Free Space	Percent Allocated
Windows Storage (3)							
DataDisksStoragePool	Storage Pool	BISQLVM1Dev	BISQLVM1Dev	BISQLVM1Dev	1.61 TB	0.00 B	
LogDisksStoragePool	Storage Pool	BISQLVM1Dev	BISQLVM1Dev	BISQLVM1Dev	1.61 TB	0.00 B	
TempDBDisksStoragePool	Storage Pool	BISQLVM1Dev	BISQLVM1Dev	BISQLVM1Dev	1.07 TB	0.00 B	

Last refreshed on 6/16/2017 11:02:33 AM

VIRTUAL DISKS
DataDisksStoragePool on BISQLVM1Dev

Name	Status	Layout	Provisioning	Capacity	Allocated	Volume
DataVirtualDisks		Simple	Fixed	1.61 TB	1.61 TB	G:

PHYSICAL DISKS
DataDisksStoragePool on BISQLVM1Dev

Slot	Name	Status	Capacity	Bus	Usage	Chas
	Msft Virtual Disk (BISQLVM1Dev)		550 GB	SAS	Automatic	Integ
	Msft Virtual Disk (BISQLVM1Dev)		550 GB	SAS	Automatic	Integ
	Msft Virtual Disk (BISQLVM1Dev)		550 GB	SAS	Automatic	Integ

DataVirtualDisk
Details

Name: DataVirtualDisk

Property: Interleave

Value: 262144

Stripe size should be 256KB for a data warehousing workload

VM Disks

(8/8)

End Result:
3 volumes

G:\ - Data

L:\ - Log

T:\ - TempDB

The screenshot shows the Windows Disk Management console. At the top, a table lists the volumes. Below this, the details for five disks are shown, each with a corresponding volume selected.

Volume	Layout	Type	File System	Status	Capacity	Free Space	% Free
(C:)	Simple	Basic	NTFS	Healthy (S...	127.00 GB	99.05 GB	78 %
Data (G:)	Simple	Basic	NTFS	Healthy (P...	1646.87 GB	1646.66 GB	100 %
Log (L:)	Simple	Basic	NTFS	Healthy (P...	1646.87 GB	1646.68 GB	100 %
TempDB (T:)	Simple	Basic	NTFS	Healthy (P...	1097.87 GB	1097.70 GB	100 %
Temporary Storage (D:)	Simple	Basic	NTFS	Healthy (P...	56.00 GB	51.65 GB	92 %

Disk	Volume	Capacity	File System	Status
Disk 0	(C:)	127.00 GB	NTFS	Healthy (System, Boot, Active, Crash Dump, Primary Partition)
Disk 1	Temporary Storage (D:)	56.00 GB	NTFS	Healthy (Page File, Primary Partition)
Disk 6	TempDB (T:)	1097.87 GB	NTFS	Healthy (Primary Partition)
Disk 7	Data (G:)	1646.87 GB	NTFS	Healthy (Primary Partition)
Disk 11	Log (L:)	1646.87 GB	NTFS	Healthy (Primary Partition)

Legend: ■ Unallocated ■ Primary partition

Encryption – Data at Rest

Storage Service Encryption (SSE) ← associated with the storage account

Disk Encryption (BitLocker) ← inside the VM

More info:

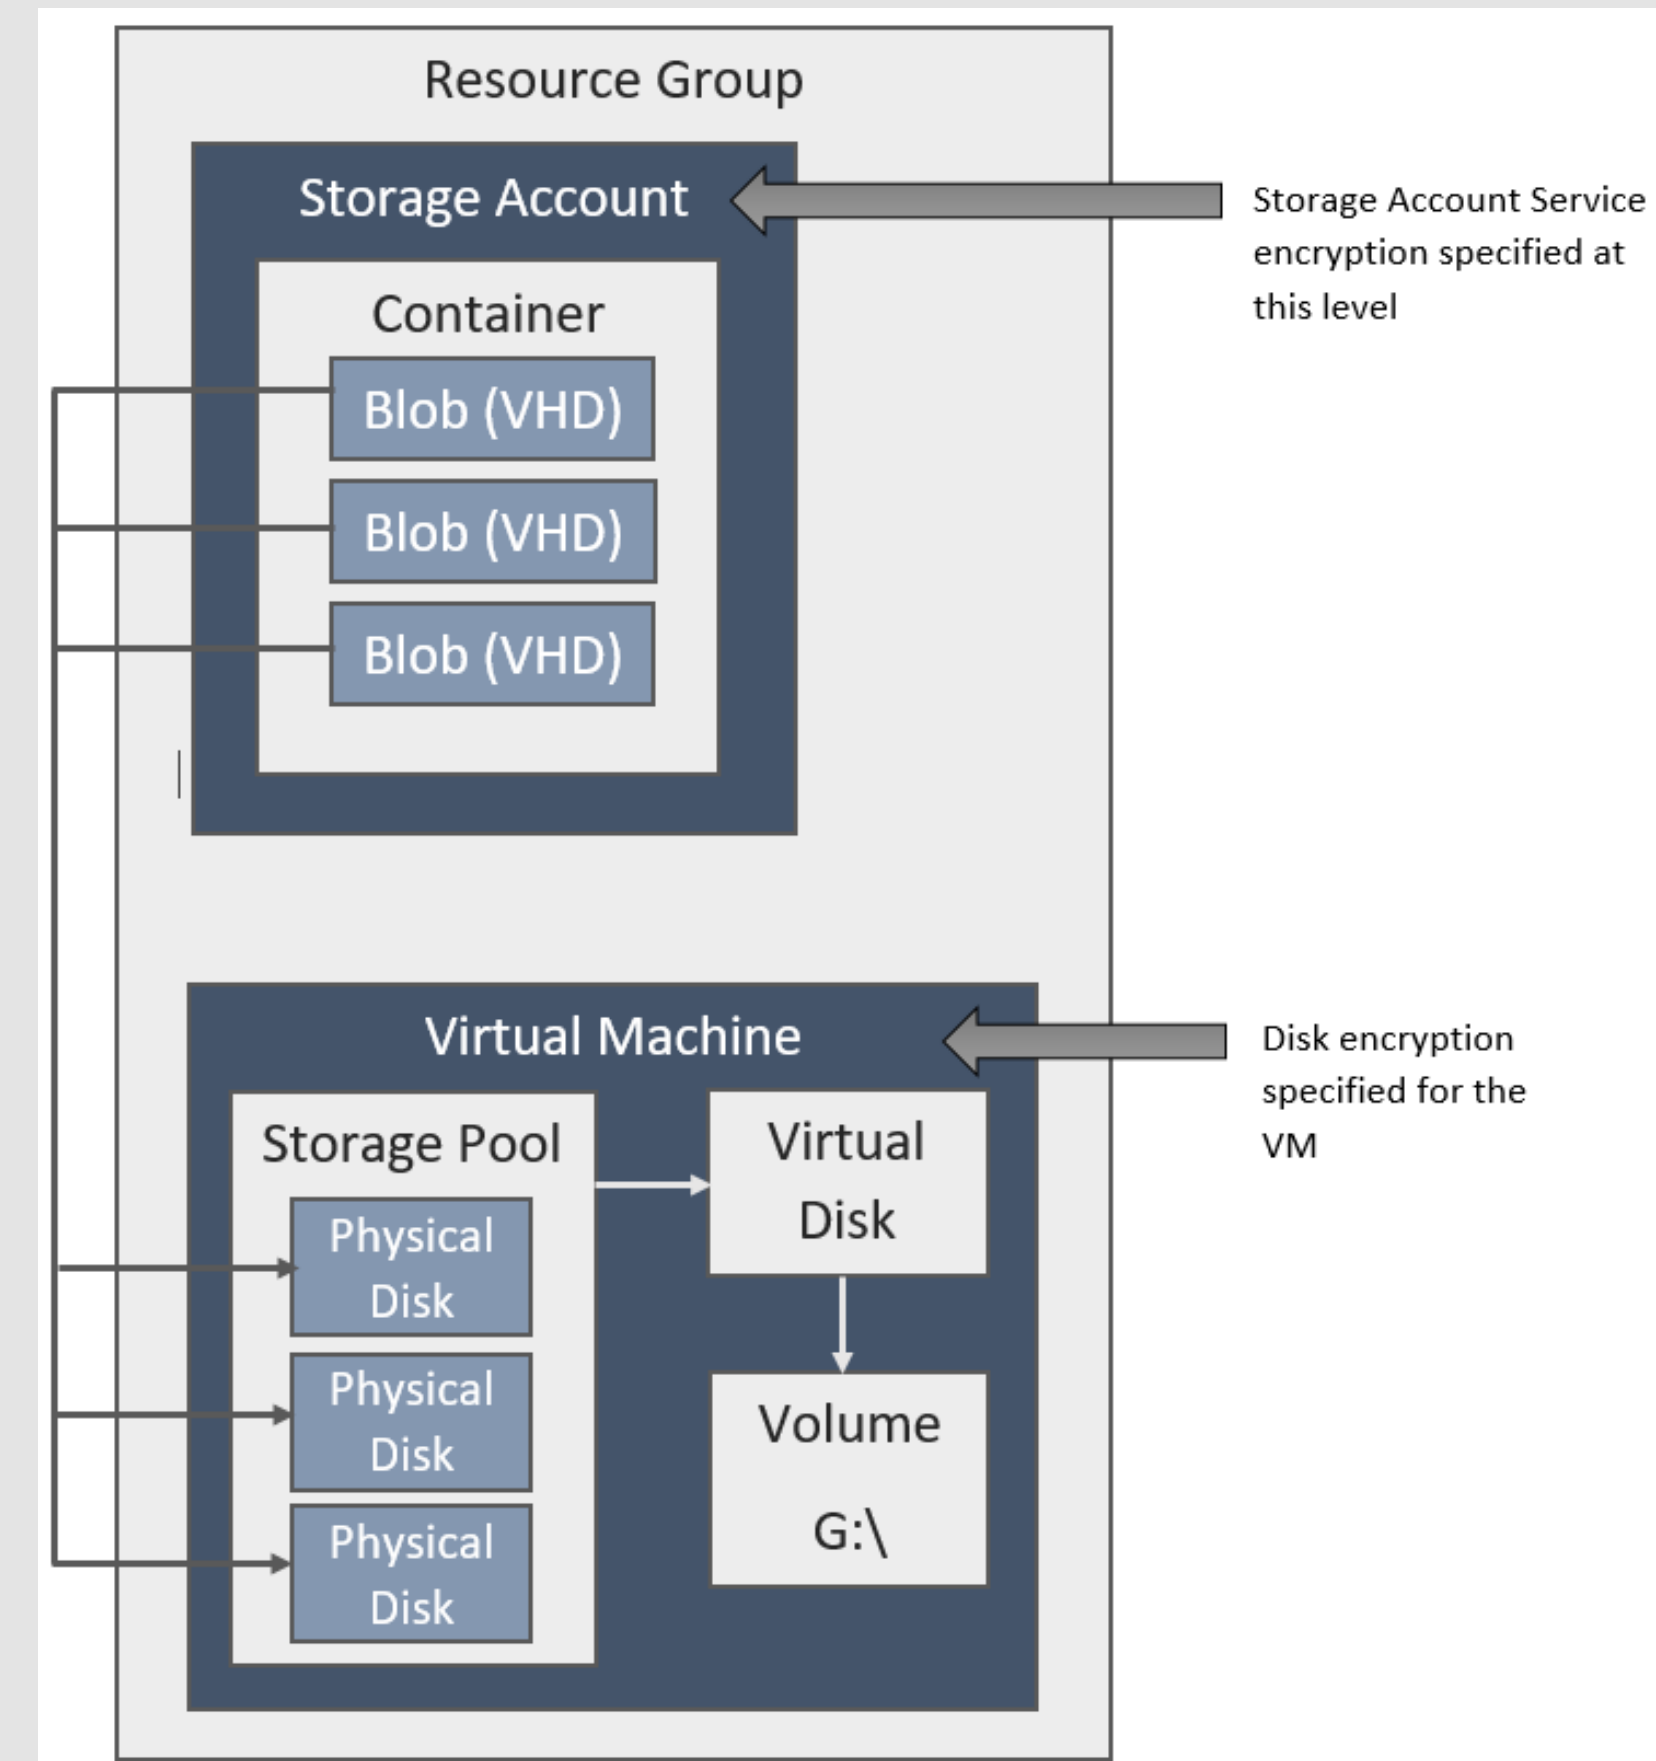
<http://www.sqlchick.com/entries/2017/2/5/deciding-on-encryption-at-rest-for-an-azure-virtual-machine>

<http://www.sqlchick.com/entries/2017/2/9/setting-up-azure-disk-encryption-for-a-virtual-machine-with-powershell>

Other options to consider for encrypting the database:

TDE (Transparent Data Encryption)

AE (Always Encrypted)



Disk Encryption

(1/2)

Step 1: Create a Service Principal

Step 2: Provision Azure Key Vault

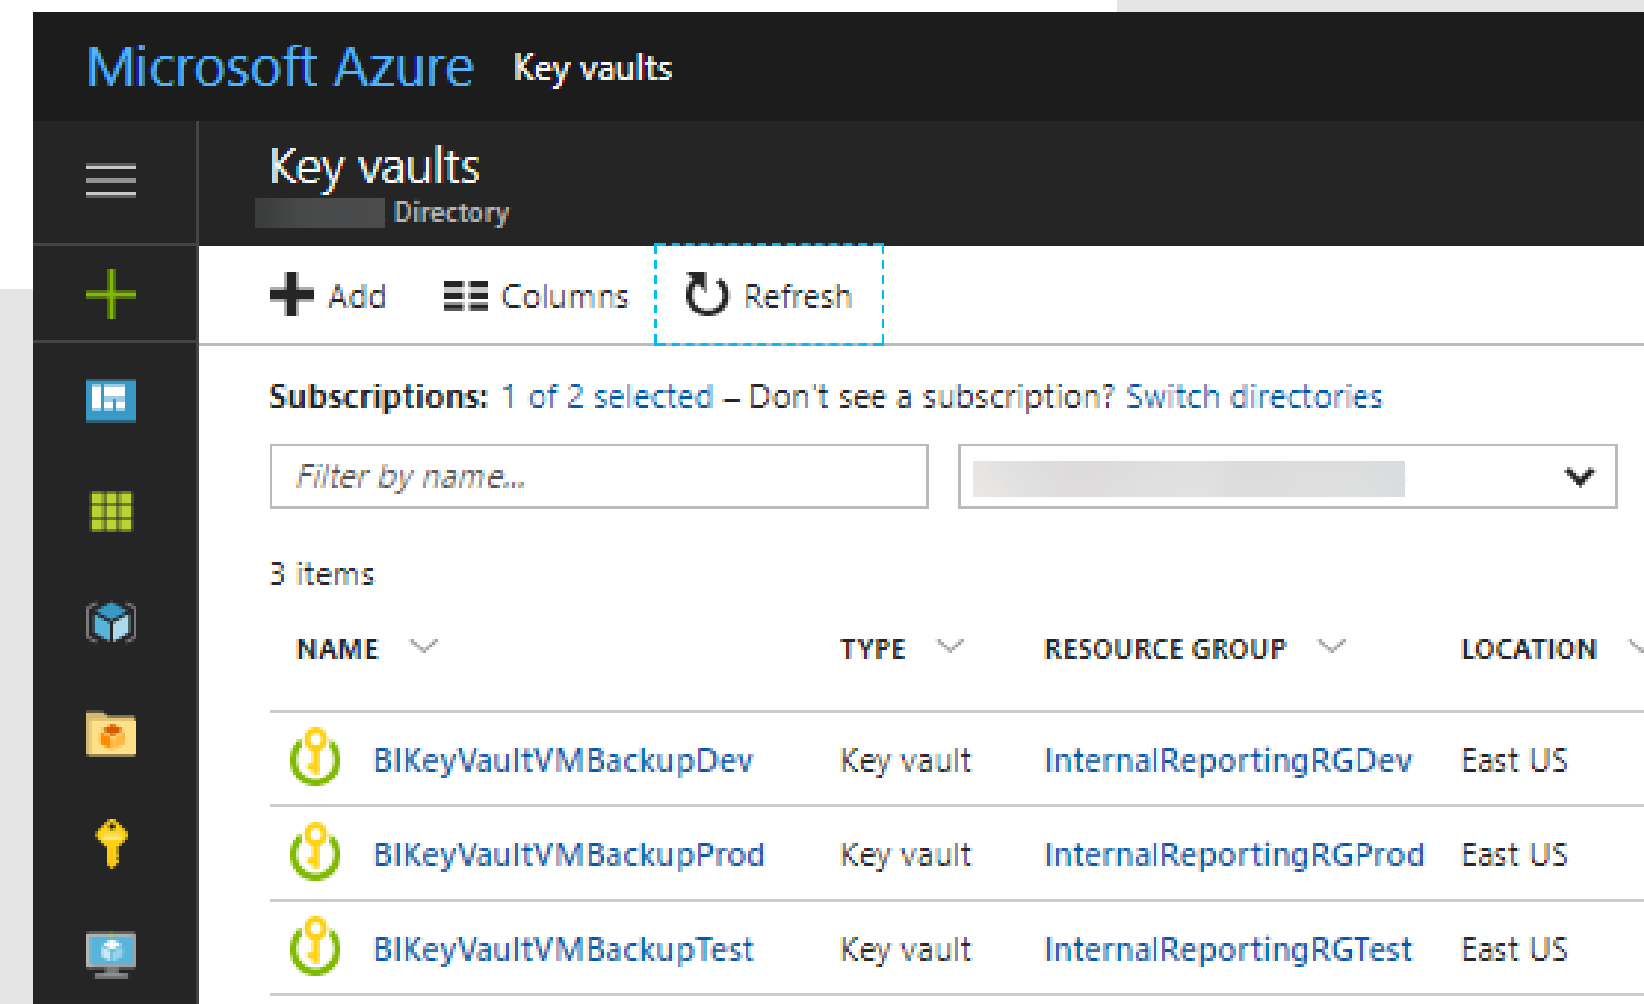
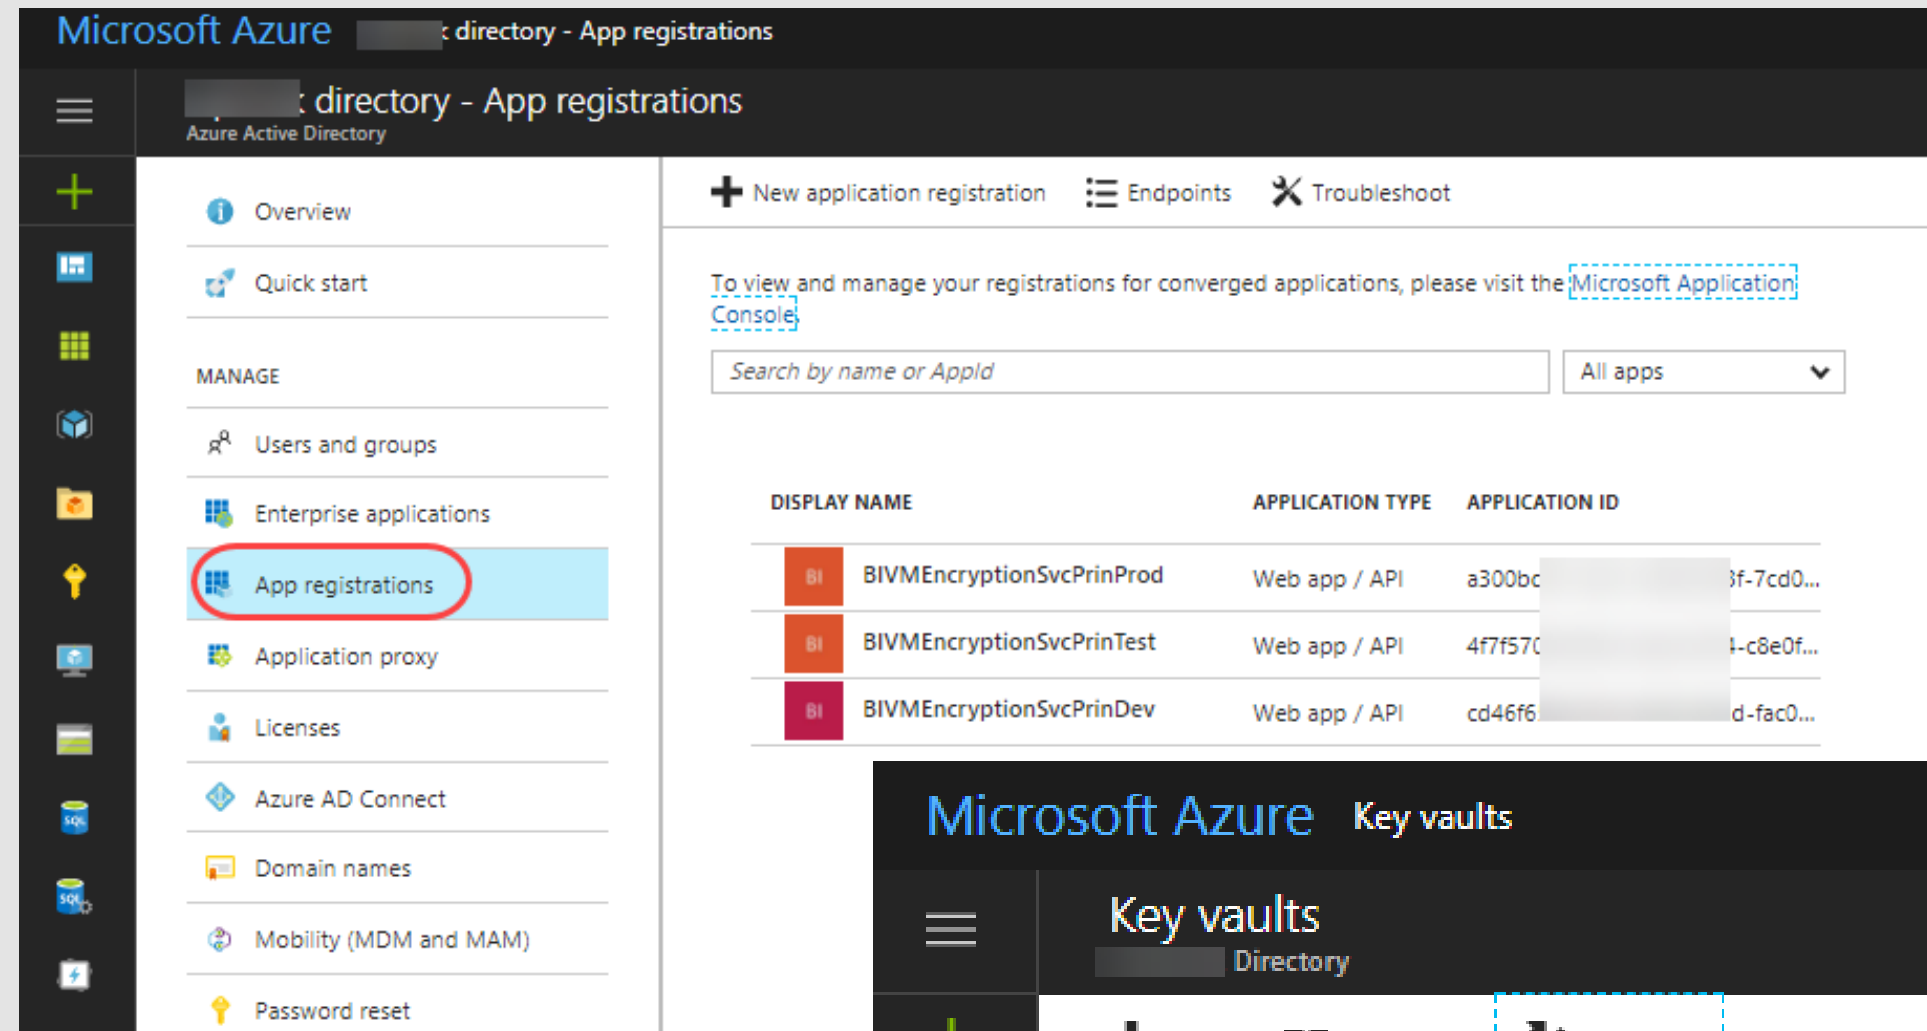
Step 3: Enable the Service Principal to communicate with Key Vault

Step 4: Enable Disk Encryption

PowerShell scripts here:

<http://www.sqlchick.com/entries/2017/2/5/deciding-on-encryption-at-rest-for-an-azure-virtual-machine>

<http://www.sqlchick.com/entries/2017/2/9/setting-up-azure-disk-encryption-for-a-virtual-machine-with-powershell>



Disk Encryption

(2/2)

End Result: All drives are encrypted.

A new Bek volume is created which is where the encryption key is read from:

BISQLVM1Dev - Disks
Virtual machine

Search (Ctrl+/) Edit

Overview
Activity log
Access control (IAM)
Tags
Diagnose and solve problems

SETTINGS
Availability set
Disks
Extensions
Network interfaces

OS disk		SIZE	STORAGE ACCOUNT TYPE	ENCRYPTION	HOST CACHING
NAME					
BISQLVM1Dev			Standard_LRS	Enabled	Read/write

Data disks		SIZE	STORAGE ACCOUNT TYPE	ENCRYPTION	HOST CACHING
LUN	NAME				
1	bisqlvm1datadisk1dev	550 GiB	Standard_LRS	Enabled	Read-only
2	bisqlvm1datadisk2dev	550 GiB	Standard_LRS	Enabled	Read-only
3	bisqlvm1datadisk3dev	550 GiB	Standard_LRS	Enabled	Read-only
4	bisqlvm1logdisk1dev	550 GiB	Standard_LRS	Enabled	None
5	bisqlvm1logdisk2dev	550 GiB	Standard_LRS	Enabled	None
		550 GiB	Standard_LRS	Enabled	None
		550 GiB	Standard_LRS	Enabled	Read-only
		550 GiB	Standard_LRS	Enabled	Read-only

+ Add data disk

Disk 2
Basic
47 MB
Online

Bek Volume (E:)
47 MB NTFS
Healthy (Primary Partition)

More info: <https://blogs.msdn.microsoft.com/mast/2016/11/27/azure-disk-encryption-how-to-recover-bek-file-from-azure-key-vault/>

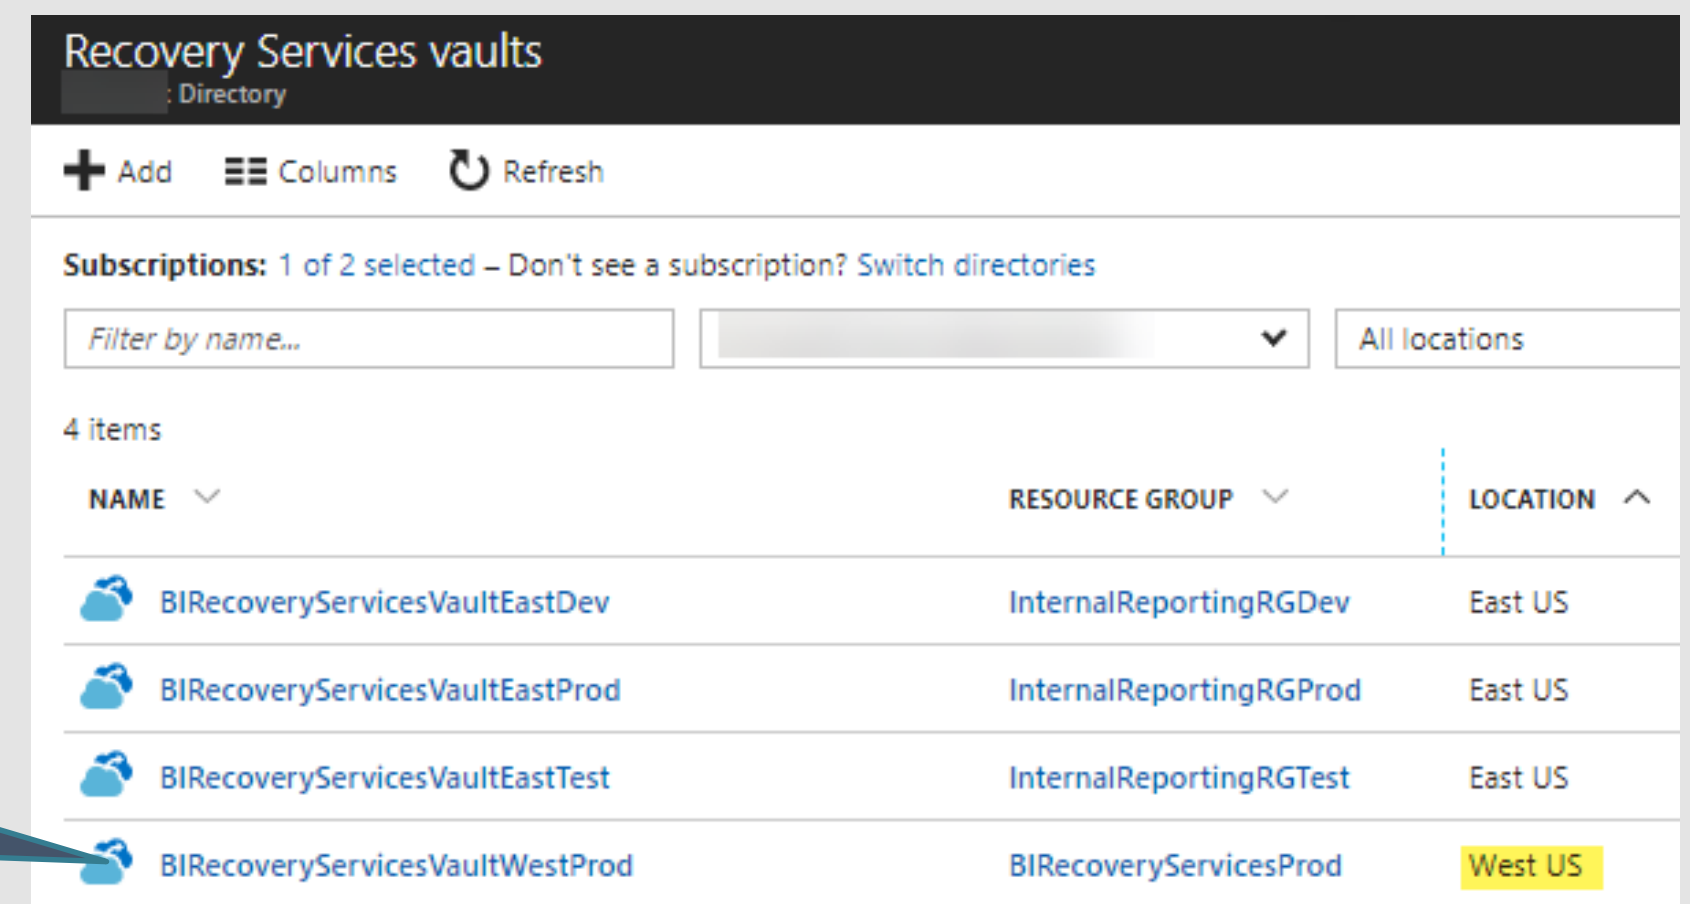
Disaster Recovery

We (the Analytics Team) don't want to incur the cost or complexity of supporting an availability set. Therefore, we are doing the following for DR purposes:

1. Documented steps for re-provisioning & redeploying
2. Use of Premium storage in Production – this gets us a 99.9% uptime SLA for a single instance Azure VM
<https://azure.microsoft.com/en-us/support/legal/sla/virtual-machines/>

3. For Prod: we are considering use of Azure Site Recovery (still in preview)
<https://docs.microsoft.com/en-us/azure/site-recovery/site-recovery-azure-to-azure>

Important to create your Recovery Services vault in a different region (the RG and the vault itself) than the source objects



NAME	RESOURCE GROUP	LOCATION
BIRecoveryServicesVaultEastDev	InternalReportingRGDev	East US
BIRecoveryServicesVaultEastProd	InternalReportingRGProd	East US
BIRecoveryServicesVaultEastTest	InternalReportingRGTest	East US
BIRecoveryServicesVaultWestProd	BIRecoveryServicesProd	West US

Automation & Scheduling

Daily & Weekly Schedule

Analytics Environment Scheduling
All times EST

Production:

12am	1am	2am	3am	4am	5am	6am	7am	8am	9am	10am	11am	12pm	1pm	2pm	3pm	4pm	5pm	6pm	7pm	8pm	9pm	10pm	11pm				
Engine+SSIS VM shut down					Auto start																		Engine+SSIS VM shut down				
												EST standard business hours															
										UK standard business hours																	
ADF job						Daily ETL	Hourly ETL & SSAS processing											Auto patching (Fri)	SSIS maint (Daily)	SQL backup (Daily)	VM bck to vault (Daily)						

Dev and Test:

12am	1am	2am	3am	4am	5am	6am	7am	8am	9am	10am	11am	12pm	1pm	2pm	3pm	4pm	5pm	6pm	7pm	8pm	9pm	10pm	11pm											
Both Dev & QA machines shut down								Manual start (when in use)																						Both Dev & QA machines shut down				
								Manual (on-demand) ETL & SSAS processing											Auto patching (Fri)	SSIS maint (Daily)	SQL backup (Daily)	VM bck to vault (Fri)												

Auto patching: requires VM to be on.

VM backup to vault: not not require VM to be on.

Automatic Shutdown

The screenshot displays the Azure portal interface for a virtual machine named 'BISQLVM1Dev'. The left-hand navigation pane is visible, with the 'Auto-shutdown' option under the 'SCHEDULES' section highlighted with a red circle. The main content area shows the configuration for the auto-shutdown feature, which is currently 'Enabled'. The 'Scheduled shutdown' is set to '9:00:00 PM' in the '(UTC-05:00) Eastern Time (US & Canada)' time zone. A notification preference is set to 'No' for sending a notification 15 minutes before the shutdown. A 'Webhook URL' field is present but empty. At the top of the main area, there is a promotional message: 'Like auto-shutdown? DevTest Labs has more features to control costs and streamline your development workflows. [Learn more.](#)'

Dev & Test

We are using the **auto-shutdown** feature within the VM properties (Dev/Test Labs).

Production

We are using Azure Automation to **shut down, and restart**, the 1st VM (DB engine/SSIS) overnight since our ETL runs 1x/hour during business hours only. This saves ~33% of compute cost.

The 2nd VM (SSAS) stays on 24/7.

Automatic Patching

We are using the Windows and SQL auto-patching feature. The VM “Manage Updates” feature is in private preview (July 2017).

SQL Server configuration

SQL connectivity level

SQL connectivity: Private (within Virtual Network)

Port: 1433

SQL Authentication: Disabled

Automated patching

Automated patching: Enable

Maintenance schedule: Friday

Maintenance start hour (local time): 22:00

Maintenance window duration (minutes): 60

Automated backup

Automated backup: Enabled

Retention period (days): 7

Storage account: bisqlvm1bckstdstrqdev

Encryption: Enabled

Backup system databases: Enabled

Configure backup schedule: Manual

Full backup frequency: Daily

Full backup start time (local VM time): 01:00

Full backup time window (hours): 1

Log backup frequency (minutes): 60

Login Properties - NT Service\SQLIaaSExtension

Select a page

Server Roles

Server role is used to grant server-wide security privileges to a user.

Server roles:

- bulkadmin
- dbcreator
- diskadmin
- processadmin
- public
- securityadmin
- serveradmin
- setupadmin
- sysadmin

Administrators Properties

General

Administrators

Description: Administrators have complete and unrestricted access to the computer/domain

Members:

- \Domain Admins
- ybenton
- vmcoates
- NT SERVICE\SQLIaaSExtension (S-1-5-80-2048003730-25558169...)

Changes to a user's group membership are not effective until the next time the user logs on.

OK Cancel Apply Help

More info: <https://docs.microsoft.com/en-us/azure/virtual-machines/windows/sql/virtual-machines-windows-sql-server-agent-extension>

VM Backups

The VM backups are managed in the Recovery Services Vault.

For the VM backups to be encrypted, disk encryption is required (discussed in the previous section).

The screenshot displays the Azure Recovery Services console for a backup item named 'bisqlvm1dev'. The interface is split into two main sections: 'Essentials' and 'Restore points'.

Essentials Section:

- Recovery services vault:** BIRecoveryServicesVaultEastDev
- Subscription name:** [Redacted]
- Subscription ID:** [Redacted]
- Item type:** Azure virtual machine
- Last backup status:** Success
- Last backup time:** 2/7/2017, 10:15:42 AM
- Latest restore point:** 2/7/2017, 10:16:41 AM (1 hour(s) ago)
- Oldest restore point:** 2/7/2017, 10:16:41 AM (1 hour(s) ago)
- Backup policy:** BIBackupPolicyWeekly

Restore points Section:

- Filtered for last 30 days
- Consistency options: CRASH CONSISTENT, APPLICATION CONSISTENT, FILE-SYSTEM CONSISTENT
- Filter items... [Input field]
- All restore points [Dropdown menu]

TIME	CONSISTENCY	ENCRYPTION	
2/7/2017, 10:16:41 AM	Application Consistent	Encrypted	...

SQL Backups: From VM

We are using the scheduled SQL Server backups feature.

The screenshot shows the Azure Blob service interface. On the left, the 'Essentials' section has a search bar for containers. The main area shows a container named 'backupcontainer' with a list of backup files. The files are named with a prefix 'EnterpriseDW_de602e2435424913adc7d836787f47c7_20170615180359-04.bak' and similar identifiers.

Be sure that this storage account is set to be encrypted – the “automaticbackup” container contains certificates & keys.

The screenshot shows the 'SQL Server configuration' window in a virtual machine named 'BISQLVM1Dev'. The 'Automated backup' section is expanded, showing the following settings:

Setting	Value
Automated backup	Enabled
Retention period (days)	10
Storage account	'bisqlvm1bckstdstrqdev'
Encryption	Enabled
Backup system databases	Enabled
Configure backup schedule	Manual
Full backup frequency	Daily
Full backup start time (local VM time)	01:00
Full backup time window (hours)	1
Log backup frequency (minutes)	60

A red box highlights the 'Full backup frequency' and 'Full backup start time (local VM time)' settings. A callout bubble points to the '01:00' value, indicating it is 9pm EST.

SQL Backups: From SQLDB

Backups for SQLDB depend on the pricing tier selected:

Full backups: weekly

Differential: every few hours

T-log: every 5-10 minutes

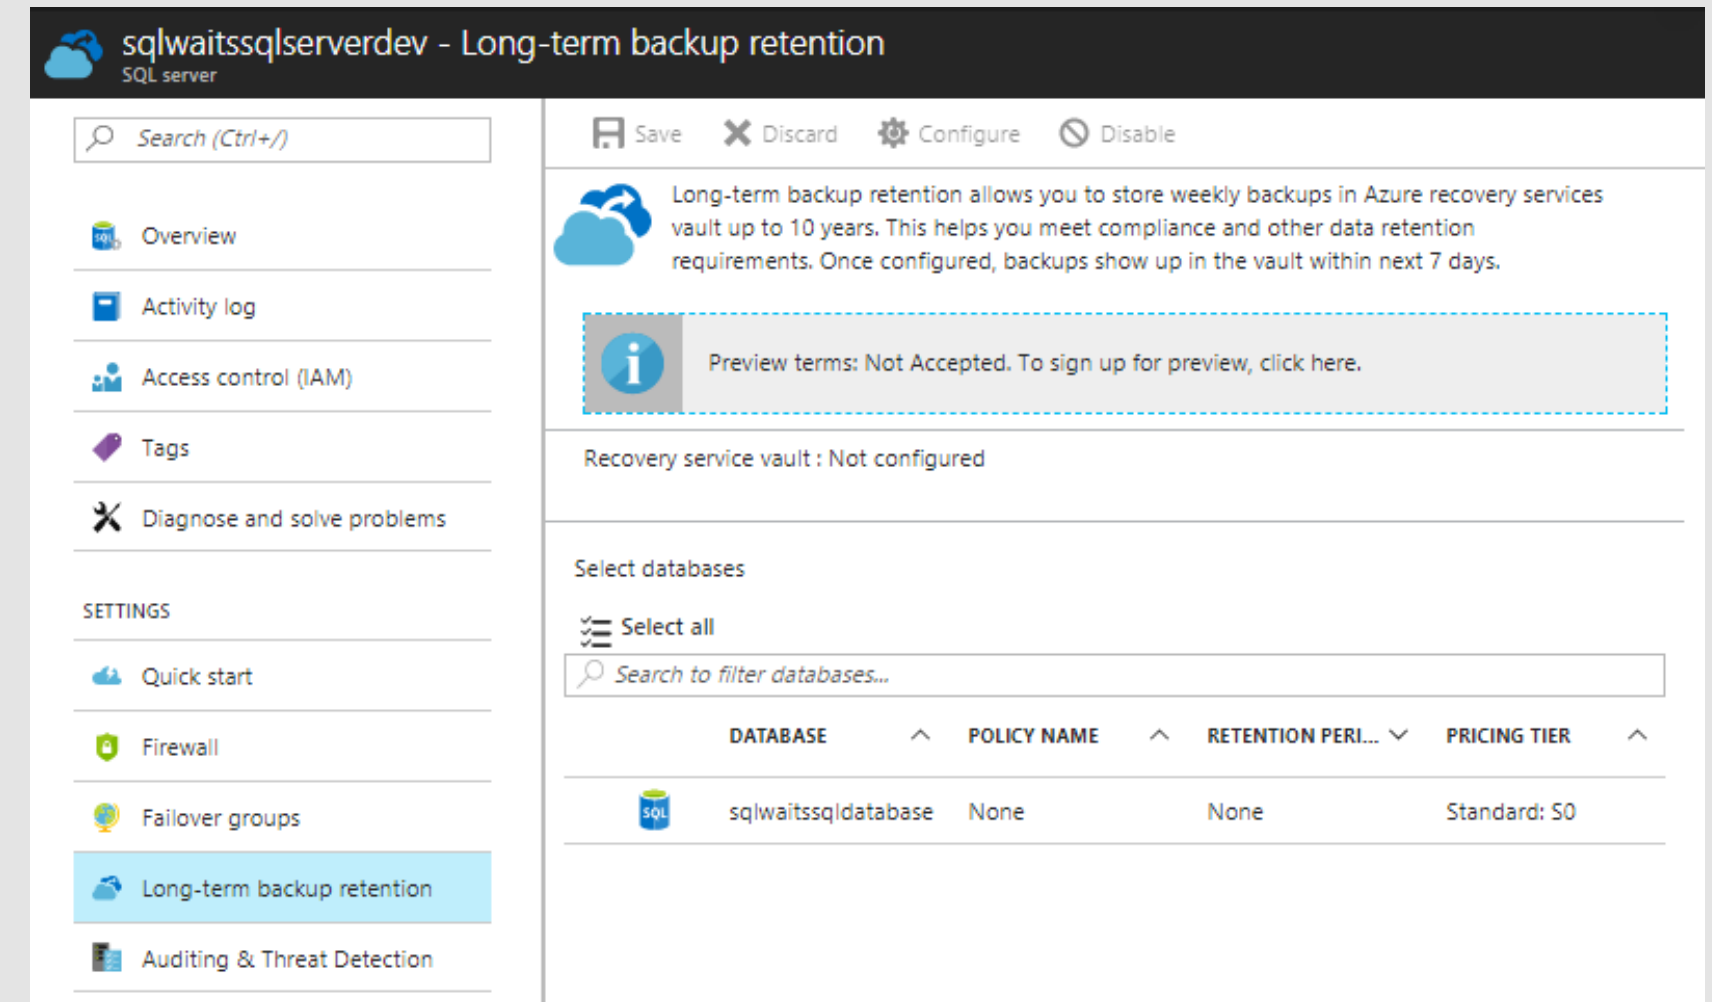
Basic service: retained 7 days

Standard service: retained 35 days

Premium service: retained 35 days

Longer-term backup retention:

- Set up a process to back up the .bacpac file (schema + data), or
- Recovery Services vault (still in Preview):



sqlwaitssqlserverdev - Long-term backup retention

SQL server

Search (Ctrl+)

Save Discard Configure Disable

Long-term backup retention allows you to store weekly backups in Azure recovery services vault up to 10 years. This helps you meet compliance and other data retention requirements. Once configured, backups show up in the vault within next 7 days.

Preview terms: Not Accepted. To sign up for preview, click here.

Recovery service vault : Not configured

Select databases

Select all

Search to filter databases...

DATABASE	POLICY NAME	RETENTION PERI...	PRICING TIER
sqlwaitssqldatabase	None	None	Standard: S0

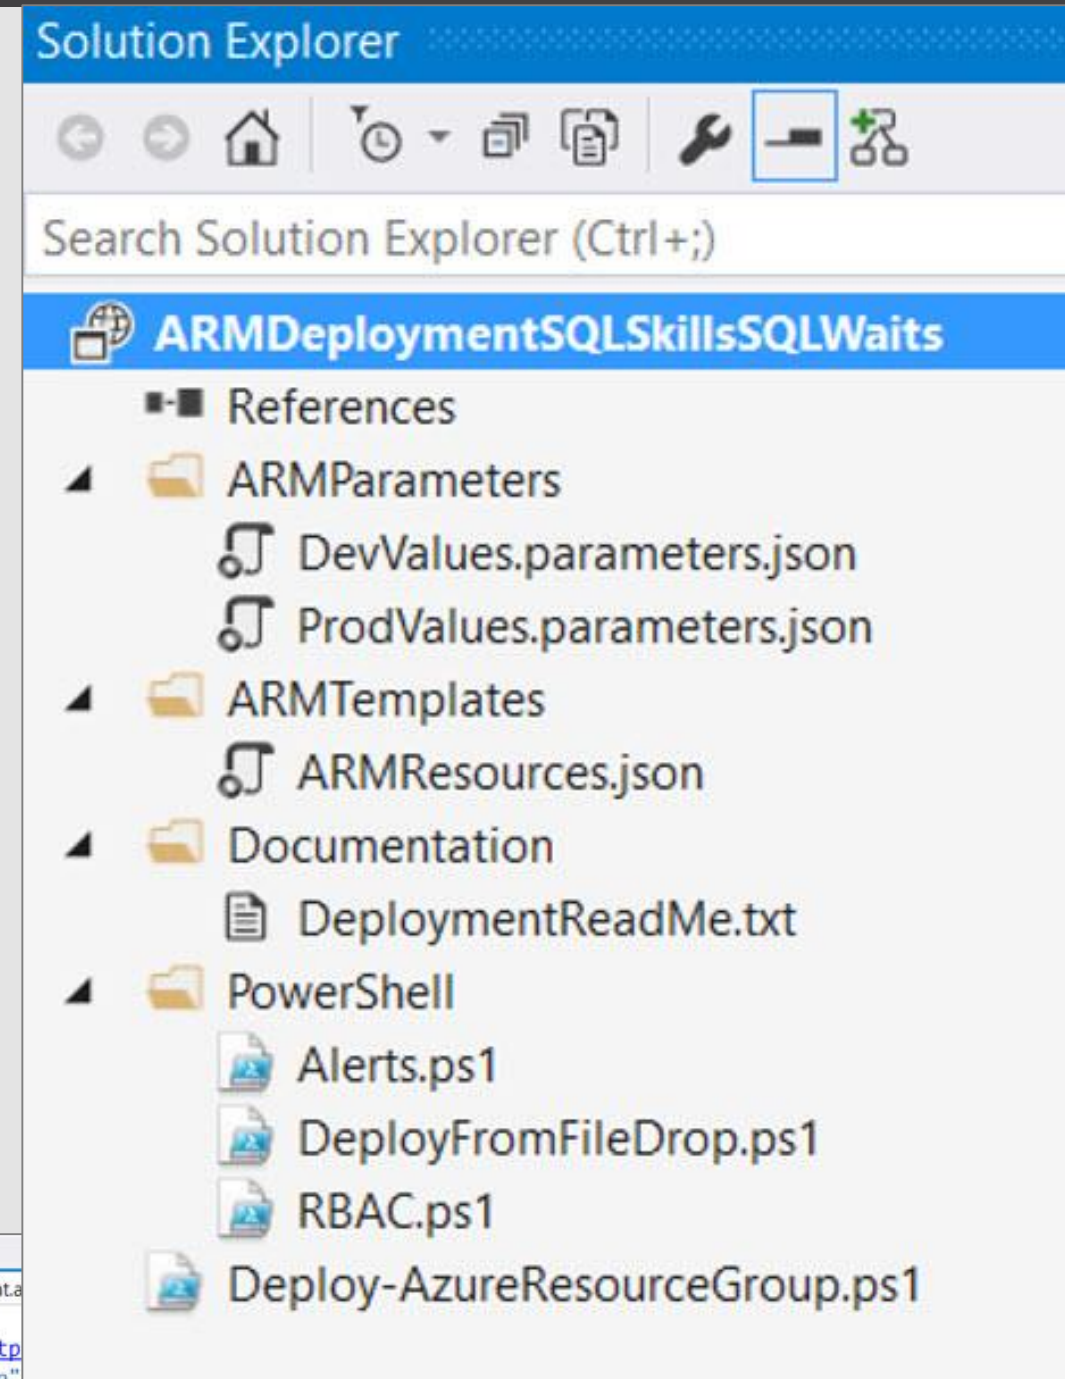
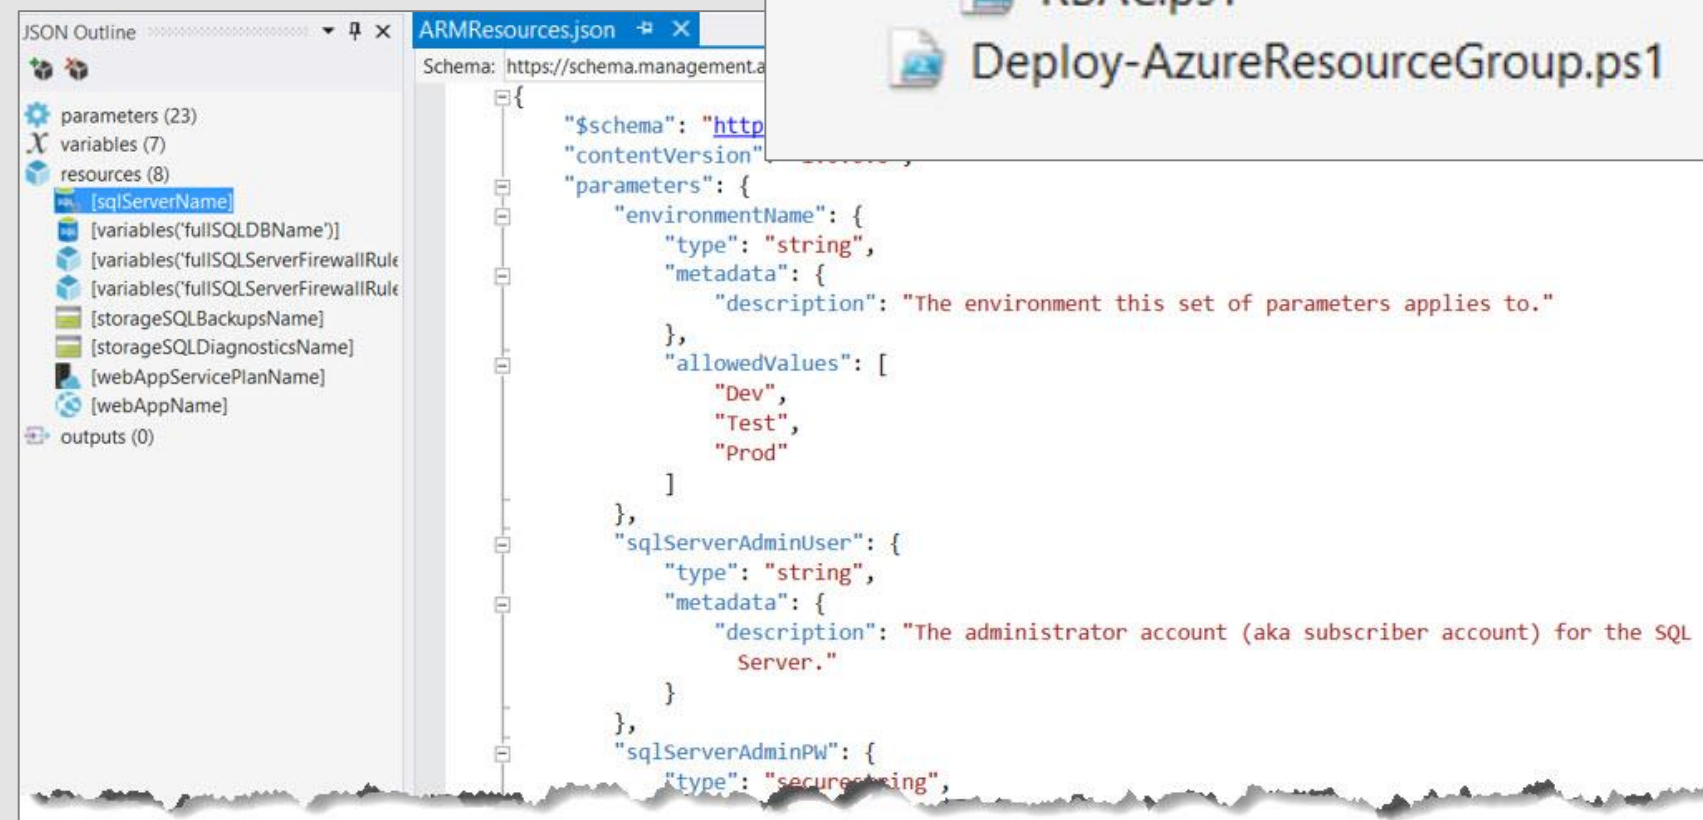
More info: <https://blogs.msdn.microsoft.com/mast/2013/03/03/different-ways-to-backup-your-windows-azure-sql-database/>

ARM Templates

Goal is for deployments to be modular & repeatable.

We have invested time in learning how to do ARM templates so we can make a good decision when to use them.

We have our “infrastructure as code” checked into source control within a Visual Studio project.



Blog post on ARM deployments:
<https://blogs.sentryone.com/melissacoates/getting-started-with-automated-arm-deployments-in-azure/>

PowerShell

We are being selective about automation.

Balance of: wanting infrastructure in source code vs. little need for deployment repeatability.

We have some PowerShell scripts for deployment purposes, such as:

- Assign tags
- Enable disk encryption
- Provision key vault
- Deploy ARM template

```
#Input Area
$subscriptionName = 'InsertSubscriptionName'
$resourceGroupName = 'SQLSkillsWaitsLibraryRGDev'
$deploymentName = 'SQLSkillsSQLWaitsDeployment'
$templateFilePath = '0:\Deployments\SQLWaits\2017-06-07\ARMResources.json'
$parameterFilePath = '0:\Deployments\SQLWaits\2017-06-07\DevValues.parameters.json'

#ARM template deployment
New-AzureRmResourceGroupDeployment `
  -Name $deploymentName `
  -ResourceGroupName $resourceGroupName `
  -TemplateFile $templateFilePath `
  -TemplateParameterFile $parameterFilePath
```

Blog post on ARM deployments:

<https://blogs.sentryone.com/melissacoates/getting-started-with-automated-arm-deployments-in-azure/>

Recommendations for ARM and PowerShell

Use ARM for:

- Deployment of resources to Test and Prod
 - Recognition of dependencies
 - Parallel deployment of resources (faster than PowerShell or CLI)

Use PowerShell for:

- Execution of ARM template
- Management of resources (ex: start/stop a VM - in conjunction with Azure Automation)
- Selective items you would prefer in PowerShell vs. ARM (ex: Tags or Metrics Alerts)

Monitoring the Solution

Monitoring

We are currently using a combination of:



Azure Monitor (Metrics, Activity Log, Diagnostic Log)



Azure Monitor Alerts



Azure Blob Storage (Metrics & Activity Log Data)



Azure Security Center



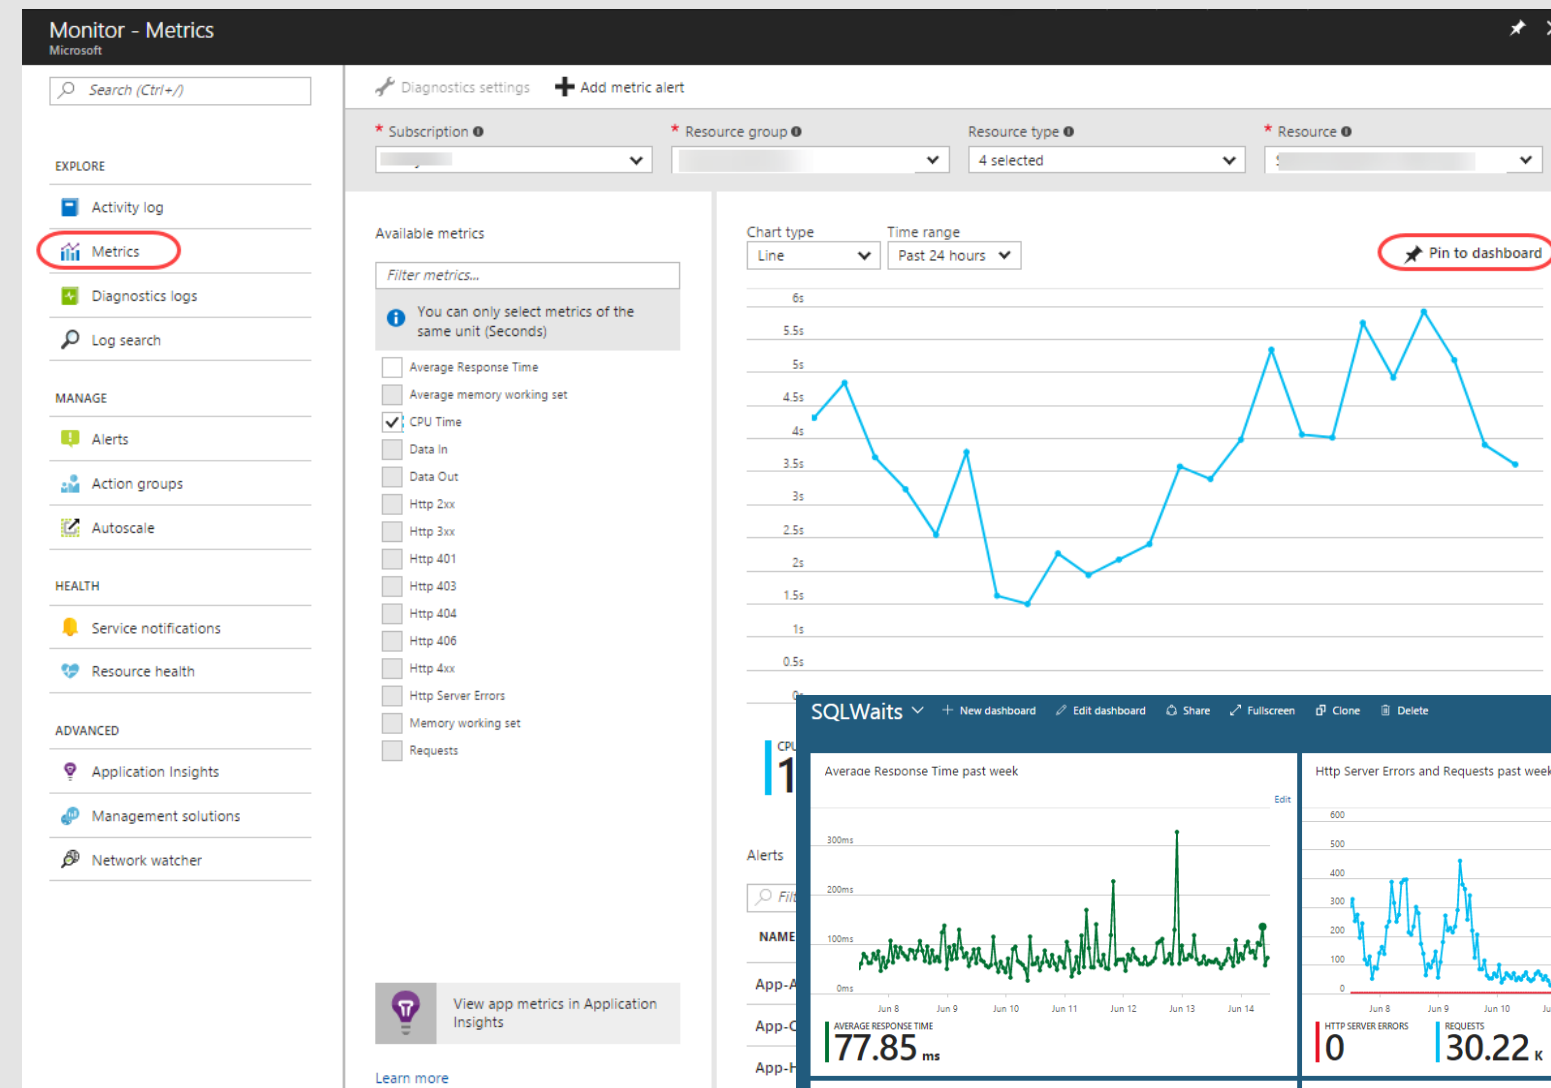
SentryOne software

We have not needed Microsoft OMS (Operations Management Suite) as of yet, but we may grow into it.

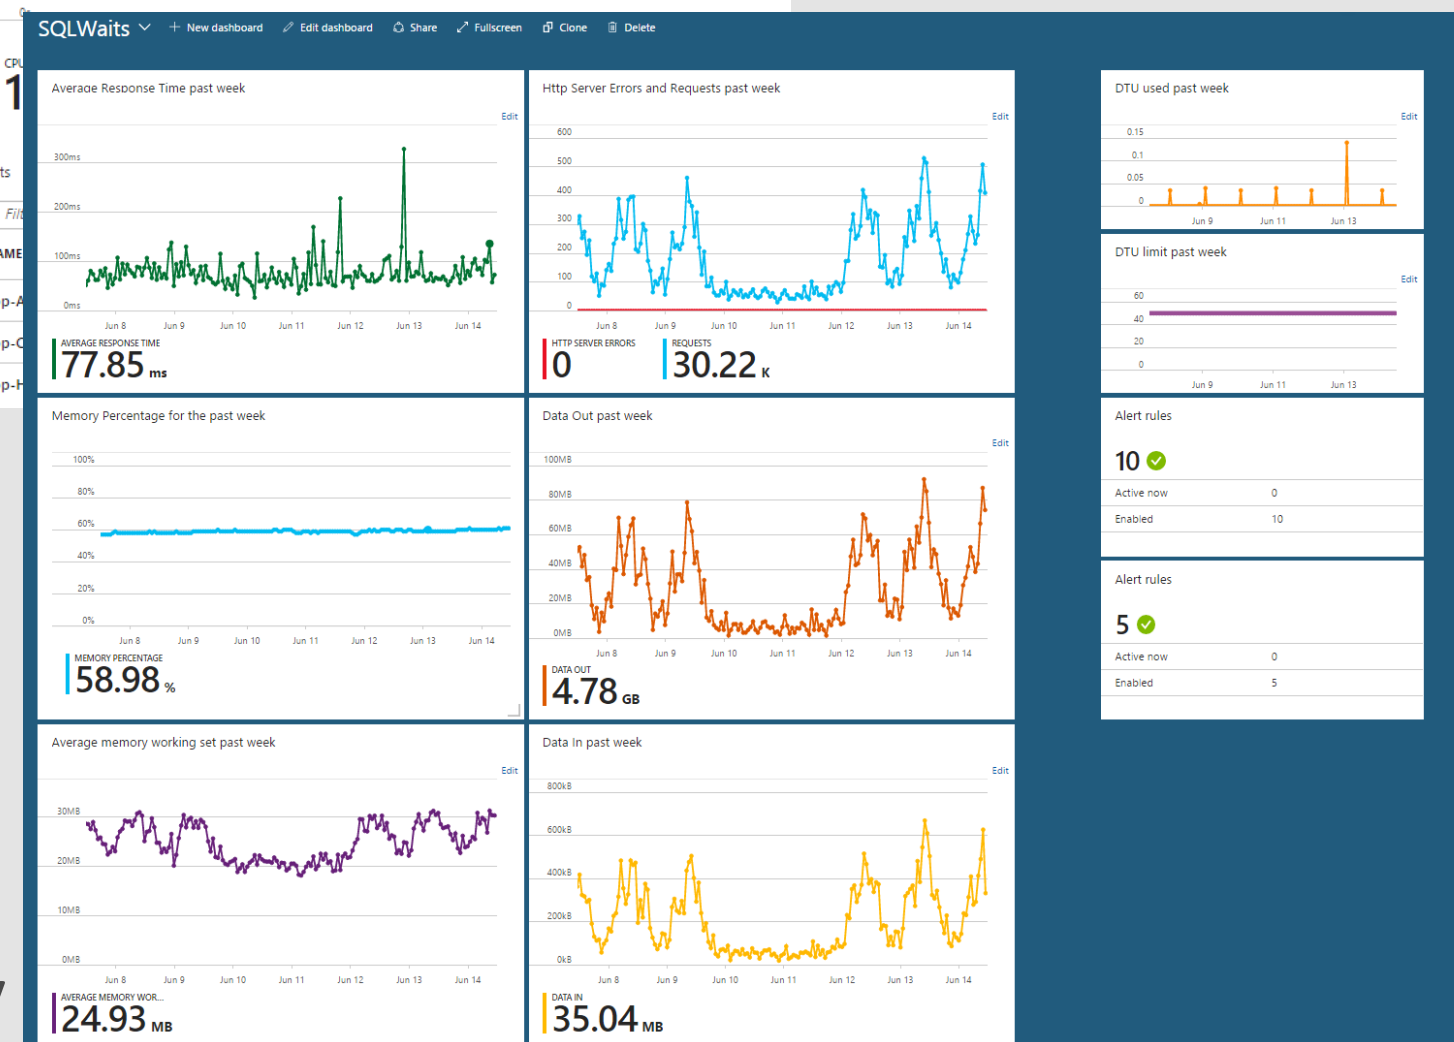
Azure Monitor

2 Activity Log

- ✓ Captures create, update & delete for resources
- ✓ Still called 'Operational Logs' when sent to blob storage



OPERATION NAME	STATUS	TIME	TIME STAMP	SUBSCRIPTION	EVENT INITIATED BY
Execute	Succeeded	17 h ago	Tue Jun 13 20...	Analytics	VS DevTest Lab
Execute	Succeeded	17 h ago	Tue Jun 13 20...	Analytics	VS DevTest Lab
Execute	Succeeded	19 h ago	Tue Jun 13 20...	Analytics	VS DevTest Lab
Deallocate	Succeeded	19 h ago	Tue Jun 13 20...	Analytics	VS DevTest Lab
Incident	Resolved	1 d ago	Tue Jun 13 20...	Analytics	@sentryone.com
ListKeys	Succeeded	1 d ago	Tue Jun 13 20...	Analytics	@sentryone.com
Start	Succeeded	1 d ago	Tue Jun 13 20...	Analytics	@sentryone.com
ListApplicableSchedules	Failed	1 d ago	Tue Jun 13 20...	Analytics	@sentryone.com
Execute	Succeeded	1 d ago	Tue Jun 13 20...	Analytics	VS DevTest Lab
Deallocate	Succeeded	1 d ago	Tue Jun 13 20...	Analytics	VS DevTest Lab
Write VirtualMachines	Succeeded	2 d ago	Mon Jun 12 2...	Analytics	@sentryone.com
Write VirtualMachines	Failed	2 d ago	Mon Jun 12 2...	Analytics	@sentryone.com
Resolved	Resolved	2 d ago	Mon Jun 12 2...	Analytics	Microsoft.Insights/alertRules
Activated	Activated	2 d ago	Mon Jun 12 2...	Analytics	Microsoft.Insights/alertRules



1 Metrics

- ✓ Performance counters in 1-minute frequency

3 Diagnostic Log

- ✓ Emitted by each resource

Azure Monitor: Alerts

Metric Alert

- Email or webhook specified within the alert
- Always applies to one single resource

Activity Log Alert

- Can use an Action Group for SMS, email, or webhook
- Is a resource itself assigned to a RG
- Can span resource groups & resources

There are currently no alerts for the diagnostics logs.

Monitor - Alerts
Microsoft

Search (Ctrl+/)

Columns **+ Add metric alert** **+ Add activity log alert**

* Subscription **Analytics** Source **All sources**

Filter alerts...

NAME	STATUS	CONDITION
App-AvgResponseTime-Prod	Active	Average Response Time > 4 Seconds
App-CPUTime-Prod	Active	CPU Time > 60 Seconds
App-HttpServerErrors-Prod	Active	Http Server Errors > 0 Count
ASP-CPUct-Prod	Active	CPU Percentage > 80 Percent
ASP-MemoryPct-Prod	Active	Memory Percentage > 85 Percent
SQLDB-BlockedByFirewall-Prod	Active	Blocked by Firewall > 0 Count
SQLDB-CPUpercentage-Prod	Active	CPU percentage > 80 Percent
SQLDB-DataIOPct-Prod	Active	Data IO percentage > 80 Percent
SQLDB-Deadlocks-Prod	Active	Deadlocks > 0 Count
SQLDB-DTULimit-Prod	Active	DTU limit > 50 Count
SQLDB-DTUused-Prod	Active	DTU used > 85 Count
SQLDB-FailedConnections-Prod	Active	Failed Connections > 0 Count
SQLDB-LogIOPct-Prod	Active	Log IO percentage > 80 Percent
SQLDB-SessionPct-Prod	Active	Sessions percentage > 80 Percent
SQLDB-WorkersPct-Prod	Active	Workers percentage > 80 Percent

Azure Monitor: Action Groups

Used with Activity Log Alerts only right now (not supported with Metric Alerts).

Action types:

- SMS
- Email
- Webhook

Monitor - Action groups
Microsoft

Search (Ctrl+/)

EXPLORE

- Activity log
- Operation log (classic)
- Metrics
- Diagnostics settings
- Log search

MANAGE

- Alerts
- Action groups**
- Autoscale

Columns + Add action group

* Subscription ⓘ
Microsoft Azure Sponsorship

NAME	SHORT NAME	RESOURCE GROUP	STATUS	ACTIONS
BI Action Group - Critical	BIAG-Crit	MonitoringAndAlertsRGProd	Enabled	1 SMS, 1 email(s), 0 webhook(s)
BI Action Group - NonCritical	BIAGNotCrit	MonitoringAndAlertsRGProd	Enabled	0 SMS, 1 email(s), 0 webhook(s)

Generally only Prod resources have any alerting configured.

Critical alerts: SMS text + e-mail
Non-critical alerts: e-mail only

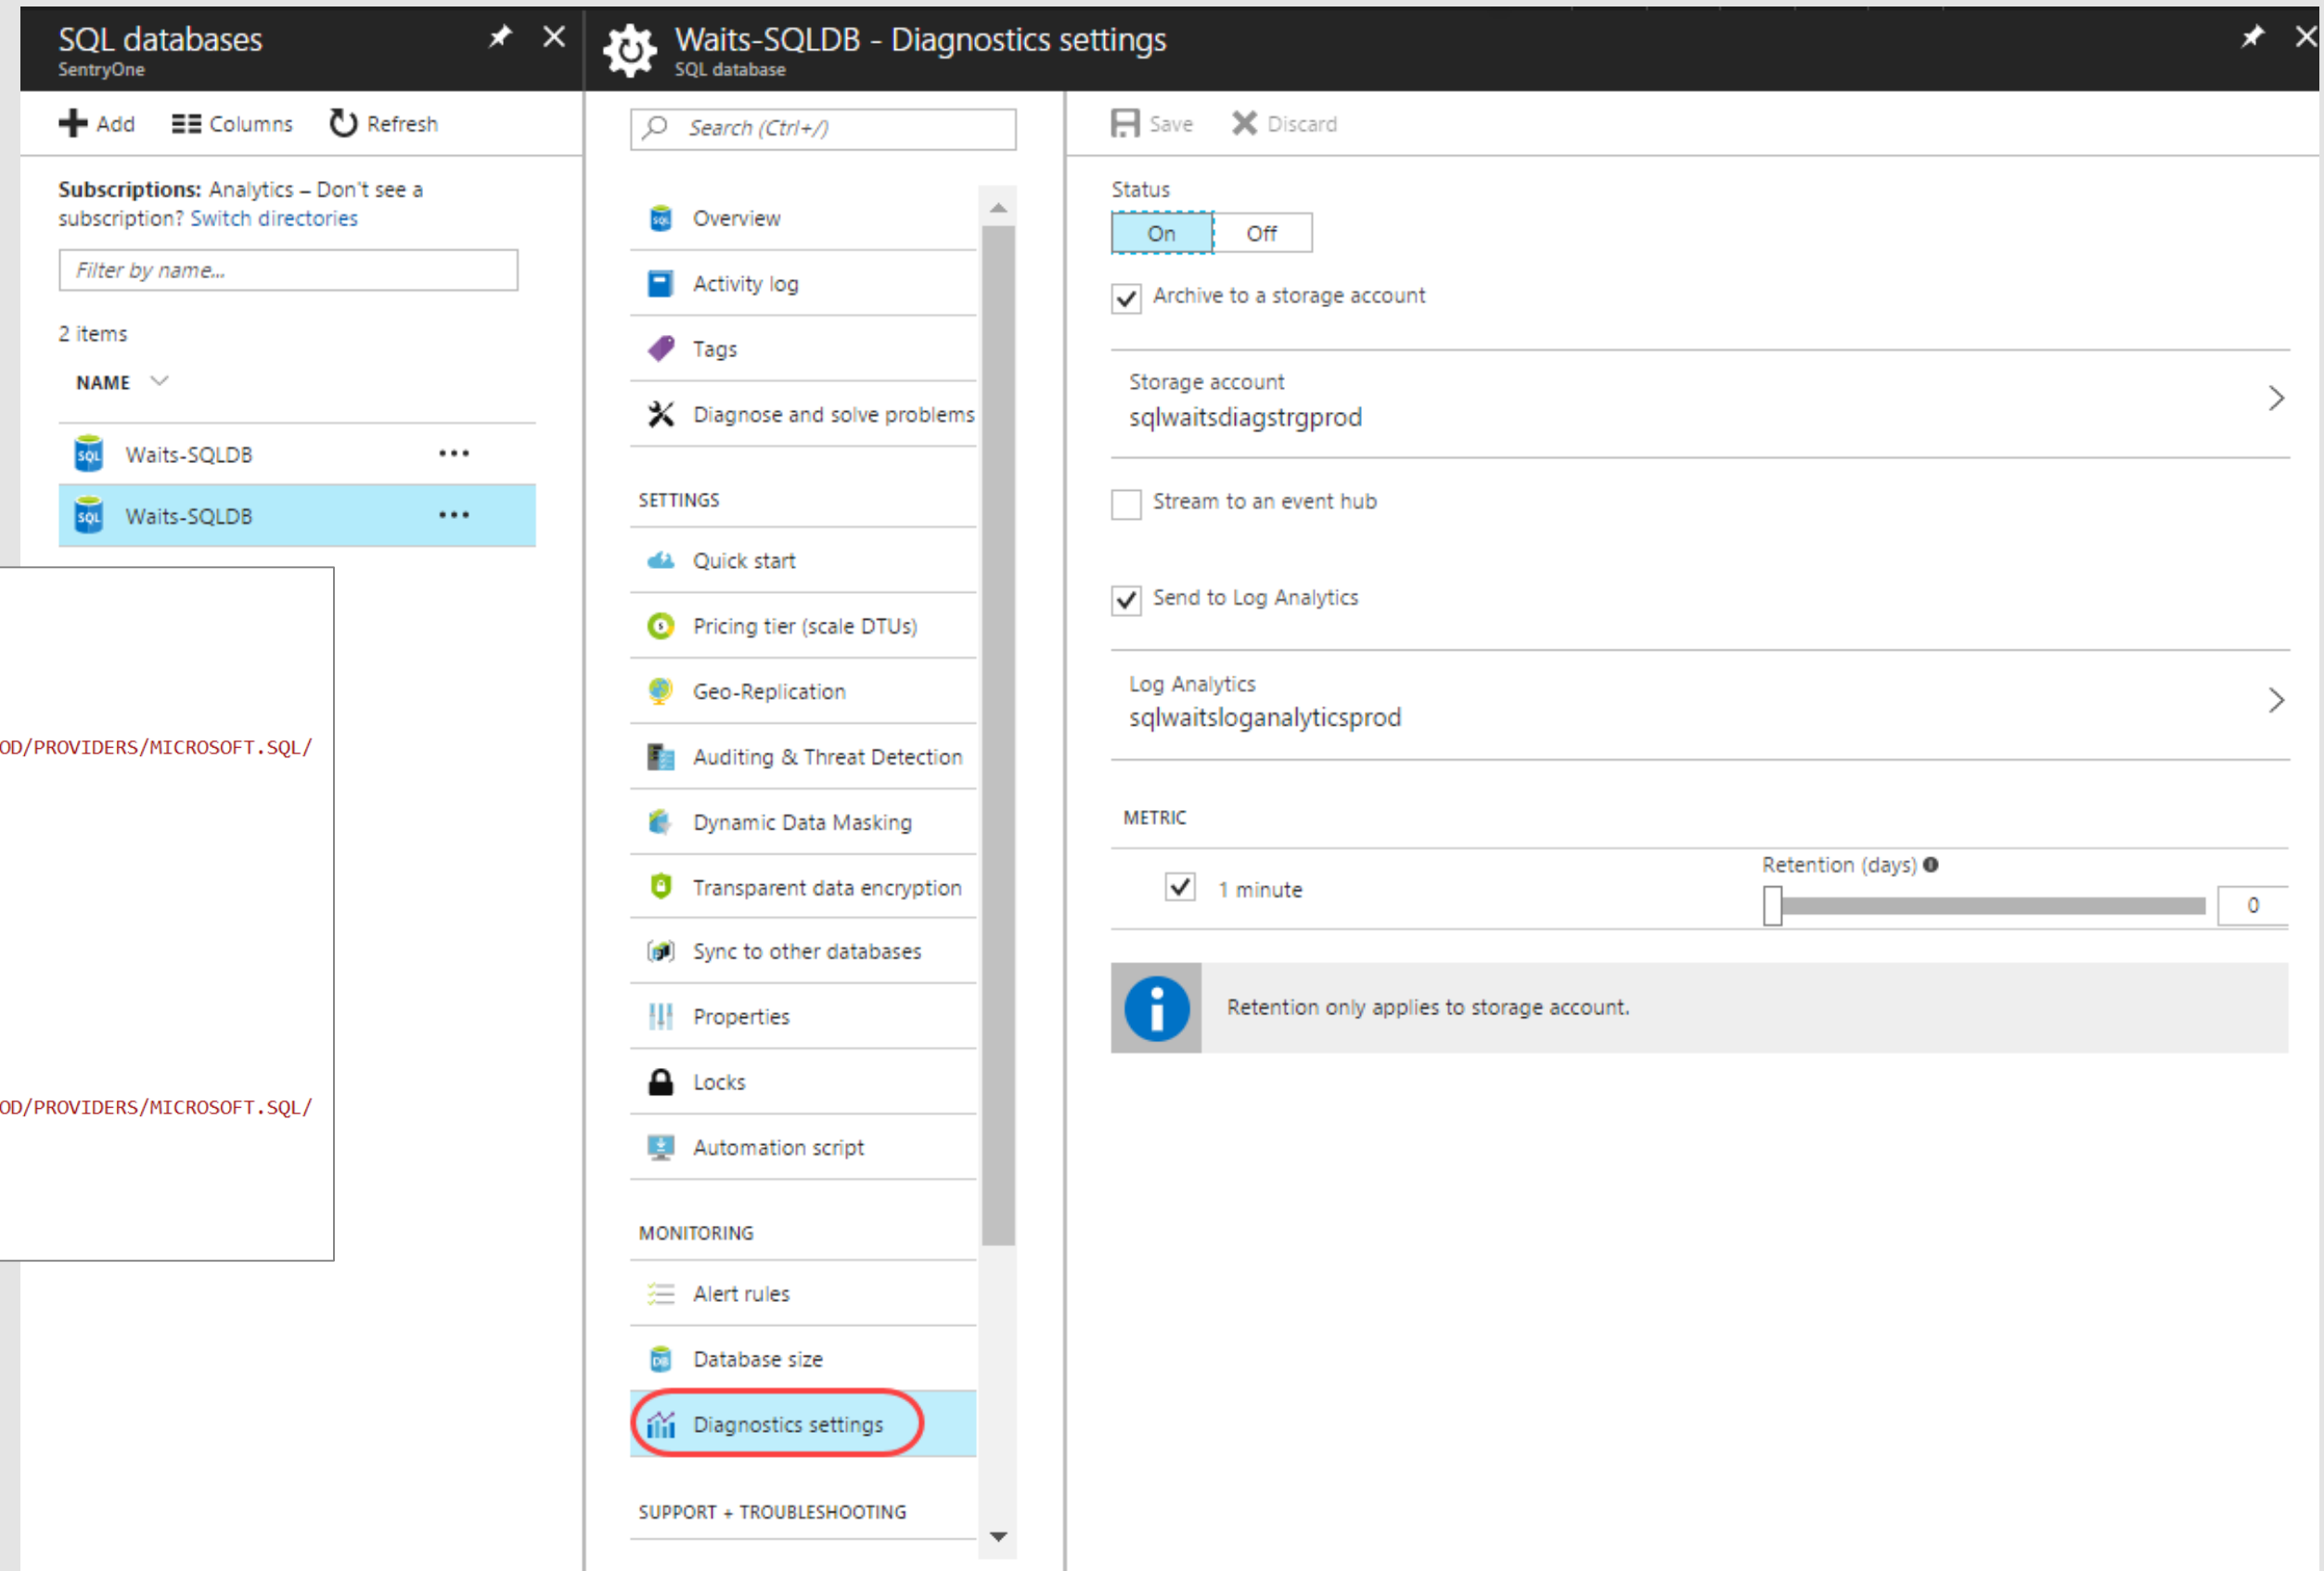
Azure Diagnostics: SQL Database (PaaS)

We are sending diagnostics to:

- Storage Account (JSON data)
- Log Analytics

```
{
  "count": 1,
  "total": 24247296,
  "minimum": 24247296,
  "maximum": 24247296,
  "average": 24247296,
  "resourceId": "/SUBSCRIPTIONS/[REDACTED]/RESOURCEGROUPS/SQLWAITSRGPROD/PROVIDERS/MICROSOFT.SQL/SERVERS/SQLWAITSSERVERPROD/DATABASES/WAITS-SQLDB",
  "time": "2017-05-10T14:17:00.0000000Z",
  "metricName": "storage",
  "timeGrain": "PT1M"
},
{
  "count": 1,
  "total": 1,
  "minimum": 1,
  "maximum": 1,
  "average": 1,
  "resourceId": "/SUBSCRIPTIONS/[REDACTED]/RESOURCEGROUPS/SQLWAITSRGPROD/PROVIDERS/MICROSOFT.SQL/SERVERS/SQLWAITSSERVERPROD/DATABASES/WAITS-SQLDB",
  "time": "2017-05-10T14:17:00.0000000Z",
  "metricName": "connection_successful",
  "timeGrain": "PT1M"
}
```

Partitioned in Storage as low as:
Resource Type > Year > Month > Day
> Hour > Minute



Azure Auditing: SQL Database (PaaS)

Auditing is set at both the Server and the Database level for SQLDB.

Send to: the 'diagnostics' blob storage account associated with this particular project.

The screenshot displays the Azure portal interface for configuring Auditing & Threat Detection on a SQL database. The left-hand navigation pane shows the 'Auditing & Threat Detection' option highlighted with a red circle. The main content area features a top navigation bar with 'Save', 'View audit logs', and 'Feedback' options. A prominent information banner states: 'Threat Detection costs \$15/server/month. It will be free for the first 60 days. If Blob Auditing or Threat Detection are enabled on the server, they will always apply to the database, regardless of the database settings.' Below this, the 'Server-level Auditing' and 'Server-level Threat Detection' are both shown as 'Enabled'. The 'Auditing' section has a toggle set to 'ON' and 'Auditing type' set to 'Blob'. The 'Threat Detection' section has a toggle set to 'ON' and 'Threat Detection types' set to 'All'. Under 'Storage details', the storage account 'sqlwaitstdiagstdstrgdev' is listed. The 'Send alerts to' field contains the email address 'AzureSupport@sentryone.com' with a green checkmark, and there is an unchecked checkbox for 'Email service and co-administrators'.

Azure Diagnostics: Virtual Machine (IaaS)

Virtual machines and Virtual machines (classic) can now be managed together in the combined list below.

Subscriptions: 1 of 2 selected – Don't see a subscription? [Switch directories](#)

Filter by name...

2 items

NAME
BISQLVM1Dev
BISQLVM2Dev

Search (Ctrl+/)

Save Discard

Overview Performance counters Logs Crash dumps Sinks Agent Boot diagnostics

Performance counters

Collecting data for these counters:

- CPU
- Memory
- Disk
- Network
- SQL Server

[Configure performance counters](#)

Event logs

Collecting data for these logs:

- Application: Critical, Error, Warning
- Security: Audit failure
- System: Critical, Error, Warning

[Configure event logs](#)

Directories

Not configured.

[Configure directories](#)

Crash dumps

Not collecting memory dumps.

[Configure crash dumps](#)

Sinks

Diagnostic data is not being sent to any sinks.

[Configure sinks](#)

Agent

Diagnostic data is being sent to this storage account:

[bisqlvm1diagstdstrgdev](#)

[Configure agent](#)

Boot diagnostics

Boot diagnostics is enabled.

[View boot diagnostics](#)

[Configure boot diagnostics](#)

Performance Counters

SQL Server counters are not gathered by default – need to explicitly select them.

Agent

The Azure Diagnostics Agent doesn't prompt for storage account – will need to change (i.e., if you have a naming convention for diagnostics storage), and then delete the auto-created storage account.

Azure Security Center

Security Center - Security policy

Search (Ctrl+)

GENERAL

- Overview
- Security policy**
- Quickstart
- Welcome

PREVENTION

NAME	INHERITANCE	DATA COLLECTION
[Redacted]	---	On

Security policy

Save Discard

Data collection

Collect data from virtual machines

On Off

Choose a storage account per regi...
All configured

Policy components

- Prevention policy
- Email notifications
- Pricing tier

Prevention policy

Show recommendations for

- System updates On Off
- OS vulnerabilities On Off
- Endpoint protection On Off
- Disk encryption On Off
- Network security groups On Off
- Web application firewall On Off
- Next generation firewall On Off
- Vulnerability Assessment On Off
- Storage Encryption On Off
- SQL auditing & Threat detection On Off
- SQL Encryption On Off

Security Center - Recommendations

Search (Ctrl+)

GENERAL

- Overview
- Security policy
- Quickstart
- Welcome

PREVENTION

- Recommendations**
- Partner solutions
- Compute
- Networking

Filter

DESCRIPTION	RESOURCE	STATE	SEVERITY
Enable advanced security for subscriptions	1 subscriptions	Open	High
Enable data collection for subscriptions	1 subscriptions	Resolved	High
Add a Next Generation Firewall	BISQLVM1IPDev	Open	High
Enable Network Security Groups on subnets	default	Open	High
Enable Auditing & Threat detection on SQL servers	sqlserversandbox	Open	High
Enable Auditing & Threat detection on SQL databases	2 SQL databases	Open	High
Enable encryption for Azure Storage Account	2 storage accounts	Open	High
Restrict access through Internet facing endpoint	BISQLVM1Dev	Open	Medium
Add a vulnerability assessment solution	BISQLVM1Dev	Open	Medium
Enable Transparent Data Encryption	HappyCars	Open	Medium
Provide security contact details	1 subscriptions	Open	Medium

More info:

<https://docs.microsoft.com/en-us/azure/security-center/security-center-recommendations>

SentryOne Monitoring



Support for monitoring:

- ✓ SQL Server
(on-prem, or in Azure or a cloud VM)
- ✓ Azure SQL Database
- ✓ Azure SQL Data Warehouse
- ✓ SQL Server Analysis Services
(on-prem, or in Azure or a cloud VM)
- ✓ Windows Computer
- ✓ Tintri VMstore
- ✓ VMware Host

Key Takeaways and Q&A

Key Takeaways

The **pace of change** in Azure is frequent—plan to keep up with announcements. New features can change a previous decision really fast.

Look at **using PaaS** unless you really need an IaaS solution. Deploy the simplest solution that gets the job done.

Don't shortcut the initial planning phase – do a **POC** if you can to learn. Don't assume a feature is definitely available or works as you expect it to.


For unfamiliar features or services, consider creating a **sandbox** area to **test & learn** on first.

Plan to invest some time learning **ARM templates & PowerShell**, if you haven't already.

Pay attention to **efficiency and opportunities for cost savings**. Developers & Administrators have a huge impact on cost.

Thank You!

To download a copy of this presentation:
SQLChick.com “Presentations & Downloads” page



Melissa Coates

Lead Analytics Architect,
SentryOne

Blog: sqlchick.com

Twitter: [@sqlchick](https://twitter.com/sqlchick)

*Creative Commons License:
Attribution-NonCommercial-NoDerivative Works 3.0*

