Qlik Sense on Azure

Deployment Guide

Contents

Abstract

Microsoft Azure offers a wide range of integrated cloud based services (i.e. analytics, computing, database, mobile, networking, storage, etc.). Qlik Sense is a platform that delivers visual analytics right to where they deliver the most value – the point at which decisions are made. Qlik's platform is widely deployed in Azure Cloud infrastructure to support visual and analytical analysis at various levels within Organizations. Running Qlik Sense solutions on Azure Virtual Machines is a great option for Organizations that want to deliver SaaS (software as a service) solutions to internal and/or external users delivering high-performance operations on any number of data sources. SaaS solutions based on Qlik's platform and delivered through Azure Cloud infrastructure allow individuals to easily create and share applications with a wide number of users across multiple type of devices.

This Deployment Guide provides an overview on how to deploy and implement Qlik Sense platform on the Azure Cloud platform including architectural considerations and configuration steps that explain how Qlik Sense can leverage various Azure services such as Azure SQL Database, Azure SQL Data Warehouse, Azure HD Insight, Azure Storage, Azure Cloud Services, Azure Active Directory and Azure Virtual Network to run scalable and reliable analytical solutions.

Introduction

Azure provides a cloud infrastructure with all the services and computing power needed to provide global solutions that can reliably perform as needed, regardless of unexpected changes in demand and concurrency across various levels of any Organization. Qlik® Sense is a next-generation visual analytics platform that empowers everyone to see the whole story that lives within their data. It supports a full spectrum of BI use cases including self-service visualization for everyone, centrally deployed guided analytics apps and dashboards, embedded analytics, and reporting, all within a governed framework that drives enterprise scalability and trust for IT giving everyone the ability to answer not just 'what happened', but 'why' and 'what it is likely to happen next'. Qlik is the first and only analytics software vendor that can do this by providing a fully integrated, simple and intuitive visual analytics platform, powered by Qlik's patented in-memory associative data indexing engine that delivers analytics and therefore insight right to the point of decision. Qlik delivers a complete analytics solution:

- Simplified visual data preparation
- Associative engine for business discovery
- Cloud-based upstream data sourcing and delivery
- Self-service & governed data visualization
- Centralized guided analytics
- Collaborative analytics and reporting
- Embedded analytics in any application or web page

Qlik: A Primer

Qlik is a Business Intelligence platform which provides “Managed Analysis” and Data Visualization capabilities allowing people to navigate and filter freely on any of the data elements in a rich multi-sheet dashboard. All visible objects are recalculated after a user performs selections (highlighted in green) causing associated data to be clearly distinguishable (highlighted in white) from non-associated data (highlighted in grey). This is usually referred as ‘the associative experience’ and it is the foundation for the data discovery capabilities within Qlik. Basically, the ability to understand the data and make new discoveries in a visual way. Qlik is able to deliver visual analytics in a variety of ways, each a category in its own right.

- **Self-Service Data Visualization** empowers all users to create and explore visual analytics in a self-service manner, whether it’s a local spreadsheet or large datasets from multiple systems-of-record. For some, this could mean starting from scratch to look at ad-hoc combinations of data to investigate a hunch, or confirm a theory. For others, it simply means re-using pre-constructed and re-usable analytics applications, with the ability to modify to suit their individual needs. The advantage with Qlik is that Qlik solutions don’t trade off agility for trust and scalability. Users have easy-to-use, drag-and-drop self service capabilities, while still operating in a framework of security and trusted data models that lead to trusted insights.

- **Guided Analytics** provides data driven applications and dashboards -typically developed by IT or technical business analysts-, for consumption by others within their day-to-day operational processes. An example of guided analytics app would be an interactive dashboard where users can search and filter information in a structured manner, and use pre-defined navigation options to move to different views to support better decision making with less complexity for the end user.

- **Embedded Analytics** provides the ability to insert analytics and visualizations into any workflow or application. This furthers the goal of delivering analytics and insight exactly where and when it is needed, to improve understanding and decision making throughout the organization.

- **Custom analytics application** provides the ability to create custom mashups that leverage out-of-the-box as well as custom design visualizations for various types of analytics. These custom applications and mashups can be easily deployed as standalone solutions and/or as part of a larger application while still providing the same associative experience as a traditional Qlik application.

- **Collaborative Analytics (Reporting)** deliver information and insights to people in a repeatable, portable format even when they are outside the BI system. Reporting is one of the oldest forms of BI and the reports are typically static, with limited or no interactivity, and are often printed for use off-line (i.e. financial reports). Today, reporting provides limited value, but when combined with visual analytics linked to the actual underlying data it drives increased efficiencies, higher collaboration and delivers the power of visual analytics to all of the business stakeholders. The ability to connect reporting with dynamic interactive visualization gives your data a story to tell, rather than being locked in a static report.

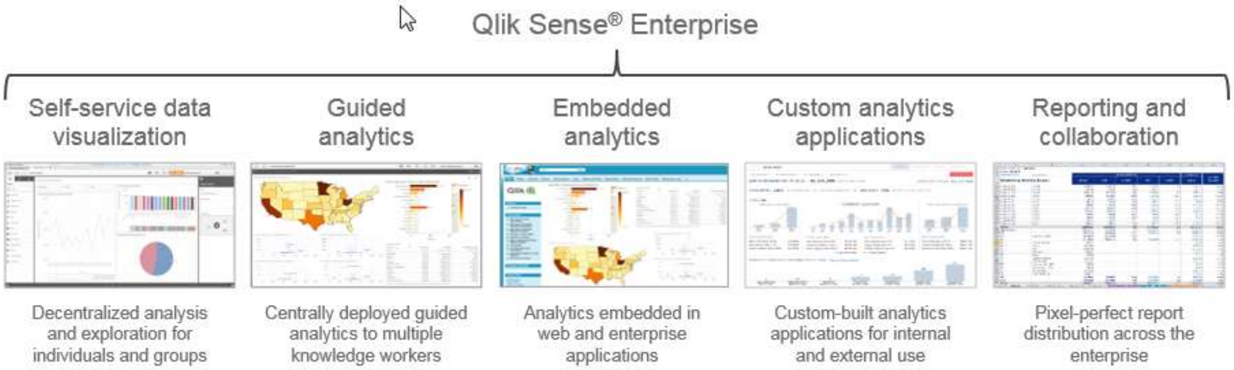
Qlik Sense® Enterprise
Self-service data visualization
Guided analytics
Embedded analytics
Custom analytics applications
Reporting and collaboration
Decentralized analysis and exploration for individuals and groups
Centrally deployed guided analytics to multiple knowledge workers
Analytics embedded in web and enterprise applications
Custom-built analytics applications for internal and external use
Pixel-perfect report distribution across the enterprise

Qlik Sense Server is a suite of component services that are installed as part of the Server. While such components can be installed across various nodes, this document covers the scenario where all of those components are installed within the same node. In other words, a single node architecture. The image below depicts some of the main component services that are part of the Qlik Sense Service.

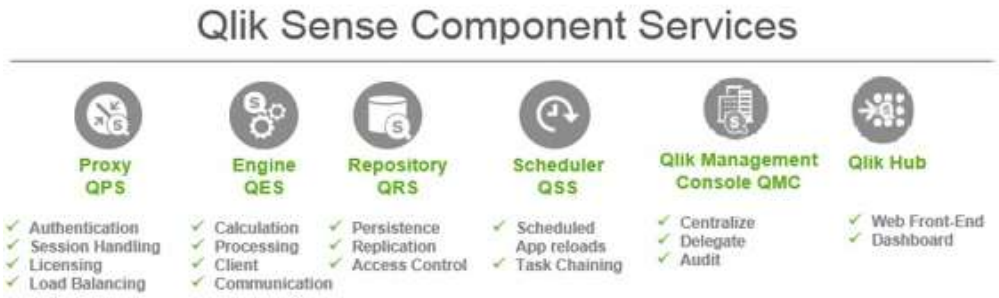

In order to learn more about such component services and how to architect a Qlik Sense environment, visit the following links:

http://help.qlik.com/en-US/sense/3.0/Subsystems/PlanningQlikSenseDeployments/Content/Server/Server-Architecture.htm

http://help.qlik.com/en-US/sense/3.0/Subsystems/PlanningQlikSenseDeployments/Content/Server/Server-Architecture-Services.htm

Qlik on Azure – High Level Architecture

Cloud based solutions are prevalent in almost every single Industry and Qlik's customers are eager to implement Qlik Sense environments on Azure infrastructure so they can rapidly deploy new applications in an simple, cost-effective and scalable manner. The difficult and lengthy procurement process for on-premises equipment makes it hard to meet deployment timelines, but customers can spin up new Azure infrastructure on demand in minutes. For those customers, the Azure based approach is ideal for a SaaS solutions that need to integrate with a number of different other back and front office solutions. This guide will help customers pursuing this approach so that they understand the different aspects of Azure and Qlik that will be involved in the setup and configuration of the environment.

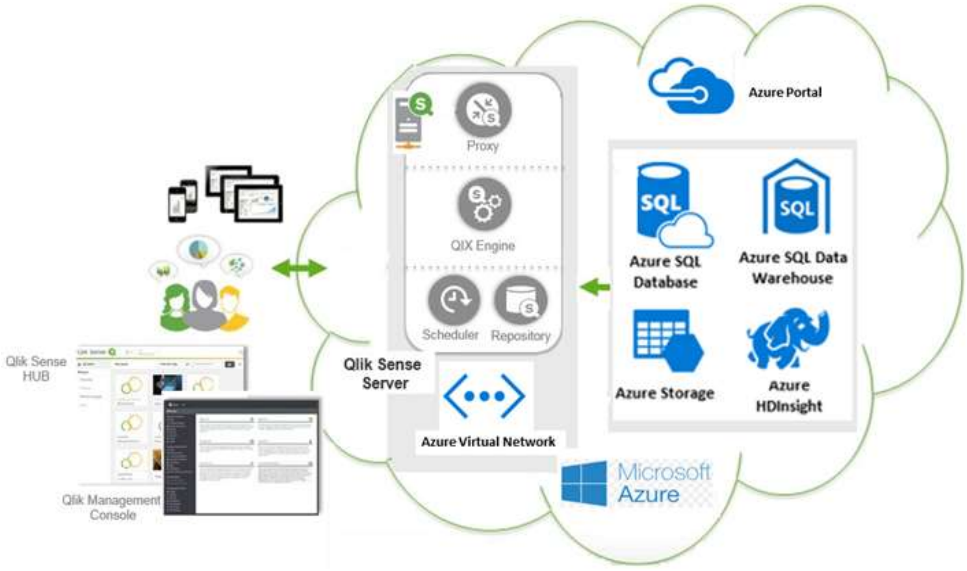

Some of the characteristics that tend to be similar across Organizations that go with Qlik and Azure implementations are:

I. **Pre-requisites**:

- Tight timeline to develop and deploy Qlik Sense applications
- Appropriate Licensing from Qlik is available
- Enterprise Architecture guidance from IT (i.e. integration with other back and front office applications, etc.)

II. **Competencies**

- Good understanding of Azure architecture and services (i.e. etc.) within the team
- Internal resources that can maintain and manage the Qlik and Azure environments are available
- Ability to setup multiple environments (i.e. Development, Testing, Disaster Recovery, etc.) within Azure and appropriate time for testing the infrastructure is available

III. **Constraints**

- Volume and frequency of data updates will not interfere with the cloud-based approach
- (i.e. clear strategy on data procurement and nightly incremental feeds)

Azure is an excellent platform for running many advanced data systems in the cloud. Characteristics of cloud computing solutions built on Azure, such as scalability, security and tight synergy with Business Users are key benefits when architecting and running Qlik applications. A general understanding of these characteristics can help you make the best architecture decisions for your analytical system. Understanding of this guide will be easier if you are already familiar with the following topics:

I. Azure Cloud Services and Azure Virtual Network

II. Qlik Sense Server

III. Qlik Sense Configuration

IV. Windows Server 2012

V. Windows Server Active Directory and DNS

VI. Azure SQL Data Warehouse

While this document covers general steps for Qlik Sense product installations and software configuration tasks, careful consideration on configuration topics is strongly recommended when you are planning and deploying Qlik Sense Solutions on Azure Cloud in Production environments. For more advanced scenarios, consult with Qlik Professional Services and/or one of Qlik's Solution Integrator partners. For generic Qlik Sense software configuration guidance and best practices, consult the Qlik product documentation available at http://help.qlik.com. If you are new to Azure, see the 'Get Started' section of the Azure documentation (which provides high level references and links to specifics areas of documentation) at: https://azure.microsoft.com/en-us/get-started/

Installation

In order to install Qlik Sense on Azure, you can download Qlik Sense Enterprise from http://www.qlik.com/download (a Qlik account will be needed in order to access the download site) and install it in an Azure Virtual Machine (to learn about Azure Virtual Machines, visit: https://azure.microsoft.com/en-us/services/virtual-machines/). If an Azure Virtual Machine doesn't exist yet, it first has to be created. To do so, there are couple of ways on how to accomplish this. The first option is to leverage an ARM (Azure Resource Manager) template. The second one would be to manually create the Azure Virtual Machine.
In order to leverage an ARM template to create the Azure Virtual Machine, the ARM template has to be created first. The ARM template can be created manually or by leveraging an existing example. Within the Azure development community, there are a number of open source samples that can be used for such purpose. For example, the following URL link under Github provides a free open source example of an ARM template that could be used to create a Windows Azure Virtual Machine that then it could be used to install Qlik Sense Enterprise and create a Qlik Sense environment running on Azure Cloud infrastructure. https://github.com/krist00fer/qlik

The page will give you access to the source code (JSON files) as well as information on how to deploy the ARM template by the Azure command line (CLI) and/or Powershell script and some generic information too. The page also provides a button ("Deploy to Azure") that would automatically import the ARM template into the Azure Portal (credentials to the Azure Portal are needed in order to import the ARM template):

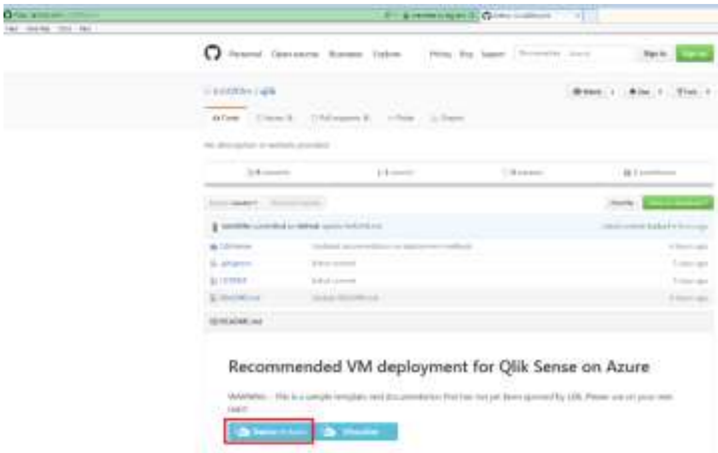

Recommended VM deployment for Qlik Sense on Azure

The page also provides the option to visualize the overall content of the ARM template, by clicking on the "Visualize" button. This functionality is particularly helpful when trying to get a high level overview on the content of the package.

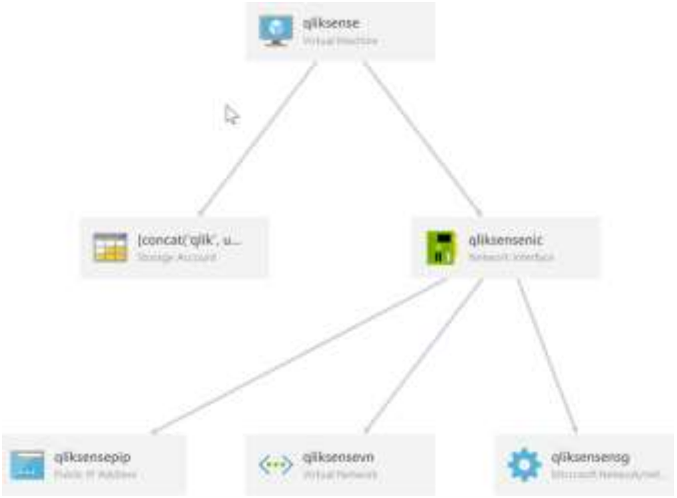

Once the ARM template has been imported into the Azure Portal, the template will prompt users for a number of different parameters and settings. The parameters that will show up will be determined by the configuration of the ARM templates. Please notice that ARM templates are very flexible and changes can be made rather easily by simply editing the JSON files. For example, using this free sample, the users would have to enter some parameters such as the name of the virtual machine, select the size of the virtual machine (this particular example has a default a Standard DS3 v2 machine type), a user name, a password, a virtual network, etc. Also, as part of the configuration of the Virtual Machine, the user can also select to either use an existing or create a new Resource Group for the Virtual Machine. All the Legal Terms associated with the template can be displayed from the Azure Portal as well. In order to then have the Virtual Machine created, just click the "Create" button at the bottom of the screen.

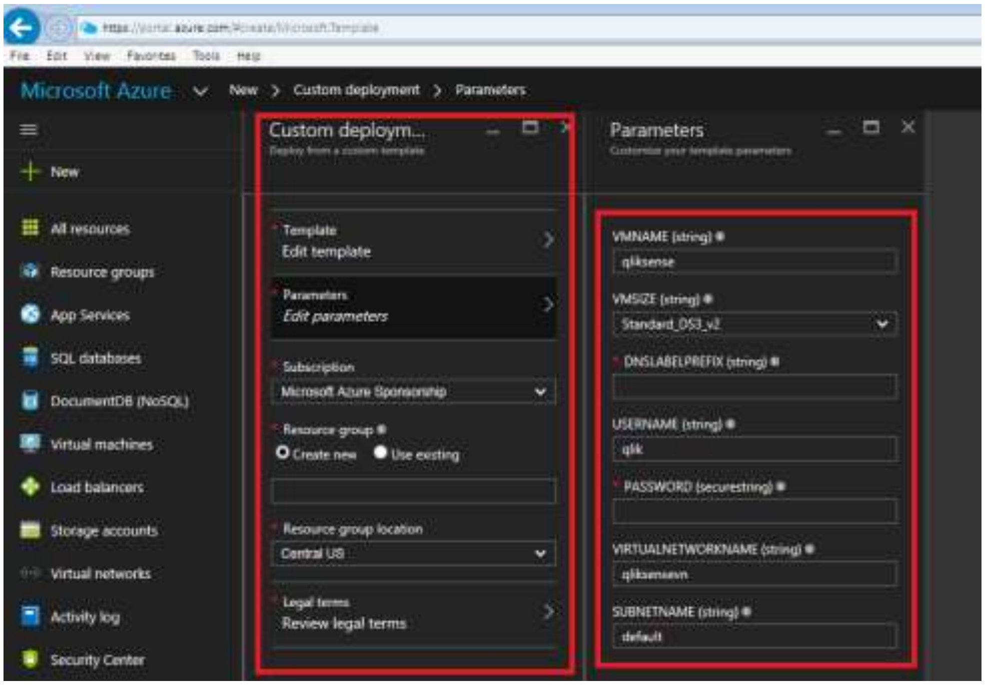

Template
Edit template

Parameters
Edit parameters

Subscription
Microsoft Azure Sponsorship

Resource group
○ Create new ● Use existing

Resource group location
Central US

Legal terms
Review legal terms

VMNAME (string)
qliksense

VMSIZE (string)
Standard_DS3_v2

DNSLABELPREFIX (string)

USERNAME (string)
qlik

PASSWORD (securestring)

VIRTUALNETWORKNAME (string)
qliksensevn

SUBNETNAME (string)
default

Another great source for content is the following URL within Github. This site contains a lot of different ARM template samples that can be leveraged accordingly. The ARM templates listed on this site can be accessed directly from the Azure Portal. Currently Qlik doesn't have any examples posted under this URL but some examples will be added in the near future: https://github.com/azure/azure-quickstart-templates

Once the Virtual Machine is available, Qlik Sense can be installed on it. The list of steps on how to do so can be found below (starting on step 35). As mentioned at the beginning of this section, the second way to create a Virtual Machine would be manually. To do so, you can follow the steps below. Please, be aware that the look and feel of the Azure Portal may vary for each customer depending on their preferences.

1. Create a Resource Group by going to https://portal.azure.com. Please, notice that in order to proceed any further you will need to have an Azure account available in order to log into Azure. On the left hand side Navigation Portal, click on Resource Group:

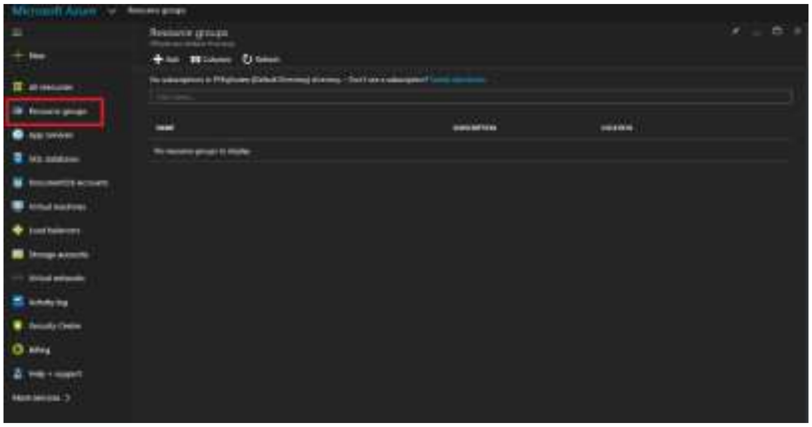

2. Click on the Add icon in order to add a new Resource Group

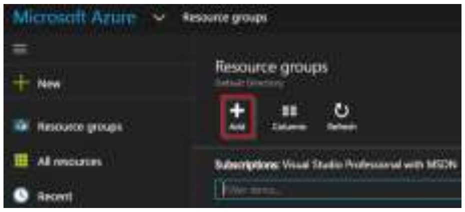

Subscriptions: Visual Studio Professional with MSDN

3. Make sure to choose a resource group name, subscription, and location that makes sense. Given that this resource group can contain many resources (e.g. multiple Qlik Sense Virtual Machines) select a name that is self-descriptive such as “QlikSense”. Notice that you can select the Resource Group location. Typically the location is the same as the location of the Virtual Machine.

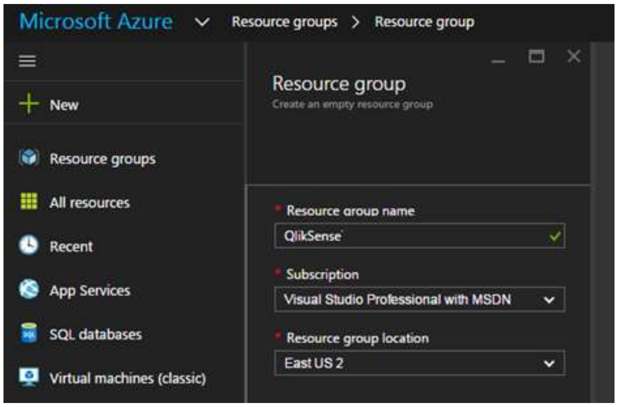

* Resource group name: QlikSense

* Subscription: Visual Studio Professional with MSDN

* Resource group location: East US 2

4. Click the “Create” button

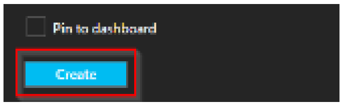

5. Create a Network Security Group by selecting “Browse” on the left navigation panel and searching for “Network Security Group” and selecting it

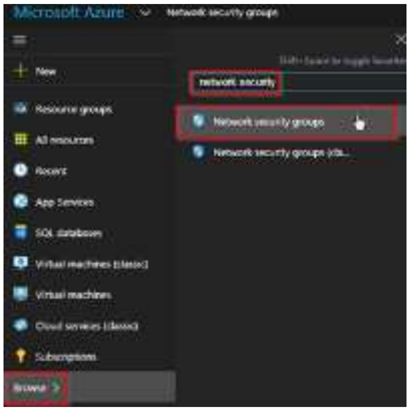

6. Click on the “Add” button

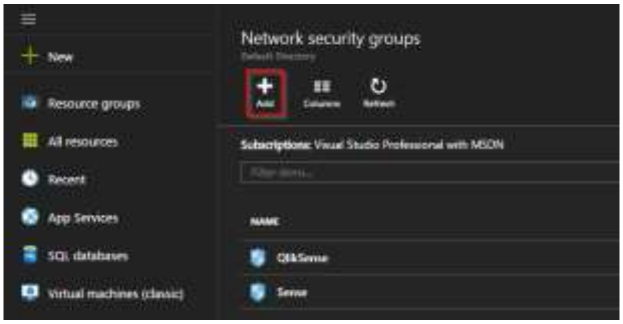

Subscriptions: Visual Studio Professional with MSDN

NAME
QlikSense
Sense

7. Select a name for the Network Security Group (i.e. “QlikSenseSecurity”) that is self-descriptive as this name will show up the same way in all the Virtual Machines that are in the “QlikSense” resource group. Select the corresponding subscription and assign it to the Resource Group that was previously created (i.e. “QlikSense”). Make sure to select the same location as the Resource Group. Click “Create”

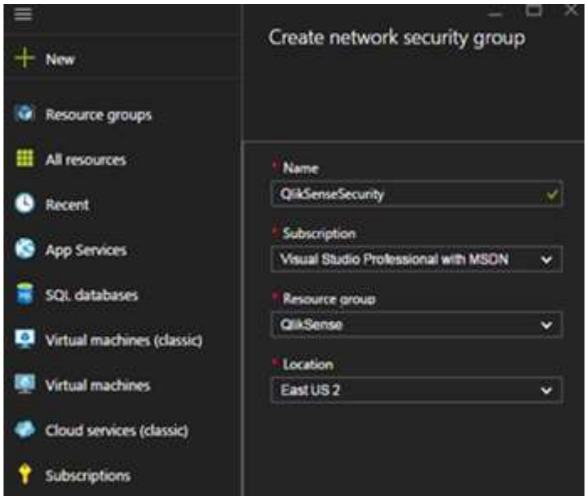

Create network security group

* Name
QlikSenseSecurity

* Subscription
Visual Studio Professional with MSDN

* Resource group
QlikSense

* Location
East US 2

8. Select the new Network Security Group. If the new group does not show up immediately, then refresh the browser.

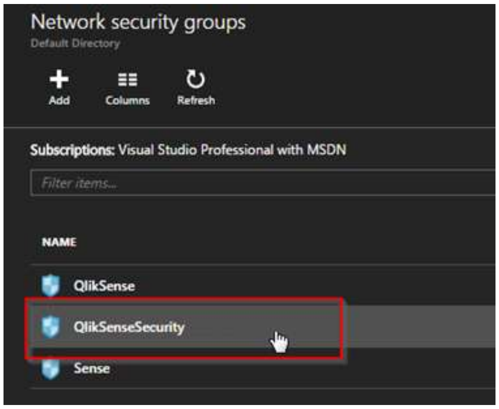

Network security groups
Default Directory

Subscriptions: Visual Studio Professional with MSDN

Filter items...

NAME
QlikSense
QlikSenseSecurity
Sense

9. All the settings associated with the Network Security Group will show up on the screen now. In the right hand side panel called "Settings" select the "Inbound Security rules"

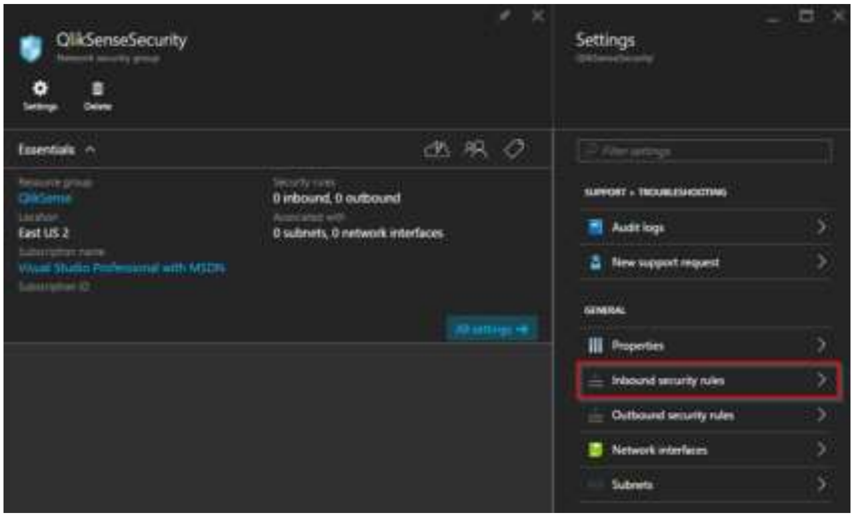

Resource group
QlikSense
Location
East US 2
Subscription name
Visual Studio Professional with MSDN
Subscription ID

Security rules
0 inbound, 0 outbound
Associated with
0 subnets, 0 network interfaces

10. Click on the "Add" button to add new port rules

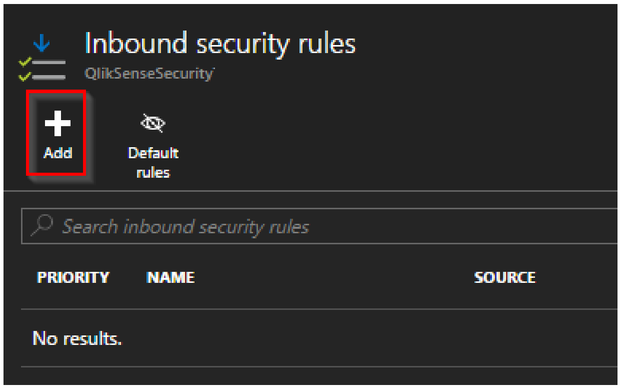

PRIORITY	NAME	SOURCE
No results.		

11. Add ports 80 (HTTP), 443 (HTTPS), 4244, 4248, and 3389. The rules should look like the screenshots below. Please, notice that the “Source” range needs to be “*”. For more information, visit the following site:

https://azure.microsoft.com/en-us/documentation/articles/virtual-networks-nsg/

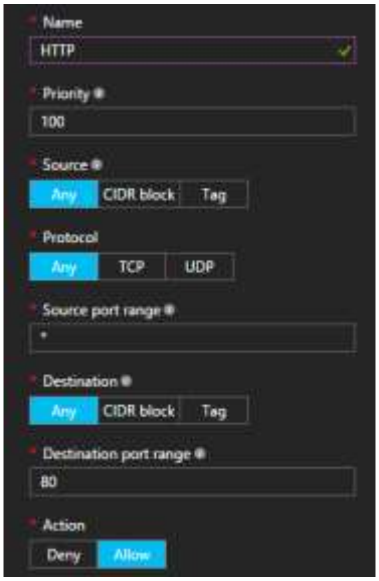

* Name: HTTP ✓

* Priority: 100

* Source: [✓ Any] CIDR block Tag

* Protocol: [✓ Any] TCP UDP

* Source port range: *

* Destination: [✓ Any] CIDR block Tag

* Destination port range: 80

* Action: Deny [✓ Allow]

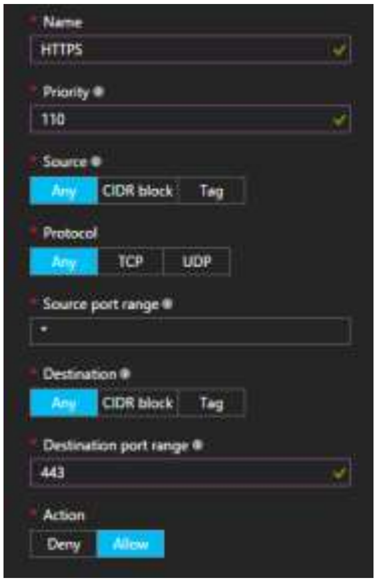

* Name: HTTPS ✓

* Priority: 110 ✓

* Source: [✓ Any] CIDR block Tag

* Protocol: [✓ Any] TCP UDP

* Source port range: *

* Destination: [✓ Any] CIDR block Tag

* Destination port range: 443 ✓

* Action: Deny [✓ Allow]

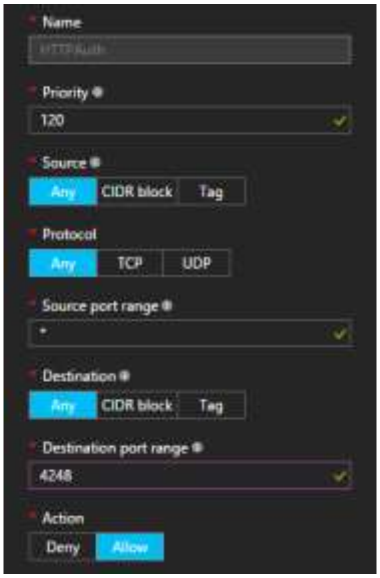

* Name: HTTPAuth

* Priority: 120 ✓

* Source: [✓ Any] CIDR block Tag

* Protocol: [✓ Any] TCP UDP

* Source port range: * ✓

* Destination: [✓ Any] CIDR block Tag

* Destination port range: 4248 ✓

* Action: Deny [✓ Allow]

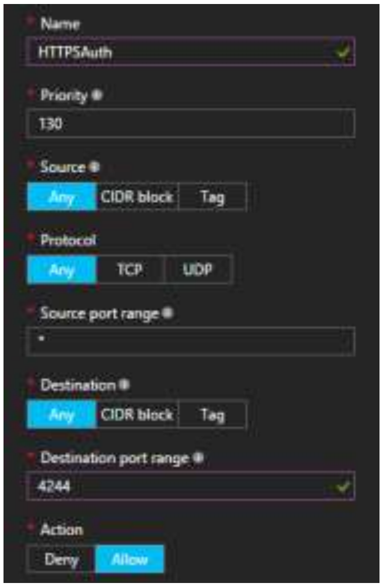

* Name: HTTPSAuth ✓

* Priority: 130

* Source: [✓ Any] CIDR block Tag

* Protocol: [✓ Any] TCP UDP

* Source port range: *

* Destination: [✓ Any] CIDR block Tag

* Destination port range: 4244 ✓

* Action: Deny [✓ Allow]

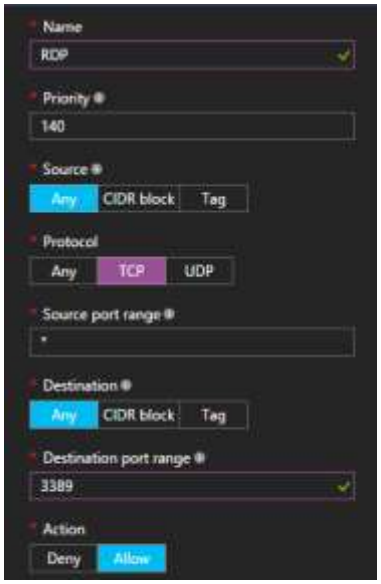

* Name: RDP ✓

* Priority: 140

* Source: [✓ Any] CIDR block Tag

* Protocol: Any [✓ TCP] UDP

* Source port range: *

* Destination: [✓ Any] CIDR block Tag

* Destination port range: 3389 ✓

* Action: Deny [✓ Allow]

12. Confirm that all the inbound rules have been added correctly

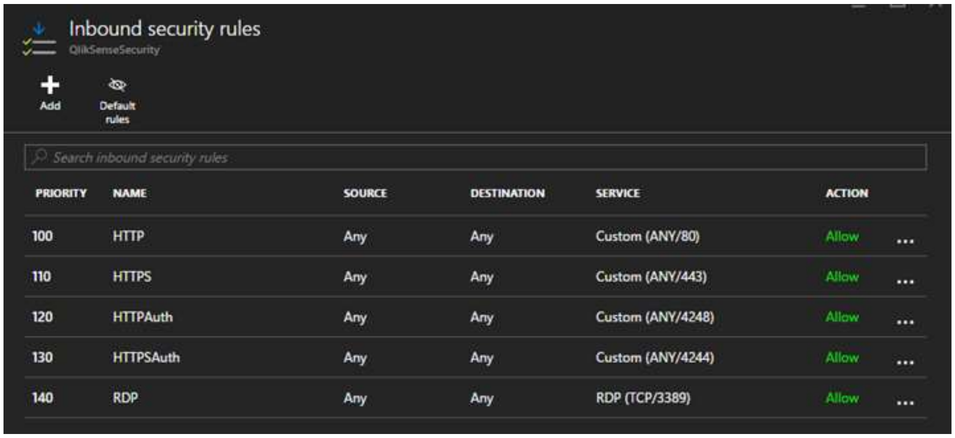

PRIORITY	NAME	SOURCE	DESTINATION	SERVICE	ACTION
100	HTTP	Any	Any	Custom (ANY/80)	Allow
110	HTTPS	Any	Any	Custom (ANY/443)	Allow
120	HTTPAuth	Any	Any	Custom (ANY/4248)	Allow
130	HTTPSAuth	Any	Any	Custom (ANY/4244)	Allow
140	RDP	Any	Any	RDP (TCP/3389)	Allow

13. Create a Storage account by browsing “Storage” on the left hand side panel. Then click on “Storage Accounts”

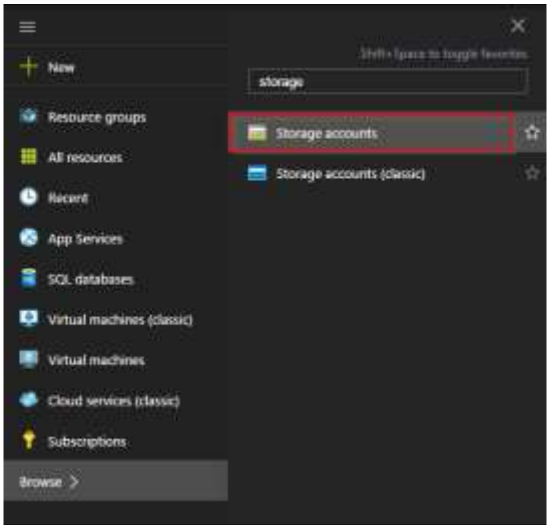

14. Click on the “Add” button

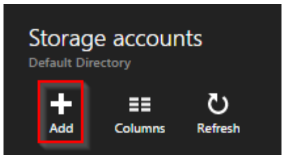

15. Enter a unique and self-descriptive name for the storage account (i.e. “qliksense”) and create the account using the following options (see screenshot below). Please, notice that for the Subscription field the value may vary. Also, it is recommended to use the same location as the Resource and Network Security Groups. Once all values have been entered, click “Create” button.

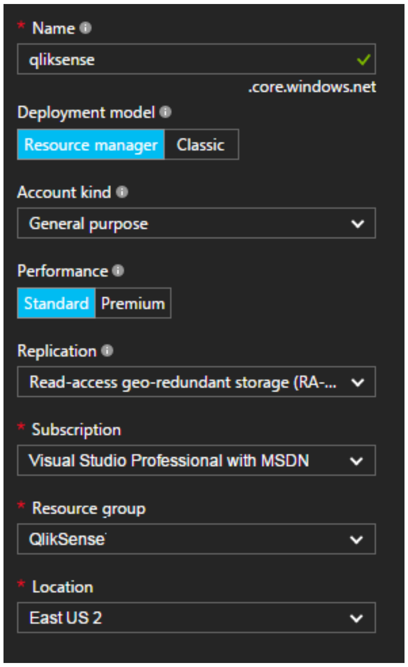

* Name
qliksense ✓
.core.windows.net

Deployment model
Resource manager | Classic

Account kind
General purpose

Performance
Standard | Premium

Replication
Read-access geo-redundant storage (RA-...

* Subscription
Visual Studio Professional with MSDN

* Resource group
QlikSense

* Location
East US 2

16. Create a Virtual Network by browsing on the left hand side panel for “Virtual Networks” and selecting the option

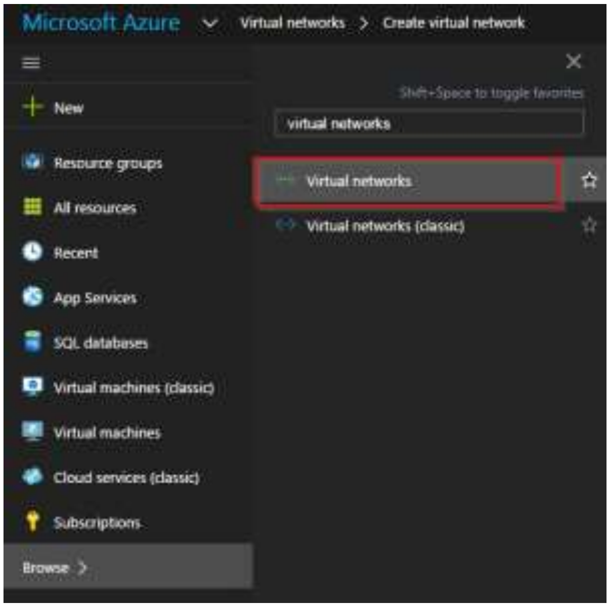

17. Click in the “Add” button

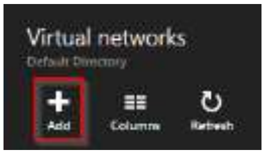

18. Enter the name for the Virtual Network and use all the default options for everything else except for the Subscription, Resource Group and Location fields. The Subscription field the value may vary depending on each customer. For the Resource Group, select the one that was previously created and for the Location select the same location as the Resource and Network Security Group. Once all the values have been set, click “Create”

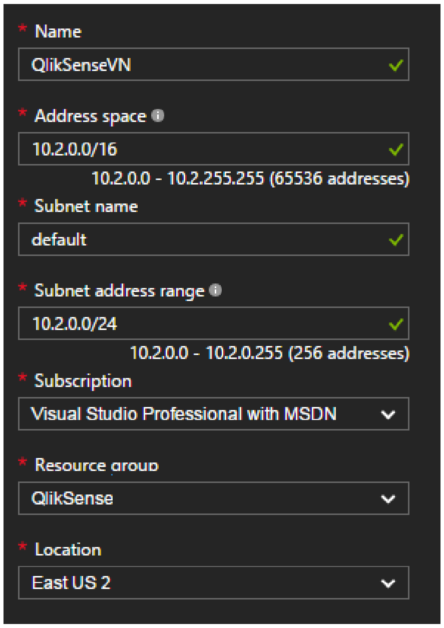

* Name
QlikSenseVN

* Address space
10.2.0.0/16
10.2.0.0 - 10.2.255.255 (65536 addresses)

* Subnet name
default

* Subnet address range
10.2.0.0/24
10.2.0.0 - 10.2.0.255 (256 addresses)

* Subscription
Visual Studio Professional with MSDN

* Resource group
QlikSense

* Location
East US 2

19. Create a Virtual Machine by clicking on the left hand side panel “Virtual Machines” and then selecting Window Server in the list of suggestions.

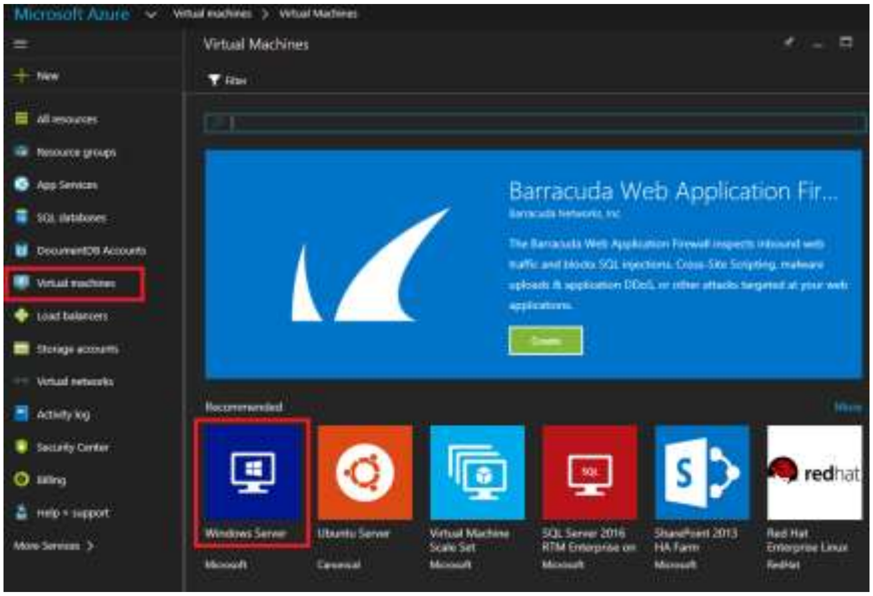

20. On the right hand side panel, select the “Windows Server 2012 R2 Datacenter” option.

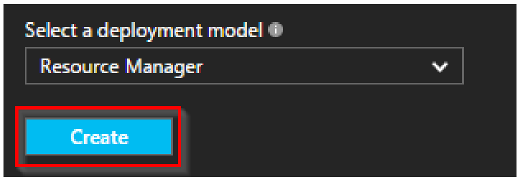

Select a deployment model
Resource Manager

21. In the new panel that will show up on the right hand side. Under the Deployment Model drop down menu, select the option called “Resource Manager” and click “Create”

22. Enter a name for the Virtual Machine as well as a username and password, the Subscription that applies, the Resource Group in question, and the location. Make sure to point to the Resource Group that was previously created and the same Location as the Resource Group.

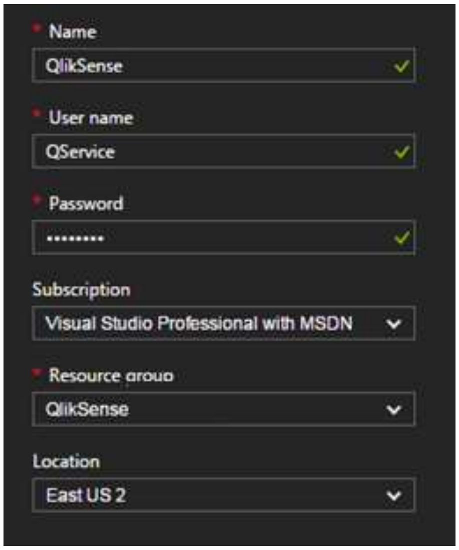

* Name
QlikSense ✓

* User name
QService ✓

* Password
•••••••• ✓

Subscription
Visual Studio Professional with MSDN

* Resource group
QlikSense

Location
East US 2

23. On the following screen where the different machine types are presented, click on the “View All” link

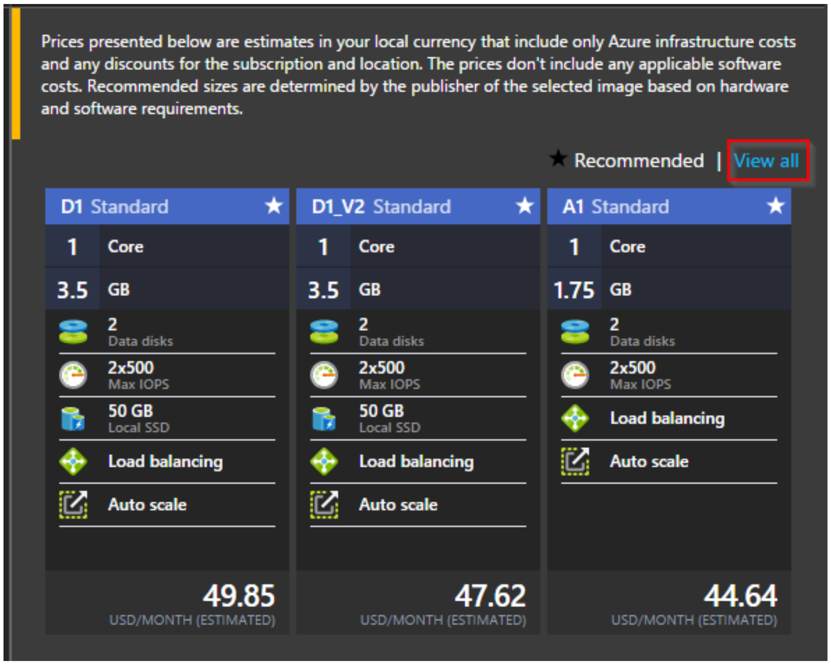

Prices presented below are estimates in your local currency that include only Azure infrastructure costs and any discounts for the subscription and location. The prices don't include any applicable software costs. Recommended sizes are determined by the publisher of the selected image based on hardware and software requirements.

★ Recommended | View all

D1 Standard ★	D1_V2 Standard ★	A1 Standard ★
1 Core	1 Core	1 Core
3.5 GB	3.5 GB	1.75 GB
2 Data disks	2 Data disks	2 Data disks
2x500 Max IOPS	2x500 Max IOPS	2x500 Max IOPS
50 GB Local SSD	50 GB Local SSD	Load balancing
Load balancing	Load balancing	Auto scale
Auto scale	Auto scale	
49.85 USD/MONTH (ESTIMATED)	47.62 USD/MONTH (ESTIMATED)	44.64 USD/MONTH (ESTIMATED)

24. There are many options that can work well to host a Qlik Sense environment. Typically, DS3 or DS12 is recommended as a starting point. Please, notice that other larger instances are available and could be used. In order to determine if a larger server is needed, please refer to the next section: “Sizing and Scalability”.

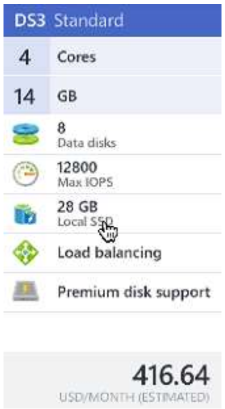

DS3 Standard
4 Cores
14 GB
8 Data disks
12800 Max IOPS
28 GB Local SSD
Load balancing
Premium disk support
416.64 USD/MONTH (ESTIMATED)

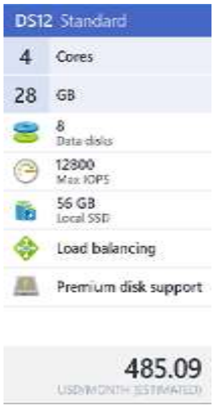

DS12 Standard
4 Cores
28 GB
8 Data disks
12800 Max IOPS
56 GB Local SSD
Load balancing
Premium disk support
485.09 USD/MONTH (ESTIMATED)

25. A new pane called “Settings” will show up. Under “Storage Account” select the storage account that was previously created (should be selected by default). Under the “Virtual Network” field select the virtual network that was previously created (should be selected by default). Leave the field “Subnet” as subnet.

26. Click on the “Public IP address” and two panes will show up. In the right-most pane, change the value in the “Assignment” field from “Dynamic” to “Static” and click “Ok”

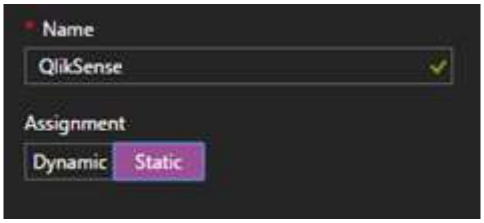

* Name
QlikSense ✓

Assignment
Dynamic [✓ Static]

27. Click “Network Security Group” and two new panels will show up. Select the security group that was previously created

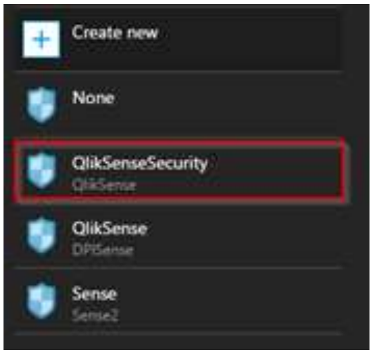

28. Leave the “Extensions” field as “No Extensions”. Under Monitoring Diagnostic option, click “Disabled” and under “Availability set leave it as “None”. The configuration should look like the screenshot below. Once all changes have been made, confirmation the validation and click “Ok”.

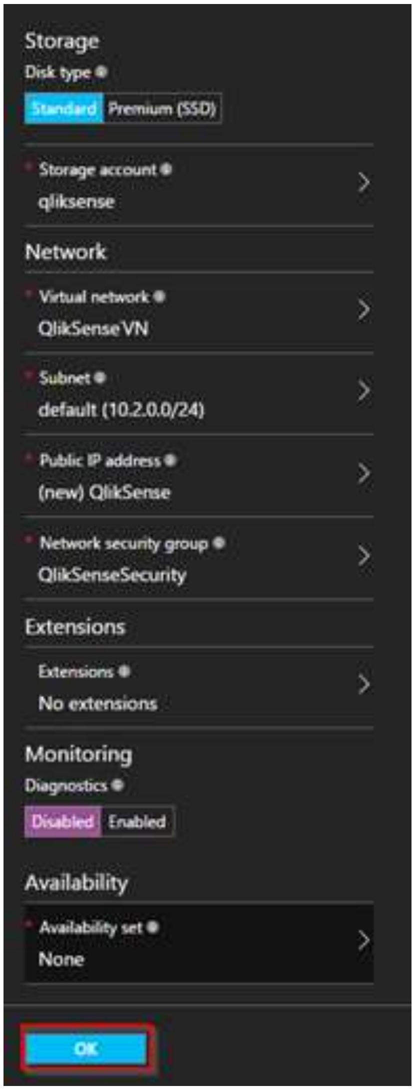

Storage

Disk type: [✓ Standard] Premium (SSD)

Storage account: qliksense

Network

Virtual network: QlikSense VN

Subnet: default (10.2.0.0/24)

Public IP address: (new) QlikSense

Network security group: QlikSenseSecurity

Extensions

Extensions: No extensions

Monitoring

Diagnostics: [✓ Disabled] Enabled

Availability

Availability set: None

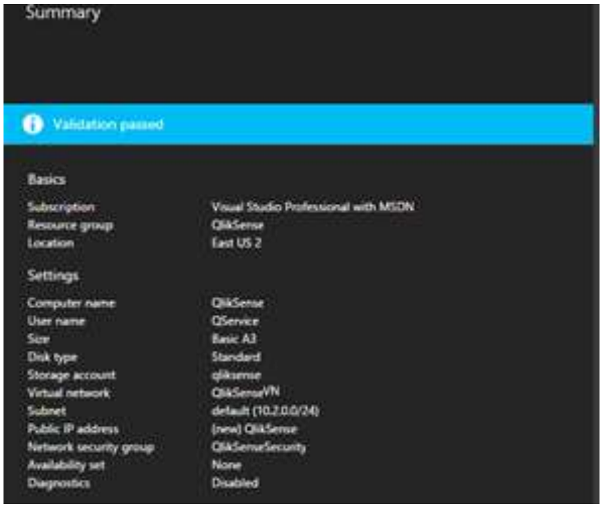

Summary

Validation passed

Basics

Subscription	Visual Studio Professional with MSDN
Resource group	QlikSense
Location	East US 2

Settings

Computer name	QlikSense
User name	QService
Size	Basic A3
Disk type	Standard
Storage account	qliksense
Virtual network	QlikSenseVN
Subnet	default (10.2.0.0/24)
Public IP address	(new) QlikSense
Network security group	QlikSenseSecurity
Availability set	None
Diagnostics	Disabled

29. The Virtual Machine will be deployed and started. This process may take few minutes.

30. Once the Virtual Machine has been deployed, the settings panel will appear automatically. Also, the settings panel would show up on the left hand side pane and could be accessed by selecting the virtual machine under the “Virtual Machines” section. Set a DNS name label in order to see the Public IP address.

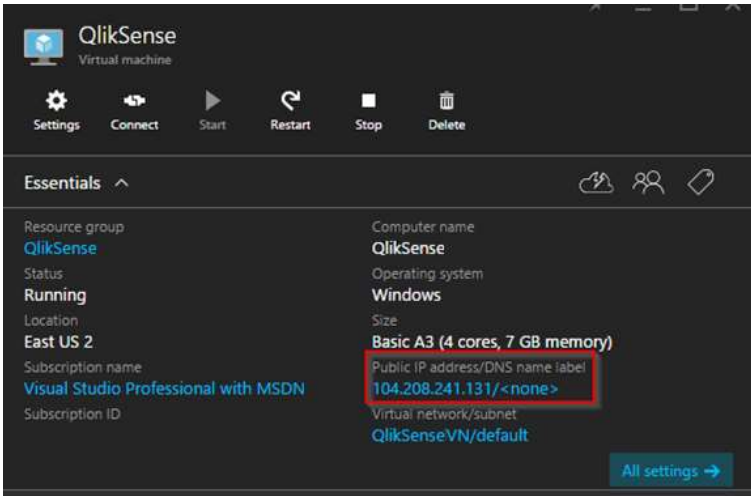

Resource group
QlikSense
Status
Running
Location
East US 2
Subscription name
Visual Studio Professional with MSDN
Subscription ID

Computer name
QlikSense
Operating system
Windows
Size
Basic A3 (4 cores, 7 GB memory)
Public IP address/DNS name label
104.208.241.131/<none>
Virtual network/subnet
QlikSenseVN/default

31. Select the “Configuration” option under the right-most “Settings” pane that will show up

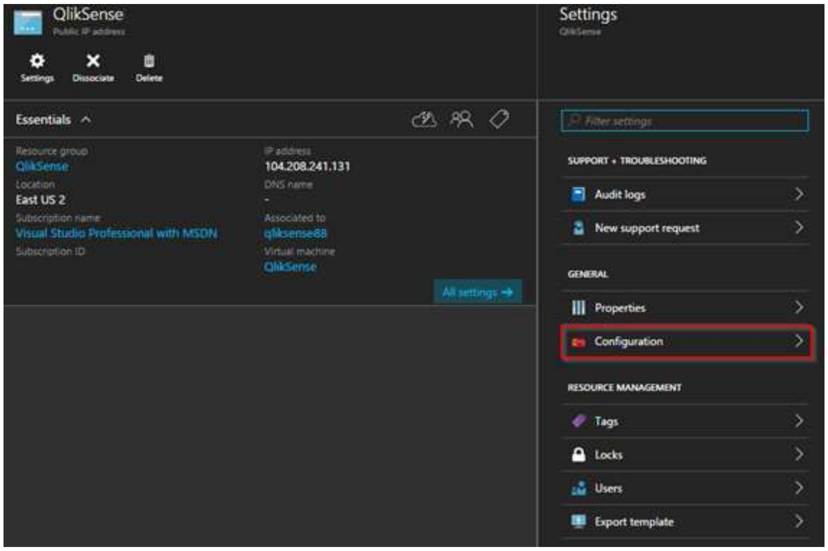

Resource group
QlikSense
Location
East US 2
Subscription name
Visual Studio Professional with MSDN
Subscription ID

IP address
104.208.241.131
DNS name
-
Associated to
qliksense88
Virtual machine
QlikSense

32. In the new pane that shows up, enter a unique name for the machine under the “DNS name label (optional)” input box

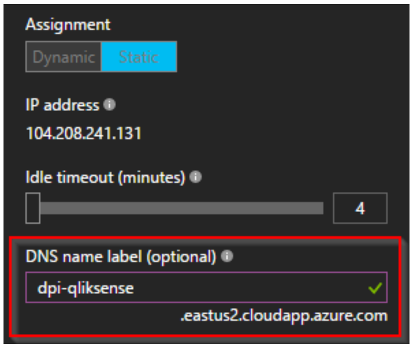

Assignment
Dynamic Static

IP address
104.208.241.131

Idle timeout (minutes)
4

DNS name label (optional)
dpi-qliksense
.eastus2.cloudapp.azure.com

33. Click on the “Save” icon to save the changes

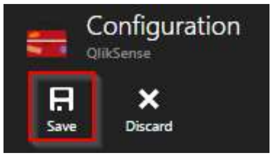

34. Connect to the Virtual Machine by selecting the machine under the “Virtual Machines” menu and clicking “Connect”

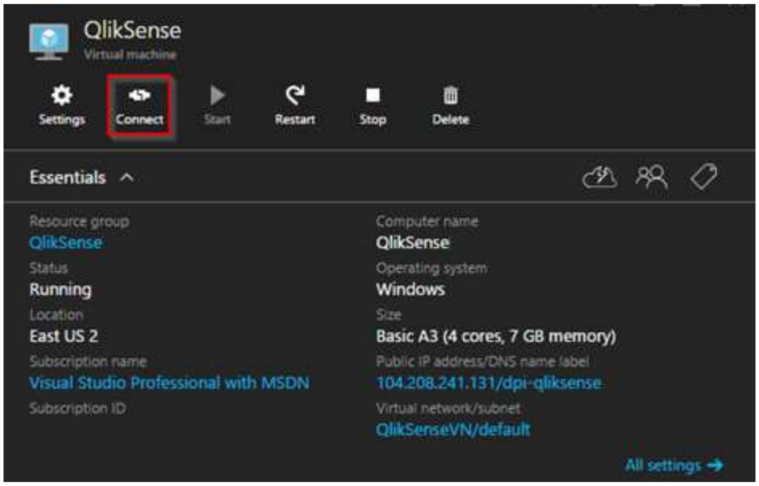

Resource group
QlikSense

Status
Running

Location
East US 2

Subscription name
Visual Studio Professional with MSDN

Subscription ID

Computer name
QlikSense

Operating system
Windows

Size
Basic A3 (4 cores, 7 GB memory)

Public IP address/DNS name label
104.208.241.131/dpi-qliksense

Virtual network/subnet
QlikSenseVN/default

35. Connect through the RDP file by double clicking on it

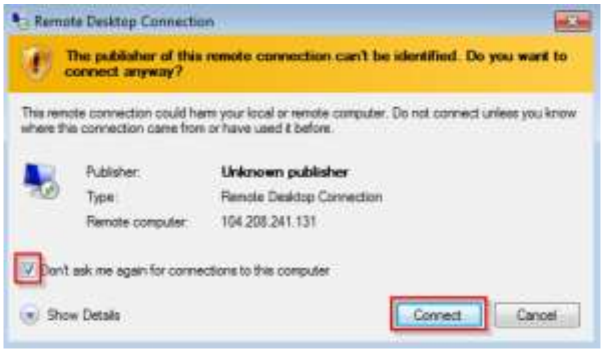

The publisher of this remote connection can't be identified. Do you want to connect anyway?

This remote connection could harm your local or remote computer. Do not connect unless you know where this connection came from or have used it before.

Publisher: **Unknown publisher**
Type: Remote Desktop Connection
Remote computer: 104.208.241.131

☑ Don't ask me again for connections to this computer

36. Enter the credentials using the IP and the username and password that were previously created

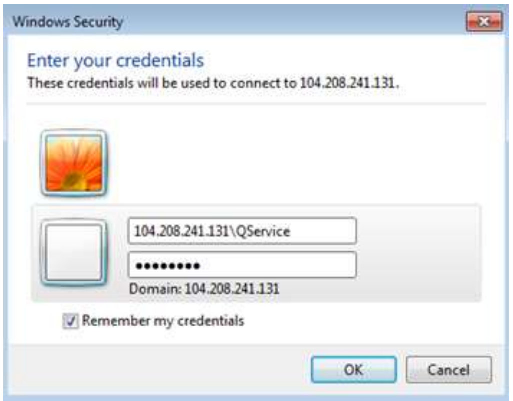

Enter your credentials
These credentials will be used to connect to 104.208.241.131.

104.208.241.131\QService
••••••••
Domain: 104.208.241.131

☑ Remember my credentials

37. Proceed through the certificate warning by clicking “Yes”

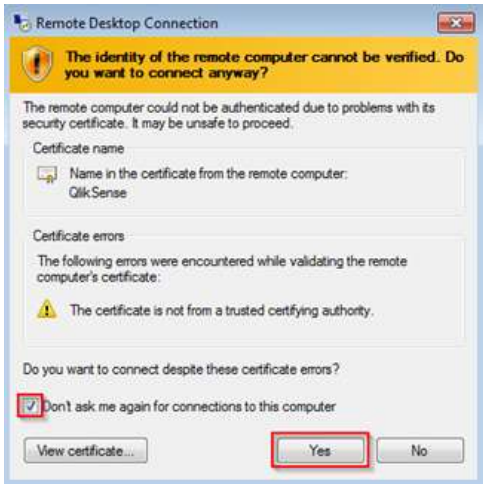

Remote Desktop Connection

The identity of the remote computer cannot be verified. Do you want to connect anyway?

The remote computer could not be authenticated due to problems with its security certificate. It may be unsafe to proceed.

Certificate name

Name in the certificate from the remote computer:
QlikSense

Certificate errors

The following errors were encountered while validating the remote computer's certificate:

The certificate is not from a trusted certifying authority.

Do you want to connect despite these certificate errors?

☑ Don't ask me again for connections to this computer

38. Once in the Virtual Machine, proceed with the installation of Qlik Sense Enterprise. Given that the new instance is not part of an existing domain, some local users need to be set up in order to use Qlik Sense as well as a Service Account to run Qlik Sense Server. In order to create a Service Account, follow the next steps. Notice that this account will be used during the installation of Qlik Sense

a. On the new instance that is going to be used as the Qlik Sense server, open the Computer Management window

b. Find the Users folder, then click Action > New User

c. Enter a user name (i.e. 'QSAdmin') and a password (i.e. 'QlikSense!')

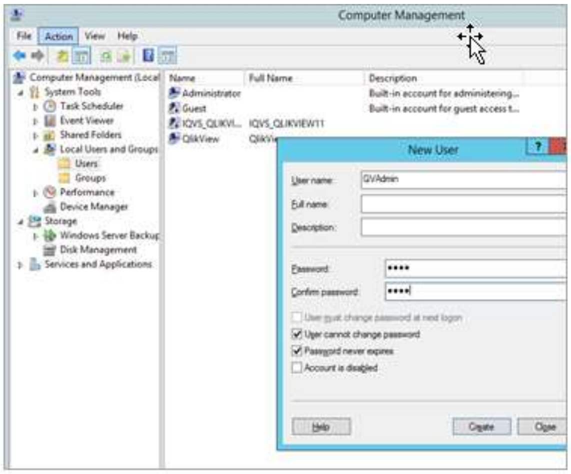

Name	Full Name	Description
Administrator		Built-in account for administering...
Guest		Built-in account for guest access t...
IQVS_QLIKVI...	IQVS_QLIKVIEW11	
QlikView	QlikVi	

New User

User name: QVAdmin

Full name:

Description:

Password: ••••

Confirm password: ••••

□ User must change password at next logon

☑ User cannot change password

☑ Password never expires

□ Account is disabled

d. Uncheck 'User must change password at next logon'

e. Check 'User cannot change password' and 'Password never expires'

f. Click Create

g. Double click on 'QSAdmin'

h. Click 'Member Of'

i. Click Add

j. Type 'Administrators'

k. Click 'Check Names'

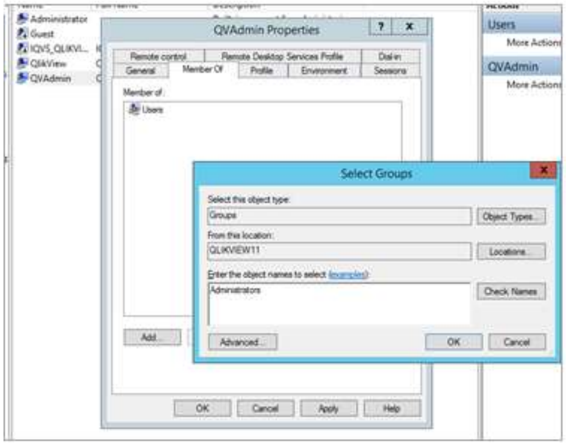

Member of:
Users

Select this object type:
Groups

From this location:
QLIKVIEW11

Enter the object names to select (examples):
Administrators

l. Click 'OK'

m. Click 'Users'

n. Click 'Remove'

o. Click 'OK'

In order to create local users, follow the next steps. These are the users that will be used to log into Qlik Sense.

a. On the new instance that is going to be used as the Qlik Sense server, open the Computer Management window

b. Find the Users folder, then click Action > New User

c. Add as many users as needed (i.e. 'QlikUser1', 'QlikUser2', 'QlikUser3', etc.)

d. Enter a Password (i.e. 'QlikSense!')

e. If applicable, uncheck 'User must change password at next logon'

f. If applicable, check 'User cannot change password'

g. If applicable, check 'Password never expires'

h. Click 'Create'

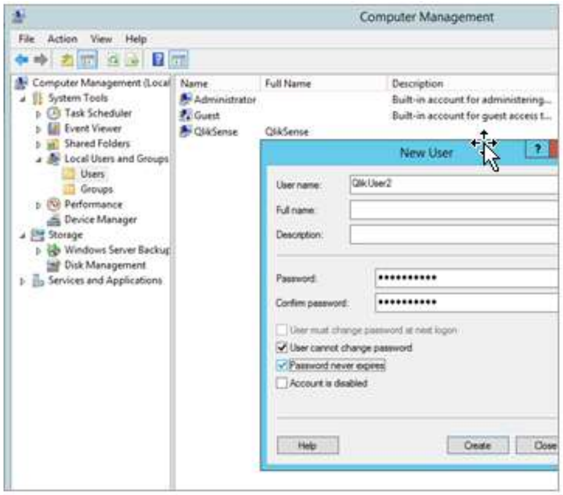

Name	Full Name	Description
Administrator		Built-in account for administering...
Guest		Built-in account for guest access t...
QlikSense	QlikSense	

User name: QlikUser2

Full name:

Description:

Password: ••••••••••

Confirm password: ••••••••••

□ User must change password at next logon

☑ User cannot change password

☑ Password never expires

□ Account is disabled

39. It is typically recommended to document the setup of all accounts in a table format to make it easy for administrators to keep control of such accounts. For example, the table below exemplifies a way in which such information can be stored. All security rules for password creation that are specific to an environment, should always be reinforced in Qlik as well.

User Name	Password	Account Settings	Windows Role Description	Function Description
QlikService	QlikSense!	• Cannot change password • Password never expires	Part of Local Administrator group	Windows Service Account to run the Windows Services for Qlik
Repository database Super user Password	aaabbbccc	Not applicable. This account is not a windows account		PostgreSQL database Supervisor
QSAdmin	QlikSense!	If applicable, uncheck 'User must change password at next logon'	Local or domain user	User who will be the Qlik Administrator
Optional Accounts				
QlikUser	abcdef	If applicable, uncheck 'User must change password at next logon'	Local or domain user	User who will have access to Qlik

40. Once all users have been created, install Qlik Sense Enterprise.

   a. Using a local Admin account, download the latest version of Qlik Sense Enterprise from http://www.qlik.com/download
   b. Run the Qlik Sense Installer
   c. Choose "Full Installation" (unless you need to install to a different path, then choose custom install)
   d. Accept the license agreement

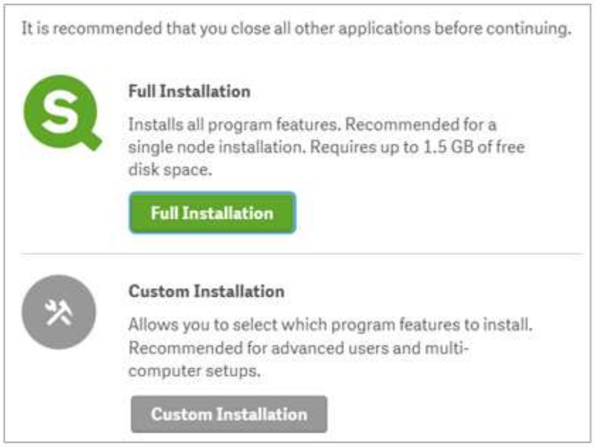

It is recommended that you close all other applications before continuing.

Full Installation

Installs all program features. Recommended for a single node installation. Requires up to 1.5 GB of free disk space.

Custom Installation

Allows you to select which program features to install. Recommended for advanced users and multi-computer setups.

e. Choose **“Central”**

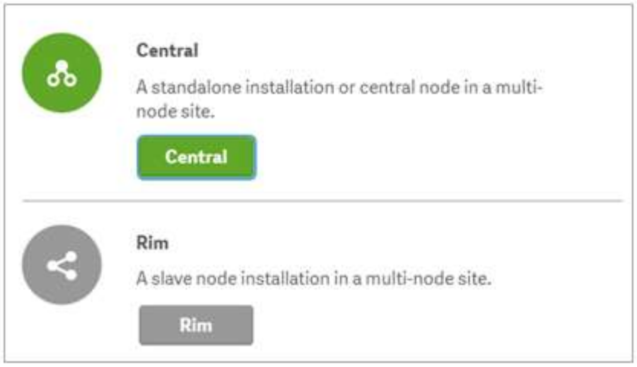

Central

A standalone installation or central node in a multi-node site.

Rim

A slave node installation in a multi-node site.

f. Enter a repository password. It can be anything, but write it down! It is recommended that the password be made the same as the service account user’s password (i.e. ‘QlikSense!’)

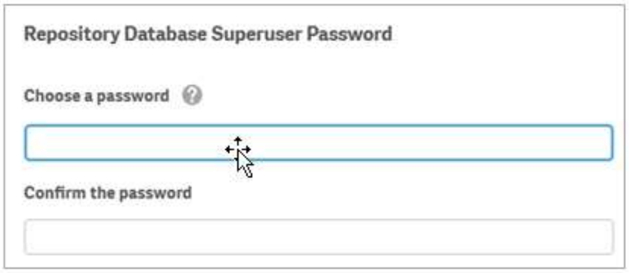

Repository Database Superuser Password

Choose a password

Confirm the password

g. Enter the service account user and password. In the case that the instance would be running as part of a then the format for the domain account would be DOMAIN\USER. Given that a local account is being used, the format is MACHINENAME\USER. Remember this is a Local Admin account.

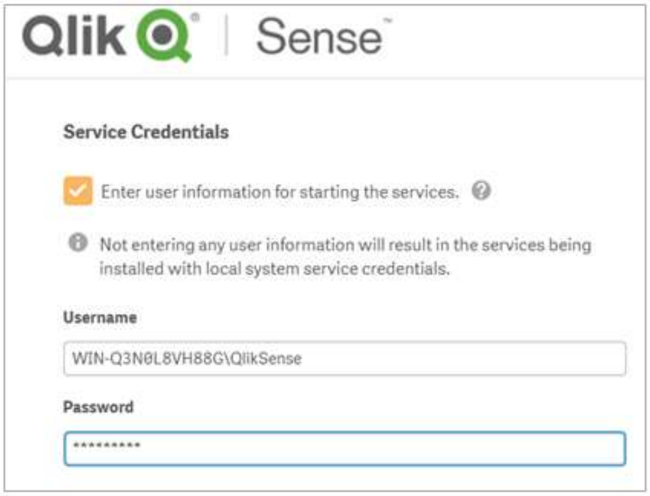

Service Credentials

☑ Enter user information for starting the services.

Not entering any user information will result in the services being installed with local system service credentials.

Username

WIN-Q3N0L8VH88G\QlikSense

Password

****

h. **The following is a critical step. Pay close attention.** Enter the machine name without the domain name. Avoid using the fully qualified name. **Do not using the CLOUDAPP.net fully qualified domain name (FQDN).** The certificates created in this step are used for ‘Intra-service’ communication and for external user access in HTTPS, hence it is critical not to use the FQDN.

For example:

IP address – 10.1.123.234

machine name – WIN-Q3N0L8VH88G

fully qualified machine name – WIN-Q3N0L8VH88G.ExampleDomain.COM

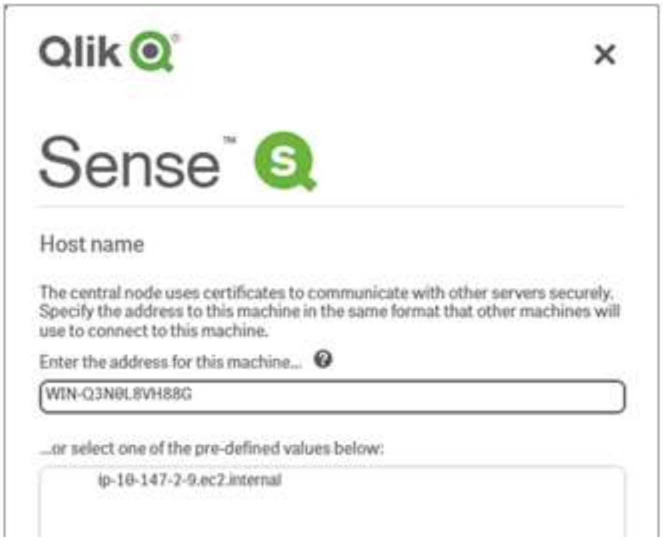

Host name

The central node uses certificates to communicate with other servers securely. Specify the address to this machine in the same format that other machines will use to connect to this machine.

Enter the address for this machine...

WIN-Q3N0L8VH88G

...or select one of the pre-defined values below:

ip-10-147-2-9.ec2.internal

i. On the next screen, click 'Install'

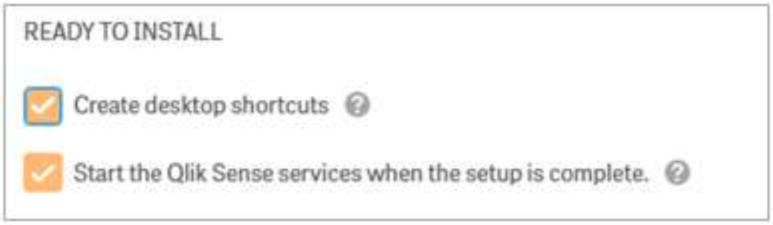

READY TO INSTALL

☑ Create desktop shortcuts

☑ Start the Qlik Sense services when the setup is complete.

j. Upon confirmation that all the services are up and running, click 'Finish'

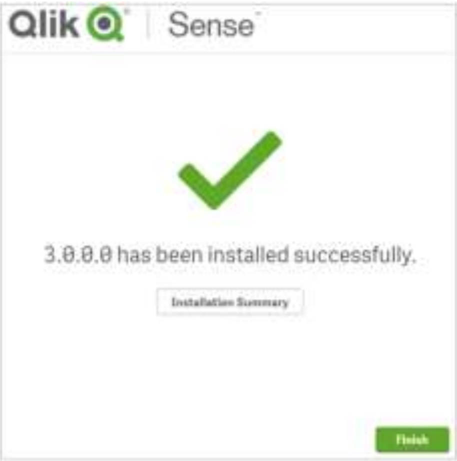

3.0.0.0 has been installed successfully.

41. Once the installation has been completed, the Qlik Sense Management Console (QMC) needs to be set up. The QMC is used for a number of different administrative and monitoring tasks including entering the license to validate the installation of Qlik Sense Server. In order to launch the QMC, open the browser and enter the following URL

    https://machinename/qmc

    If the QMC doesn't come up the first time, give it 30 seconds. It means that the services are starting. A security warning may come up. If so, click through it. If you are prompted to login to the QMC, login with the same local Administrator account used to RDC into the Qlik Sense Server.

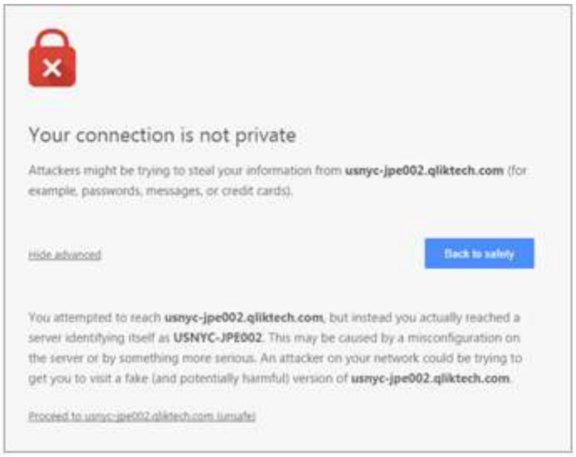

Your connection is not private

Attackers might be trying to steal your information from **usnyc-jpe002.qliktech.com** (for example, passwords, messages, or credit cards).

You attempted to reach **usnyc-jpe002.qliktech.com**, but instead you actually reached a server identifying itself as **USNYC-JPE002**. This may be caused by a misconfiguration on the server or by something more serious. An attacker on your network could be trying to get you to visit a fake (and potentially harmful) version of **usnyc-jpe002.qliktech.com**.

42. Once the QMC opens, it will prompt for the license information. Enter it and click Get LEF from server. The validation of the license will require internet access. If the server doesn't have internet access, a LEF file will be needed from a Qlik representative.

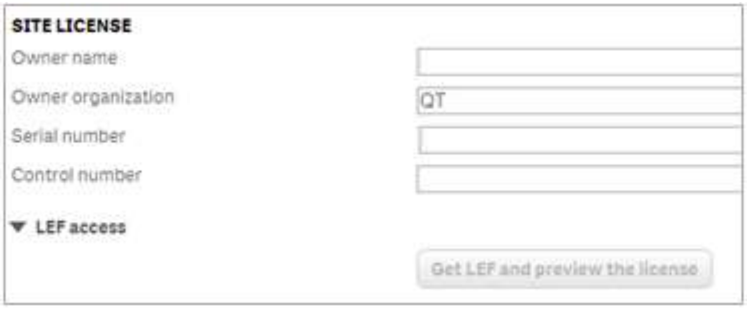

SITE LICENSE

Owner name	
Owner organization	QT
Serial number	
Control number	

Once validated, a message that tokens have been successfully allocated will show up.

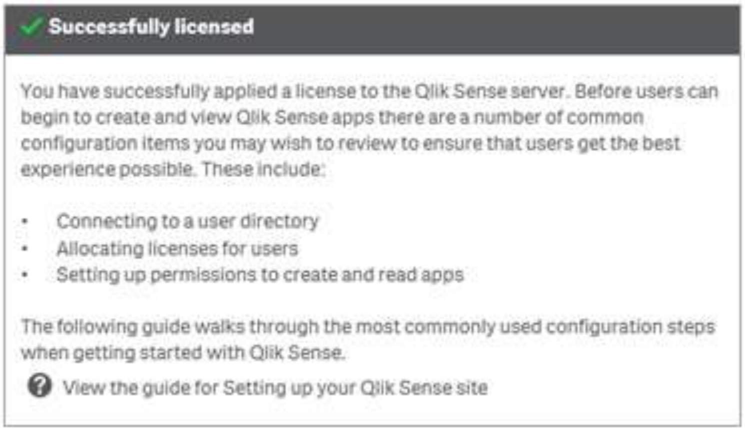

✓ **Successfully licensed**

You have successfully applied a license to the Qlik Sense server. Before users can begin to create and view Qlik Sense apps there are a number of common configuration items you may wish to review to ensure that users get the best experience possible. These include:

- Connecting to a user directory
- Allocating licenses for users
- Setting up permissions to create and read apps

The following guide walks through the most commonly used configuration steps when getting started with Qlik Sense.

43. Now that the license has been validated, the tokens need to be assigned to users accordingly. The tokens can be assigned automatically by users based on User Access Rules. In order to learn more about how to manage QMC resources, visit the following link:

http://help.qlik.com/en-US/sense/2.2/Subsystems/ManagementConsole/Content/manage-QMC-resources.htm

In order to do so, follow the next steps:

a. Click on the Start button on the top left corner of the QMC > License and Tokens > User Access Allocations

b. Click the Allocate button.

c. Select the id you are logged in with (i.e. Administrator), and click Allocate. A ‘Successfully Allocated’ message will display at the bottom of the browser. This is so the administrator always has a license. Notice that from Governance point of view, the ‘administrator’ account should never be used to configure Qlik Sense Enterprise. The configuration should be done using a Qlik Administrator account. The Qlik Sense Administration should log into Windows using that Qlik Sense Administrator account (i.e. QSAdmin).

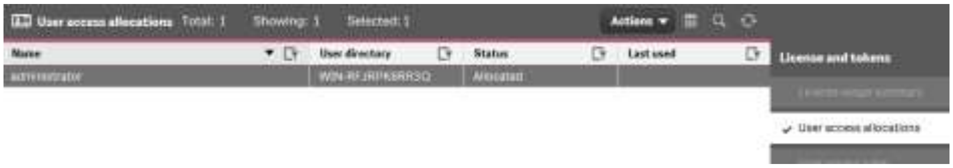

User access allocations Total: 1 Showing: 1 Selected: 1

Name	User directory	Status	Last used
administrator	WIN-RFJRPK8RR3Q	Allocated	

d. Click on the Start button on the top left corner of the QMC > License and Tokens > User Access Rules >

e. Click “Create New”

f. Click “Basic”

g. Create the rule as ‘user name like value *’

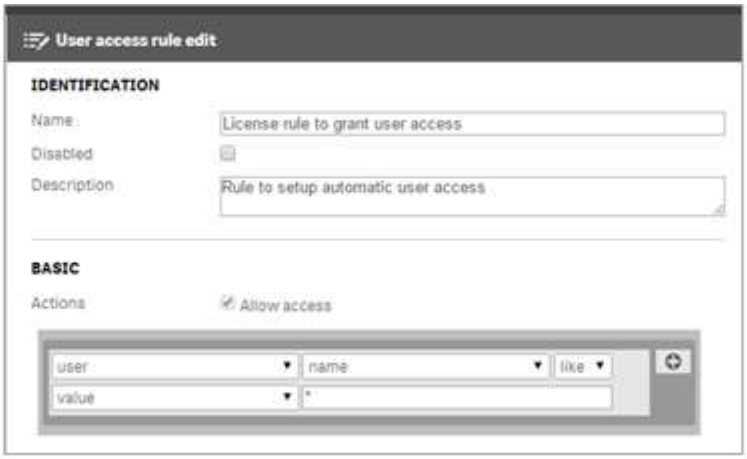

User access rule edit

IDENTIFICATION

Name: License rule to grant user access

Disabled: □

Description: Rule to setup automatic user access

BASIC

Actions: ☑ Allow access

user | name | like
value | *

h. Click Apply. Anyone connecting to Qlik Sense Server will get a license. This is useful for development and testing purposes but may not be adequate for a Production deployment. To learn more on how to manage a full Production environment, visit the following link:

http://help.qlik.com/en-US/sense/2.2/Subsystems/ManagementConsole/Content/manage-license-and-tokens.htm

i. Typically, rather than adding individual rules, a User Directory would be created and configured and the rules to assigned tokens would be based on such User Directory. In the simplest scenario, a User Directory Connector would be created. The local Directory could be called anything (i.e. QLIK):

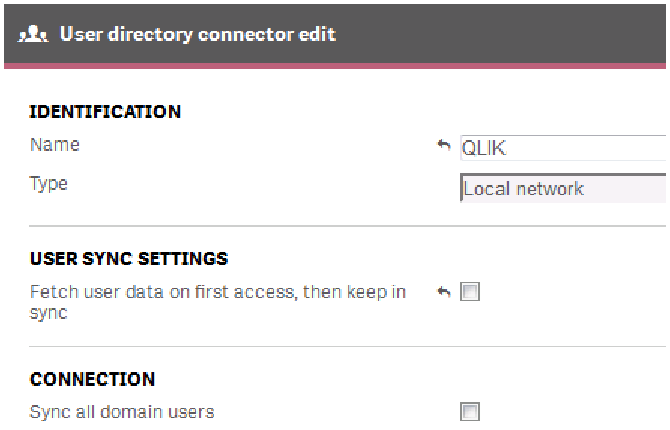

User directory connector edit

IDENTIFICATION

Name: QLIK

Type: Local network

USER SYNC SETTINGS

Fetch user data on first access, then keep in sync □

CONNECTION

Sync all domain users □

And then a User Access Rule could look like the following:

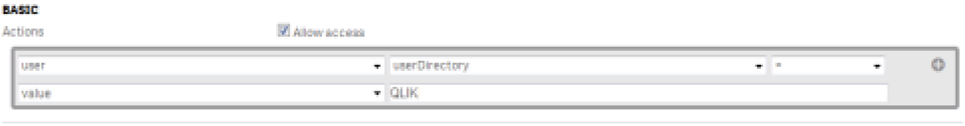

BASIC

Actions ☑ Allow access

user	userDirectory	=
value	QLIK	

In order to learn more about User Directory Connector, visit the following link:

http://help.qlik.com/en-US/sense/3.0/Subsystems/ManagementConsole/Content/create-user-directory-connector.htm

44. Using the QMC, a proxy can be set up in order to enable HTTP access on Qlik Sense Server and allow access to Qlik Sense Server over alternate URL formats. For example, to use HTTP instead of HTTPS in your browser and avoid security warnings. To learn more about the Proxies section of the QMC, visit the following link:

http://help.qlik.com/en-US/sense/2.2/Subsystems/ManagementConsole/Content/proxies-overview.htm

Follow the next steps to enable HTTP and add the machine name as part of the proxy values:

a. Open the QMC > click on the Start button on the top left corner > Proxies > Central > Edit

b. Click 'Ports' on the right hand side

c. Check Allow HTTP

d. Click Apply (a message saying that the proxy will be restarted will show up. Click OK)

e. Click on the Start button on the top left corner > Virtual Proxies > "Central Proxy (Default)" > Edit

f. Click Advanced

g. Scroll down, add values to Host white list to add the following four values

h. Add the Public IP address of the Qlik Sense Server (found the configuration of the Virtual Machine)

i. Click Add New Value

j. Add the Virtual Machine name (Hostname) of the Qlik Sense Server

k. Click Add New Value

l. Add the Public DNS name label (i.e. WIN-Q3N0L8VHHHG.cloudapp.net)

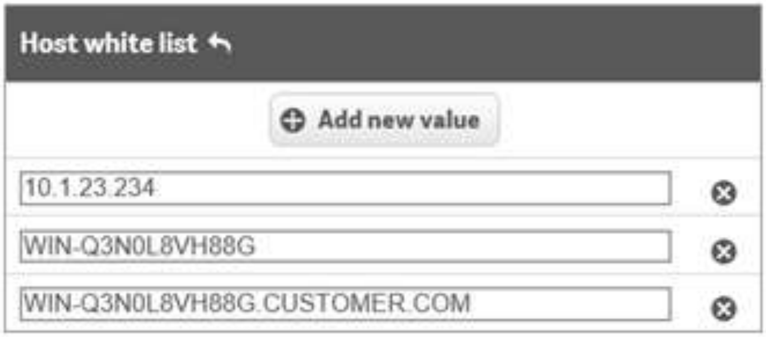

10.1.23.234

WIN-Q3N0L8VH88G

WIN-Q3N0L8VH88G.CUSTOMER.COM

m. Click Apply. A warning saying that the proxy will have to be restarted will show up. Click OK. If any DNS entries has been created (e.g., qlikbi.company.com), add this to the whitelist, too. This step may require to import own SSL certificate. Adding these values to the Websocket Origin Whitelist allows Qlik Sense to accept URLs of these formats.

45. The QMC and Qlik Sense HUB should be fully functional at this point. In order to perform a quick test, the two shortcuts that were created on the Desktop of the Virtual Machine Instance could be used. The Qlik Sense HUB is the single point of entry for all users to perform a number of different activities such as create new applications, access existing applications, etc. In order to learn more about the Qlik Sense HUB, visit the following link:

http://help.qlik.com/en-US/sense/2.2/Subsystems/Hub/Content/Hub/hub.htm?q=hub

You will know Qlik Sense Server is setup correctly if the browser gives no security warnings. Notice that the shortcuts that were automatically created will be in the form of:

- http://MachineName/hub
- http://MachineName/qmc

Which would only work locally within the Virtual Machine Virtual Network where the Qlik Sense Server instance was created. In order to enable external access using other URL formats that were configured during the Qlik Sense Server setup, it is needed to open up some ports. The following steps explain how to do so:

a. Go to the Windows Firewall with Advanced Security settings (go to Control Panel and search for Windows Firewall and then go to the Advance Settings).

b. Click on 'Inbound Rules' in the left panel, then click on 'New Rule' in the right panel.

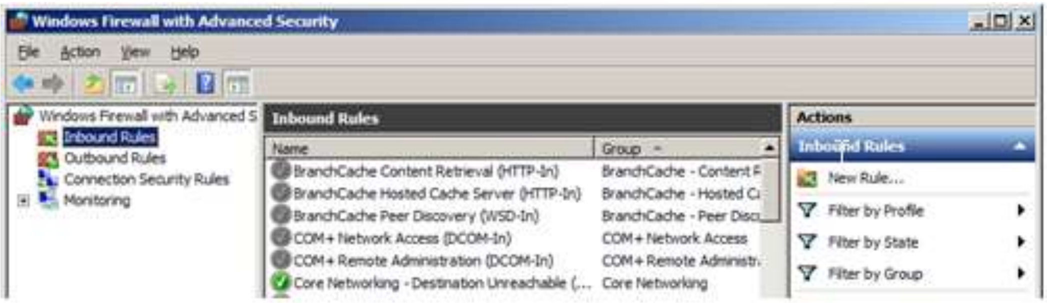

Name	Group
BranchCache Content Retrieval (HTTP-In)	BranchCache - Content F
BranchCache Hosted Cache Server (HTTP-In)	BranchCache - Hosted C
BranchCache Peer Discovery (WSD-In)	BranchCache - Peer Disc
COM+ Network Access (DCOM-In)	COM+ Network Access
COM+ Remote Administration (DCOM-In)	COM+ Remote Administr
Core Networking - Destination Unreachable (...	Core Networking

c. Select ‘Port’ as the rule type.

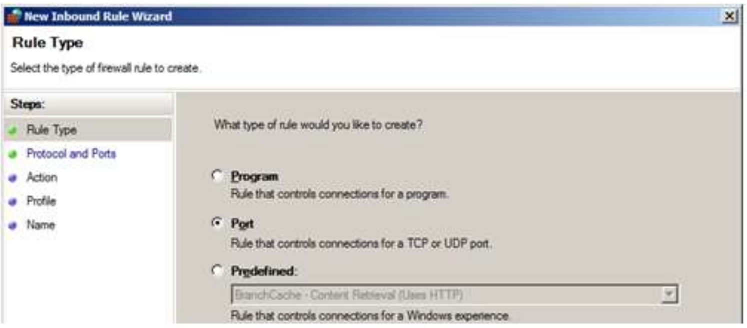

Rule Type

Select the type of firewall rule to create.

What type of rule would you like to create?

○ Program
Rule that controls connections for a program.

● Port
Rule that controls connections for a TCP or UDP port.

○ Predefined:
BranchCache - Content Retrieval (Uses HTTP)
Rule that controls connections for a Windows experience.

d. In the ‘Specific local ports’ field, enter ’80, 443, 4244, 4248’. This will allow you to access the Qlik Sense Hub and QMC. As a reminder, probably HTTP and HTTP were not both enabled, so only the appropriate port out of these two would have to be used here.

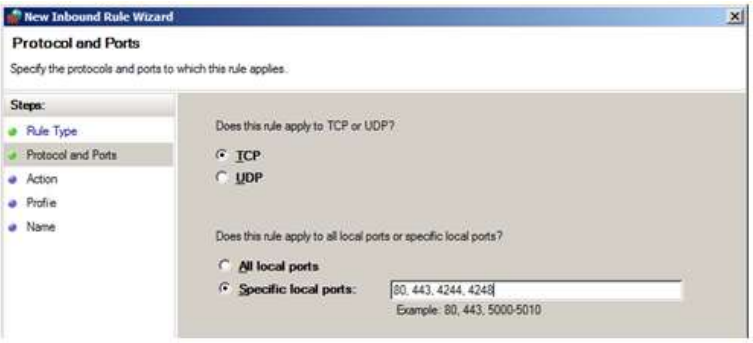

Protocol and Ports

Specify the protocols and ports to which this rule applies.

Does this rule apply to TCP or UDP?

● TCP
○ UDP

Does this rule apply to all local ports or specific local ports?

○ All local ports
● Specific local ports: 80, 443, 4244, 4248
Example: 80, 443, 5000-5010

e. Ensure that ‘Allow the connection’ is selected and click ‘Next’.

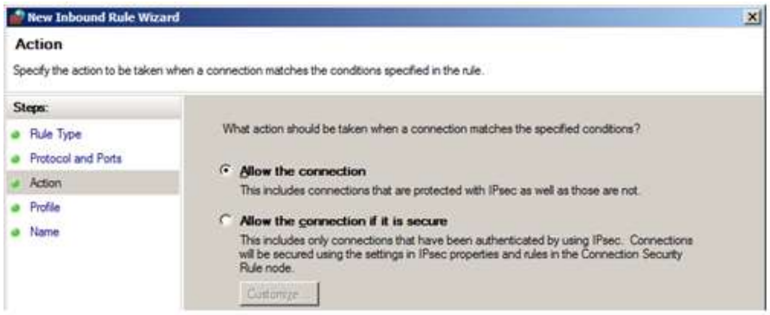

Action

Specify the action to be taken when a connection matches the conditions specified in the rule.

What action should be taken when a connection matches the specified conditions?

● Allow the connection
This includes connections that are protected with IPsec as well as those are not.

○ Allow the connection if it is secure
This includes only connections that have been authenticated by using IPsec. Connections will be secured using the settings in IPsec properties and rules in the Connection Security Rule node.

f. Ensure that all options are ticked and click ‘Next’.

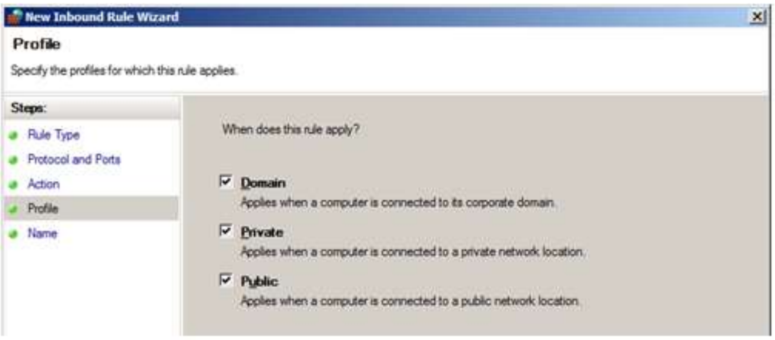

Profile

Specify the profiles for which this rule applies.

When does this rule apply?

☑ Domain
Applies when a computer is connected to its corporate domain.

☑ Private
Applies when a computer is connected to a private network location.

☑ Public
Applies when a computer is connected to a public network location.

g. Give the rule a name such as ‘Qlik Sense’ and then click ‘Finish’. This will now allow traffic on those ports through the firewall to the operating system.

h. Close Windows Firewall.

46. In order to test access to the Qlik Sense HUB and QMC from outside the Azure Qlik Sense Server instance, the following formats could be used:

    These forms would display a warning with the website security certificate. Continuing to the website and passing one of the Qlik Sense user’s credentials (i.e. Administrator, QlikUser1) would open up either the Qlik Sense HUB or QMC:

    - https://AzurePublicDSN /hub
    - https://AzureInstancePublicIP/hub
    - https://AzurePublicDSN/qmc
    - https://AzureInstancePublicIP/qmc

    These forms (in case that HTTP was enabled during the configuration) would prompt for the Qlik Sense user’s credentials (i.e. Administrator, QlikUser1) and direct users to either the Qlik Sense HUB or QMC:

    - http://AzurePublicDSN /hub
    - http://AzureInstancePublicIP/hub
    - http://AzurePublicDSN/qmc
    - http://AzureInstancePublicIP/qmc

47. If the request for access is being made from a device running iOS, notice that as of iOS version 8, untrusted certificates are not allowed. This means that the self-signed certificates that Qlik Sense Server uses out of the box are not sufficient to enable Qlik Sense HUB nor QMC access on iOS devices. To learn more about how Qlik Sense Server leverages certificates for authentication purposes, visit the following links:

    http://help.qlik.com/en-US/sense/2.2/Subsystems/PlanningQlikSenseDeployments/Content/Server/Server-Security-Authentication-Certificate-Trust.htm?q=certificate

    http://help.qlik.com/en-US/sense/2.2/Subsystems/PlanningQlikSenseDeployments/Content/Server/Server-Security-Authentication.htm

As an option, the request could be made using HTTP instead of HTTPS. Another option would be to install a customer provided certificate on the Qlik Sense Server with private key from a trusted root certificate authority (e.g. Symantec, GoDaddy, Thawte, DigiCert, or many others) and add this certificate’s security thumbprint to the thumbprint text box in the Proxy configuration. Installing a trusted certificate in Azure Virtual Machines requires additional configuration. If a Client Provided Browser Certificate is available, the certificate could be imported into the Qlik Sense Sever environment by following the next steps:

a. Generate Certificate Signing Request for Trusted Certificate. If a certificate from a Trusted Certificate Authority is not available yet, visit the following link for instructions on how to generate the CSR.

https://community.qlik.com/docs/DOC-15740

b. Import Certificate to Windows Certificate Store by following the next steps:

   I. Click Start > type MMC > Right Click > Run as Administrator

   II. Click File > Add / Remove Snap In

   III. Click Certificates > Click Add > choose My User Account

   IV. Click Certificates > Add > choose Computer Account > choose Local Computer

   V. Navigate to each folder listed below and import the customer provided certificate

      - Certificates - Current User > Trusted Root Certification Authorities > Certificates
      - Certificates (Local Computer) > Trusted Root Certification Authorities > Certificates
      - Certificates (Local Computer) > Personal > Certificates

c. Validate Existence of Private Key by double clicking on the imported certificate to view its properties. Then validate the private key exists with this certificate. If not, ask the person who provided it.

d. Import Certificate Thumbprint to Qlik Sense Proxy by navigating to the Details tab and copy the Thumbprint value. Make sure to copy all leading and trailing spaces.

   I. Open QMC > Proxies > Central Proxy > Edit

   II. On the right hand side of the screen, click Security to enable additional properties.

   III. Paste the browser thumbprint, including all spaces.

   IV. Click Apply to restart the Proxy.

e. As a precaution, also enable HTTP traffic (see Proxy Setup section) if you haven't already for testing purposes. This will ensure that if this process has issues there is still a way to access the QMC.

Scalability and Sizing

Qlik Sense's services are the cornerstones of the Qlik Sense architecture and are configured to support a variety of enterprise-ready deployment scenarios. As environments grow in terms of number of users, number and size of applications, number of data sources, etc.it is important to understand how to size the environment correctly and how to scale the environment accordingly. Resources need to be allocated correctly across the following Qlik Sense Services:

- **Engine Service** – The QIX engine, provides in-memory Associative Data Indexing and calculation supporting analysis
- **Proxy Service** – Manages authentication, handles user sessions and load balancing
- **Repository Service** –Manages Qlik Sense applications, controls access, and handles configuration
- **Scheduling Service** – Manages reloads of Qlik Sense applications and other scheduled tasks

- **Service Dispatcher** – Launch and manage the data profiling service for data load model, migration service to make sure the app runs in the installed version of Qlik (runs on the Central Node only) and Chart sharing between two users

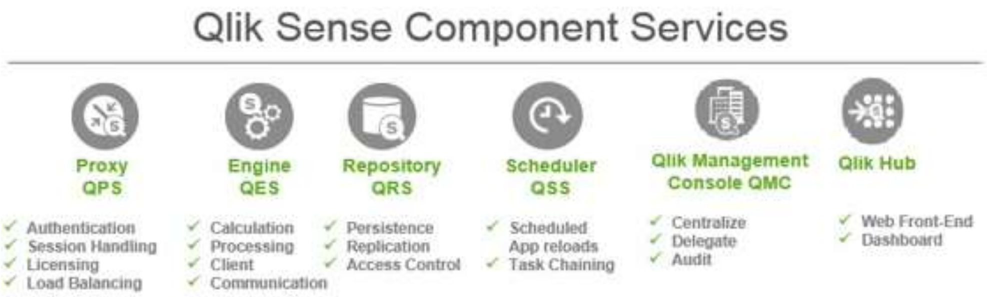

In order to provide some guidance on how to size an environment and understand better how Qlik Sense running on Azure can easily scale, a benchmarking analysis has been documented in a technical brief that is publicly available. The technical brief outlines performance benchmarks for Qlik Sense and is based on a testing methodology called the Qlik Capacity Benchmark. This series of tests is conducted by varying data volumes, users, and applications in a given environment and recording the results. This exhaustive set of permutations yields a matrix of CPU utilization, RAM utilization, and response times. This approach is different than many other scalability tests. Not only are metrics reported when a server is saturated, but also when the server is only partially utilized. This methodology provides transparency to the testing process and resulting metrics. While Qlik applications do vary in size and complexity, the methodology provides a comprehensive set of data with which customers can judge scalability and plan for deployments. For more information about Qlik Sense Architecture, visit the following links:

Qlik Sense Architectural Overview Whitepaper
http://www.qlik.com/resource-library/qlik-sense-architectural-overview

Qlik Sense Scalability Datasheet
http://www.qlik.com/resource-library/qlik-sense-scalability

*** Advanced:** Another source that is highly recommended is Qlik's YouTube channel. There is a series of 3 videos that talk more about Qlik Sense Scalability Tool:

https://www.youtube.com/watch?feature=player_embedded&v=8H0JJYyUft8

https://www.youtube.com/watch?feature=player_embedded&v=GiSF1yV9D9g

https://www.youtube.com/watch?v=3uPuCi7jzN0&feature=player_embedded

Below is the link to the Qlik® Sense Performance Benchmark Technical Brief:

http://www.qlik.com/resource-library/qlik-sense-performance-benchmark

Connectivity

As mentioned earlier on this document, Azure ecosystem is a collection of web services that collectively make up a cloud computing platform. Compared to buying and building a physical server farm, the three key benefits of Azure's cloud platform are:

- **Time to Value** – a platform that helps build enterprise applications quickly using well integrated web services
- **Flexibility** – ability to build hybrid applications leveraging assets that are on premises and cloud web services
- **Ease of use** – empowering everybody with technologies that they are already familiar with

Azure offers a number of webs services that can be used as data stores where Qlik Sense applications can pull data from in order to build applications. Such services include:

a. Azure SQL Database – SQL Server 2016 on the Cloud
https://azure.microsoft.com/en-us/documentation/services/sql-database/

b. Azure SQL Data Warehouse – scale-out, massively parallel processing, highly scalable database for both, relational and non-relational data
https://azure.microsoft.com/en-us/documentation/services/sql-data-warehouse/

c. Azure Storage – scalable cloud storage (Blob Storage, Table Storage, Azure Queues and Azure Files)
https://azure.microsoft.com/en-us/documentation/services/storage/

d. Azure HDInsight – elastic map reduce (Hadoop as Service)
https://azure.microsoft.com/en-us/documentation/services/hdinsight/

Qlik Sense server leverages various connectivity mechanisms in order to connect to different kind of sources. Whether it is an ODBC connection, OLE DB connection, REST API connection or a native connector to a specific source (i.e. SFDC, SAP, etc.), Qlik Sense will leverage that connection in order to model, read, parse and load the source data into analytical applications. Qlik Sense applications can leverage various type and number of sources simultaneously within the same application. Some common scenarios would be: (a) an application that uses both; data stored in Azure SQL Database and Azure SQL Data Warehouse, (b) an application that uses data stored in Azure data sources as well as data stored on premises (this would be called a hybrid application). In order to learn more about how Qlik Sense application connect to different data sources, visit the following link:

http://help.qlik.com/en-US/sense/2.2/Subsystems/Hub/Content/LoadData/connect-data-sources.htm?q=connectivity

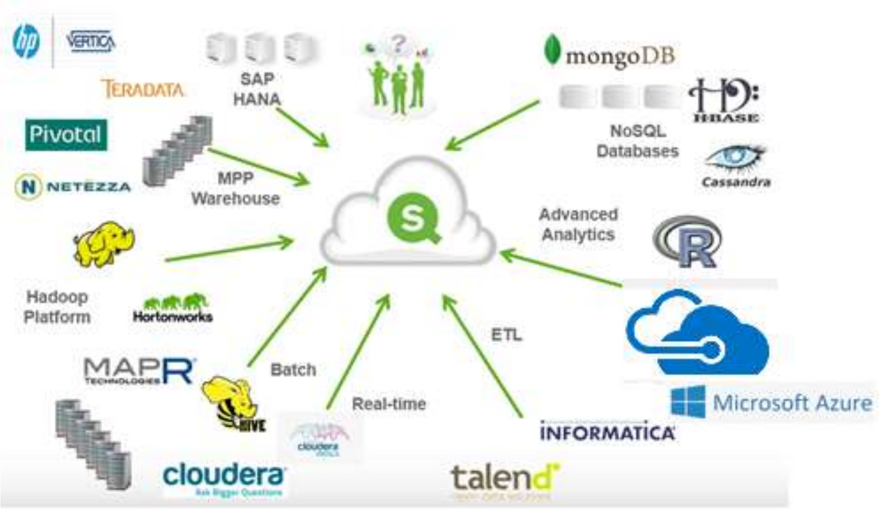

When Qlik Sense is deployed within Azure infrastructure, it is very common to see Qlik Sense applications leveraging Azure Data sources such as Azure SQL Database, Azure SQL Data Warehouse, etc. Azure SQL Data Warehouse is one of the fastest growing products within Azure. Azure SQL Data Warehouse is a cloud-based, scale-out database capable of processing massive volumes of data, both relational and non-relational. Built on our massively parallel processing (MPP) architecture, SQL Data Warehouse can handle enterprise workloads very efficiently. It combines the SQL Server relational database with Azure cloud scale-out capabilities. Being entirely elastic, Azure SQL Data Warehouse can increase, decrease, pause, or resume compute in seconds reducing the overall costs by scaling out CPU whenever needed and cutting back usage during non-peak times. Azure SQL Data Warehouse leverages the Azure platform so it is easy to deploy, seamlessly maintained, and fully fault tolerant because of automatic back-ups. As Organizations use Azure SQL Data Warehouse for their traditional Enterprise data warehousing needs, Big Data scenarios and SaaS offerings, Qlik Sense makes it very easy to consume the data that is stored and loaded in Azure SQL Data Warehouse and use it in combination with other data sources. To learn more about the integration of Qlik's platform with Azure SQL Data Warehouse, visit the following link:
https://azure.microsoft.com/en-gb/documentation/articles/sql-data-warehouse-partner-business-intelligence/

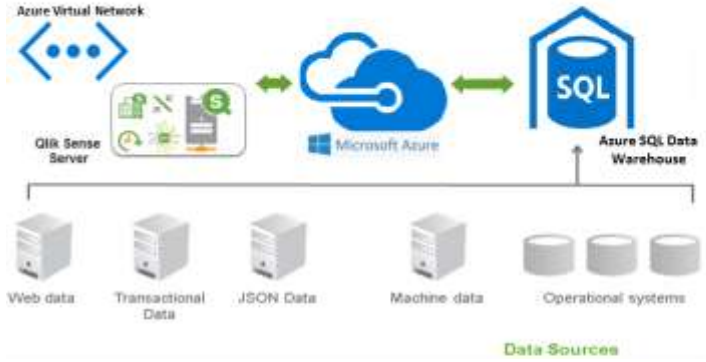

In the case that data source resides outside Azure, Qlik Sense could connect to them by leveraging a VPN connection (using an application gateway attached to the Azure Virtual Network that talks to a gateway located where the data source(s) in question reside via a VPN Tunnel). Azure ExpressRoute lets you extend on-premises networks into the Azure cloud over a dedicated private connection facilitated by a connectivity provider.

With ExpressRoute, connections to Azure cloud services, such as Office 365, and CRM Online can be established. Connectivity can be from an any-to-any (IP VPN) network, a point-to-point Ethernet network, or a virtual cross-connection through a connectivity provider at a co-location facility. ExpressRoute connections do not go over the public Internet. This allows ExpressRoute connections to offer more reliability, faster speeds, lower latencies, and higher security than typical connections over the Internet. To learn more about Azure ExpressRoute and the different options on how to configure the Azure VPN Gateway, visit the following link:

https://azure.microsoft.com/en-us/documentation/articles/vpn-gateway-topology/

As mentioned before in this section, Qlik Sense leverages native ODBC drivers when available. In the case of Azure SQL Data Warehouse, Qlik Sense establishes an ODBC connection in order to connect to Azure SQL Data Warehouse. Due to the distribution of data happening inside the Azure Platform; Qlik recommends users to follow Azure SQL Data Warehouse best practices to determine how to place organize data in an optimal way for query performance. The follow diagram represents high level, the architecture of an Azure SQL Data Warehouse and Qlik Sense environment where the SQL Data Warehouse Control Node manages and optimizes queries. It is the front end that interacts with all applications and connections. In Azure SQL Data Warehouse, the Control node is powered by SQL Database, and connecting to it looks and feels the same. Under the surface, the Control node coordinates all of the data movement and computation required to run parallel queries on the distributed data across the different Compute nodes. The Compute nodes serve as the power behind SQL Data Warehouse. They are SQL Databases that store data and process queries. When data is added, SQL Data Warehouse distributes the rows to the Compute nodes. The Compute nodes run the parallel queries on the data and after processing, they pass the results back to the Control node. To finish the query, the Control node aggregates the results and returns the final result. The data is stored in Azure Blob storage. When Compute nodes interact with the data, they write and read directly to and from blob storage. Since Azure storage expands transparently and limitlessly, SQL Data Warehouse can do the same. Since compute and storage are independent, SQL Data Warehouse can automatically scale storage separately from scaling compute, and vice-versa. Azure Blob storage is also fully fault tolerant, and streamlines the backup and restore process.

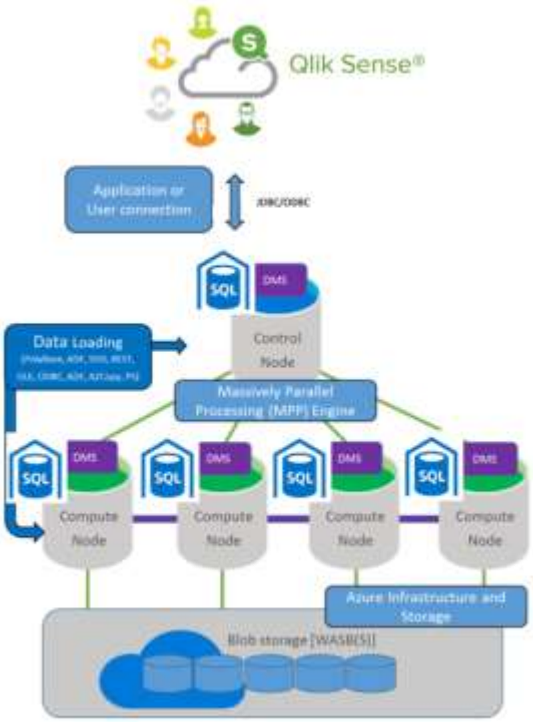

In order to establish ODBC connectivity from Qlik Sense Enterprise to Azure SQL Data Warehouse, first the Azure SQL ODBC driver needs to be downloaded and installed on the Qlik Sense Server. In order to learn more about the Azure SQL Data Warehouse ODBC client and to download it, visit the following links:

https://msdn.microsoft.com/library/jj730314.aspx (ODBC Driver Download)

https://azure.microsoft.com/en-gb/documentation/articles/sql-data-warehouse-connect-overview/ (Connectivity)

Note that the same Azure SQL ODBC driver supports connections to Azure SQL Database, Azure SQL Data Warehouse, SQL Server 2016, SQL Server 2014, SQL Server 2012, SQL Server 2008 R2, SQL Server 2008, and SQL Server 2005. Once the ODBC driver, has been installed, follow the next steps to create the ODBC DSN that will be used by Qlik Sense Server:

1. Start the “Data Sources (ODBC)” program in Windows
2. You should see a dialog box that looks like this (Drivers Tab):

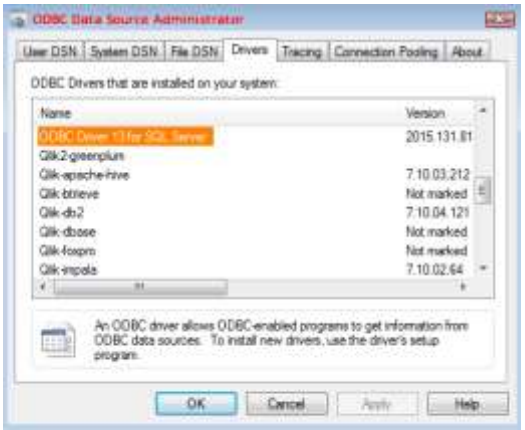

ODBC Drivers that are installed on your system:

Name	Version
ODBC Driver 13 for SQL Server	2015.131.81
Qlik2-greenplum	
Qlik-apache-hive	7.10.03.212
Qlik-btrieve	Not marked
Qlik-db2	7.10.04.121
Qlik-dbase	Not marked
Qlik-foxpro	Not marked
Qlik-impala	7.10.02.64

An ODBC driver allows ODBC-enabled programs to get information from ODBC data sources. To install new drivers, use the driver's setup program.

3. In the System DSN tab, click “Add” and a new window that looks like this will pop up. Select “ODBC Driver 13 for SQL Server” and click “Finish”. This will start the process to create a Data Source to SQL Server

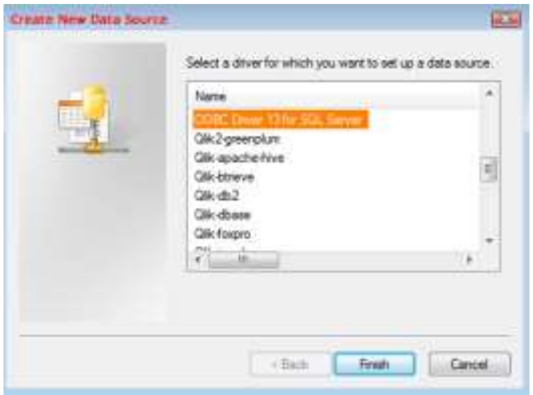

Select a driver for which you want to set up a data source.

Name
ODBC Driver 13 for SQL Server
Qlik2-greenplum
Qlik-apache-hive
Qlik-btrieve
Qlik-db2
Qlik-dbase
Qlik-foxpro

4. On the next screen enter a name for the Data Source and the name of the Azure SQL Data Warehouse server. During the configuration of the Azure SQL Data Warehouse server, a full

name is given to the database. This is the same name that should be used. The name will be in the form of: “XXXX.database.windows.net”. Click “Next”

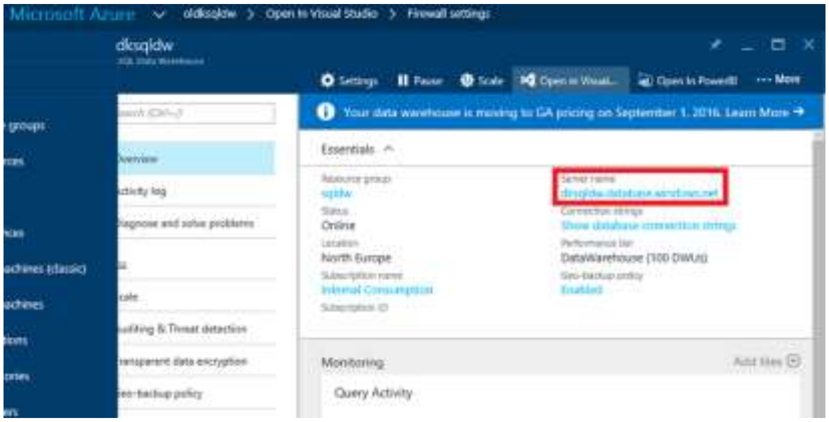

Your data warehouse is moving to GA pricing on September 1, 2016. Learn More →

Resource group: sqldw
Status: Online
Location: North Europe
Subscription name: Internal Consumption
Subscription ID
Server name: dlsqldw.database.windows.net
Connection strings: Show database connection strings
Performance tier: DataWarehouse (100 DWUs)
Geo-backup policy: Enabled

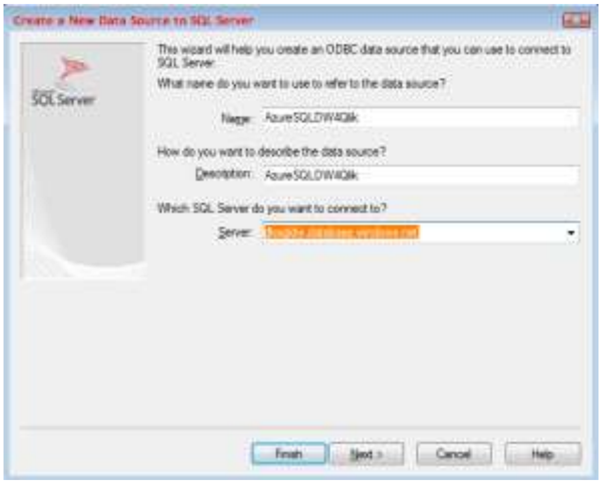

Create a New Data Source to SQL Server

This wizard will help you create an ODBC data source that you can use to connect to SQL Server.

What name do you want to use to refer to the data source?

Name: AzureSQLDW4Qlik

How do you want to describe the data source?

Description: AzureSQLDW4Qlik

Which SQL Server do you want to connect to?

Server: dlsqldw.database.windows.net

5. Depending on the authentication mechanism that is being used, select the desired option. For example, if a User ID and password have been created on the database to connect to it, use the SQL Server authentication option and enter the credentials. Click “Next”

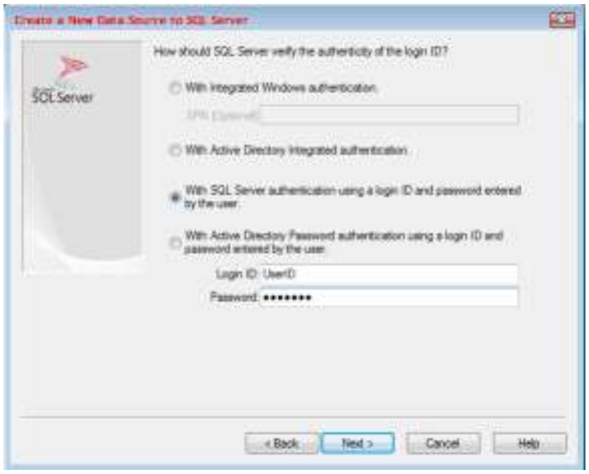

Create a New Data Source to SQL Server

How should SQL Server verify the authenticity of the login ID?

○ With Integrated Windows authentication.

SPN (Optional):

○ With Active Directory Integrated authentication

● With SQL Server authentication using a login ID and password entered by the user.

○ With Active Directory Password authentication using a login ID and password entered by the user

Login ID: UserID

Password: •••••••

6. Once authenticated, select the database in question and click “Next”

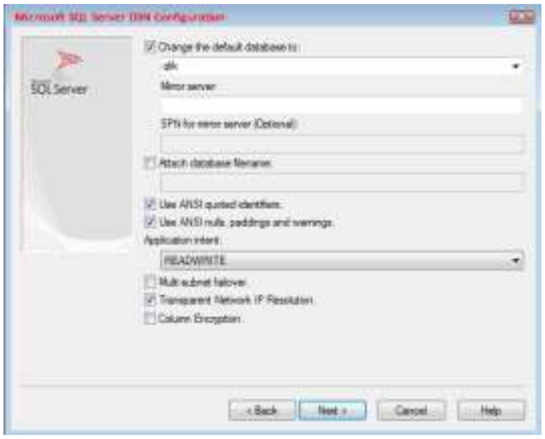

Microsoft SQL Server DSN Configuration

☑ Change the default database to:

qlik

Mirror server:

SPN for mirror server (Optional):

□ Attach database filename:

☑ Use ANSI quoted identifiers.

☑ Use ANSI nulls, paddings and warnings.

Application intent:

READWRITE

□ Multi-subnet failover.

☑ Transparent Network IP Resolution.

□ Column Encryption.

7. Select the preferred language and then click “Finish”

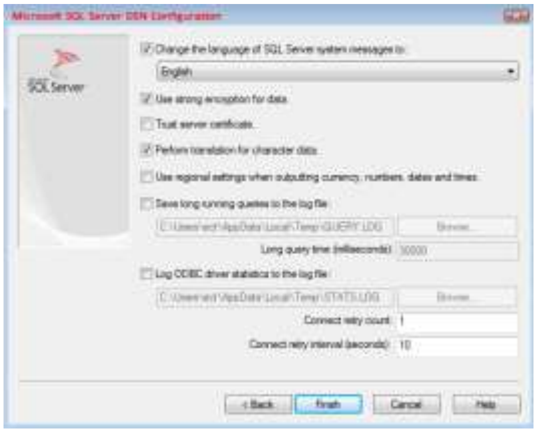

Microsoft SQL Server DSN Configuration

☑ Change the language of SQL Server system messages to:

English

☑ Use strong encryption for data

□ Trust server certificate

☑ Perform translation for character data

□ Use regional settings when outputting currency, numbers, dates and times

□ Save long running queries to the log file:

C:\Users\ect\AppData\Local\Temp\QUERY.LOG

Long query time (milliseconds): 30000

□ Log ODBC driver statistics to the log file:

C:\Users\ect\AppData\Local\Temp\STATS.LOG

Connect retry count: 1

Connect retry interval (seconds): 10

8. Make sure to test the connectivity to the database by clicking “Test Data Source”

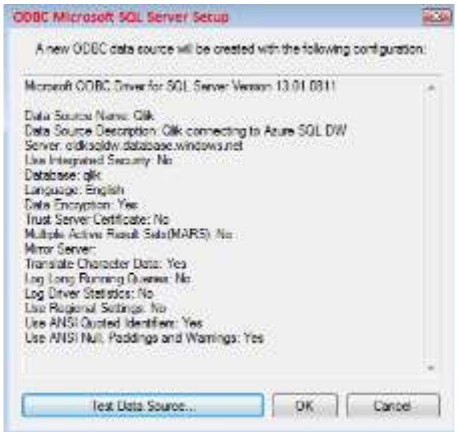

ODBC Microsoft SQL Server Setup

A new ODBC data source will be created with the following configuration:

```
Microsoft ODBC Driver for SQL Server Version 13.01.0811

Data Source Name: Qlik
Data Source Description: Qlik connecting to Azure SQL DW
Server: qlikqsqldw.database.windows.net
Use Integrated Security: No
Database: qlik
Language: English
Data Encryption: Yes
Trust Server Certificate: No
Multiple Active Result Sets(MARS): No
Mirror Server:
Translate Character Data: Yes
Log Long Running Queries: No
Log Driver Statistics: No
Use Regional Settings: No
Use ANSI Quoted Identifiers: Yes
Use ANSI Null, Paddings and Warnings: Yes
```

9. Confirm that the connection completes successfully and click “Ok”

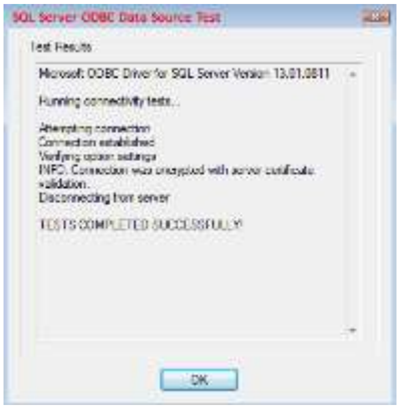

```
Microsoft ODBC Driver for SQL Server Version 13.01.0811

Running connectivity tests...

Attempting connection
Connection established
Verifying option settings
INFO: Connection was encrypted with server certificate validation.
Disconnecting from server

TESTS COMPLETED SUCCESSFULLY!
```

10. From Qlik Sense Enterprise, create a new application and connect to Azure SQL Data Warehouse using the ODBC DSN that has been created

While this document does not cover every single Azure data source, each Azure data source provides native connectivity that would be leveraged in a similar way by Qlik Sense Server. Depending on the source in question, Azure provides a number of very good reference on how to tune the environment to achieve the best performance possible from Qlik Sense Server. For example, below are some references that can be used for Azure SQL Data Warehouse. Typically, Qlik recommends to leverage the processing power of Azure SQL Data Warehouse to execute complex SQL queries (i.e. multiple sub-selects, complex join, etc.) and send the resulting data set to Qlik applications via ODBC. Here are some pointers for designing Azure SQL Data Warehouse to work well with Qlik Sense Server:

- Manage databases in Azure SQL Data Warehouse
  https://azure.microsoft.com/en-us/documentation/articles/sql-data-warehouse-overview-manage/

- Load data into Azure SQL Data Warehouse
  https://azure.microsoft.com/en-us/documentation/articles/sql-data-warehouse-overview-load/

- Leverage other services with SQL Data Warehouse
  https://azure.microsoft.com/en-us/documentation/articles/sql-data-warehouse-overview-integrate/

How to get started creating a new application

Creating new applications in Qlik Sense is a very straight forward process. Qlik Sense technology is based on the simple concept of applications. An application in the context of Qlik Sense is a file which extension is QVF. A QVF file (or application) is the container that captures the things required for users to be able to analyze, drill, visualize, explore and discover their data. In essence you can thing of a QVF as the unit within the Qlik Sense environment that uses interact with in order to analyze and discover the data and that it knows where the data sources are located and how they are being used.

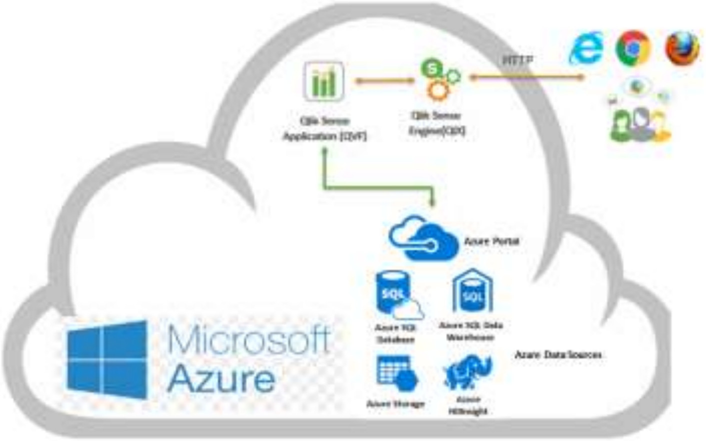

The QVF file will contain these three main things:

a. **Visualizations:** this layer of the application will be conformed by the following elements: sheets that contain visualizations, bookmarks for specific set of selections and interactive stories with live data.

Visual interface
Sheet with visualizations
Bookmarks
Stories with snapshots

b. **Data**: when a Qlik Sense application connects to a specific data source, only each one of the distinct data points within each field, table, and data source are captured as part of the data load doing this in a very efficient way. Based on the data structure of the different data sources, the application will be automatically an associative data model that enables the discovery experience with the application. Based on the associative data model, different measures and dimensions can be optionally defined with the app and leveraged throughout various visualizations and analytics. In order to learn more about the Qlik Sense associative engine as well as the data loading process, visit the following links:

http://help.qlik.com/en-US/sense/2.2/Subsystems/Hub/Content/Selections/associative-selection-model.htm?q=associative

http://help.qlik.com/en-US/sense/2.2/Subsystems/Hub/Content/Scripting/introduction-data-modeling.htm?q=data loading

c. **Script**: As the data is mapped into an application and an associative data model is automatically created, then the application also captures all the business rules, connectivity information and rationale automatically as part of the application's script. Such script can be manually modified if wanted.

Foundation
Data model
Script
Measures
Dimensions
Fx

There are many free sources that can be leveraged to learn more about how to build a new application in Qlik Sense from scratch. Here are some examples coming from three sources: Qlik Community site (prior registration on the site is required), Qlik Help site that is publically available and Qlik's You Tube channel. The first three links are strongly recommended to get initiated.

- https://community.qlik.com/docs/DOC-7155
- https://community.qlik.com/docs/DOC-7156
- https://www.youtube.com/watch?v=i5-6v3NqarE
- http://help.qlik.com/en-US/sense/2.2/Subsystems/Hub/Content/Visualizations/create-apps-visualizations.htm
- https://www.youtube.com/watch?v=WW6tNTH2H5o&list=PLW1uf5CQ_gSrO7bIgP1xF0HowDXnG1S6l&index=3

Additionally, there are more resources available at Qlik's site under the Training section:
http://www.qlik.com/services/training/qlik-sense-training

Reference Architecture

Whether Qlik Sense is deployed to Azure manually or by leveraging Azure Marketplace, the following components need to be deployed and configured as part of the reference architecture:

- An Azure Virtual Network: Azure Virtual Network lets you provision a logically isolated section of the Azure cloud where you can launch Azure resources in a virtual network that you define. You have complete control over your virtual networking environment, including selection of your own IP address range, creation of subnets, and configuration of route tables and network gateways. In order to learn more Azure Virtual Network, visit the following link: https://azure.microsoft.com/en-us/documentation/services/virtual-network/
- **Subnet**: at least one subnet (either public or private) within the Virtual Network. In the case of just one, it would be considered the default one
- **Virtual Machine:** A Windows 2012 Server Virtual Machine instance deployed in the default subnet for installing and configuring Qlik Sense Server
- **Resource Group/Manager:** Azure Resource Manager groups all the resources within an environment/solution to make it easier to deploy, update or delete all of them in a single, coordinated operation. A Resource Group is the container that holds related resources for an application. The resource group can include all of the resources for an application, or only those resources that you group together. To learn more about Azure Resource Manager and Resource Groups visit the following links:
- https://azure.microsoft.com/en-us/documentation/articles/resource-group-overview/#resource-groups
- **Network Security groups:** a list of Access Control List (ACL) rules that allow or deny network traffic to the Virtual Machine instances in a Virtual Network. They can be associated with either subnets or individual Virtual Machine instances within that subnet. When they are associated with a subnet, the

ACL rules apply to all the Virtual Machine instances in that subnet. In addition, traffic to an individual Virtual Machine can be restricted further by associating a Network Security Group directly to a Virtual Machine. To learn more about Azure Network Security Groups, visit the following site:

https://azure.microsoft.com/en-us/documentation/articles/virtual-networks-nsg/

- **Azure Active Directory and Identity Management:** depending on the expected administration of the environment, integration with Azure Active Directory and Identity Management may be needed in order to manage fine-grained permissions for access to various Azure services involved in the deployment process. To learn more about Azure Identity Management, visit the following site:

  https://azure.microsoft.com/en-us/documentation/articles/fundamentals-identity/

- **Qlik Sense:** a Qlik Sense node (cluster in the case of the multimode deployment. Multi-node deployments are outside the scope of this document) configured based on compute and storage requirements

As stated earlier, the scope of this document is for a Qlik Sense single node deployment within a single region within Azure. Hence, one option is to launch the Qlik Sense node within a public subnet providing direct Internet access. Another option is to deploy the Qlik Sense node in a private subnet. The diagram below represents a reference architecture where Qlik Sense is within its own private subnet and various Azure data sources are leveraged as part of the environment running within their own private subnet as well.

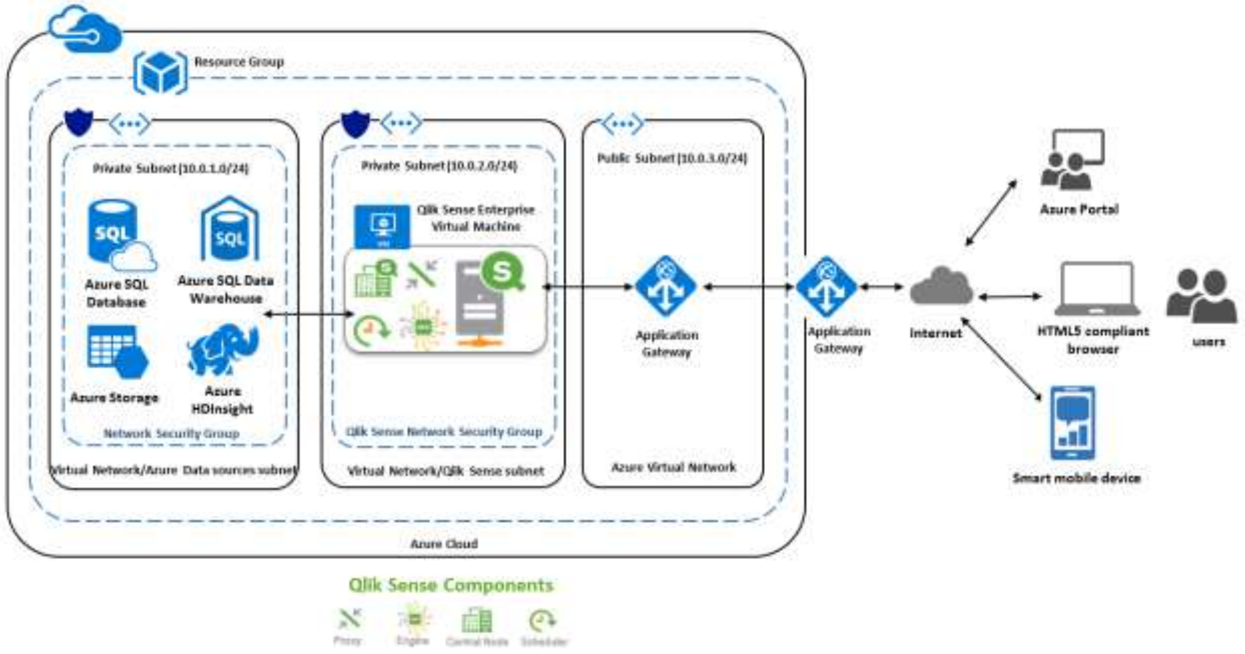

Storage and Access Patterns

As mentioned earlier in the section that talks about the installation process, part of the process of creating the Azure environment is to allocate enough storage space for the Qlik Sense node. Qlik Sense Server uses storage for a number of different things such as storing applications (QVF files), storing metadata information, storing logging information, logging users session information, etc. There are three main ways to manage the storage

1. **Standard and Premium Storage**: part of the configuration process of the Virtual Machine instance that is used for the Qlik Sense server is to determine the amount of storage use within the instance before having to add extra volumes for further storage. All Azure virtual machines have at least two disks – a Windows operating system disk (in the case of a Windows VM) and a temporary disk. The operating system disk is created from an image, and both the operating system disk and the image are actually virtual hard disks (VHDs) stored in an Azure storage account. Virtual machines also can have one or more data disks that are also stored as VHDs (to learn more about Virtual Machine storage, visit: https://azure.microsoft.com/en-us/documentation/articles/virtual-machines-windows-about-disks-vhds/) Microsoft Azure offers two types of durable storage: Premium Storage and Standard Storage. Premium Storage stores data on the latest technology Solid State Drives (SSDs) whereas Standard Storage stores data on Hard Disk Drives (HDDs). In order to leverage the storage within a Virtual Machine, a Storage account has to first be created. To learn more about SSD storage, visit the following page: https://azure.microsoft.com/en-us/blog/introducing-premium-storage-high-performance-storage-for-azure-virtual-machine-workloads/

Leveraging premium storage provides plenty of storage space (up to 35 TB - https://azure.microsoft.com/en-us/documentation/articles/storage-premium-storage/).

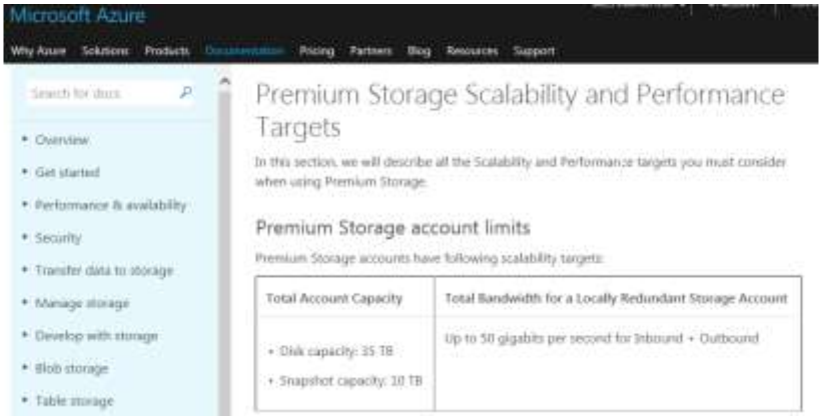

Premium Storage Scalability and Performance Targets

In this section, we will describe all the Scalability and Performance targets you must consider when using Premium Storage.

Premium Storage account limits

Premium Storage accounts have following scalability targets:

Total Account Capacity	Total Bandwidth for a Locally Redundant Storage Account
• Disk capacity: 35 TB • Snapshot capacity: 10 TB	Up to 50 gigabits per second for Inbound + Outbound

Once the storage space needed is determine (based on amount of raw data, taxonomy of the data, number of applications, expected size of applications, complexity of applications, concurrency, etc.), the storage can be allocated accordingly by choosing the right machine type. To learn more about the storage that comes with various different machines types, visit the following link: https://azure.microsoft.com/en-us/documentation/articles/virtual-machines-windows-sizes/

2. **Leverage additional Standard and/or Premium Storage**: You can add data disks to a virtual machine at any time, by attaching the disk to the virtual machine. You can use a VHD that you've uploaded or copied to your storage account, or one that Azure creates for you. Attaching a data disk associates the VHD file from your storage account with the virtual machine, by placing a 'lease' on the VHD so it can't be deleted from storage while it's still attached. The VHDs used in Azure are .vhd files stored as page blobs in a standard or premium storage account in Azure. Azure supports the fixed disk VHD format. The fixed format lays the logical disk out linearly within the file, so that disk offset X is stored at blob offset X. A small footer at the end of the blob describes the properties of the VHD. All .vhd files in Azure that you want to use as a source to create disks or images are read-only. When you create a disk or image, Azure makes copies of the .vhd files. These copies can be read-only or read-and-write, depending on how you use the

VHD. To learn more about best practices on how to attach VHD to a Virtual Machine, visit the following link:

https://azure.microsoft.com/en-us/documentation/articles/virtual-machines-windows-about-disks-vhds/

3. **Leverage Azure Storage**: Azure Storage is massively scalable, so you can store and process hundreds of terabytes of data to support the big data scenarios required by scientific, financial analysis, and media applications. Or you can store the small amounts of data required for a small business website. Azure Storage is elastic, so you can design applications for a large global audience, and scale those applications as needed - both in terms of the amount of data stored and the number of requests made against it. You pay only for what you use, and only when you use it. Azure Storage supports clients using Windows operating systems. Given that Qlik Sense is a Windows based platform a general-purpose storage account can give access to Azure Storage services such as Tables, Queues, Files, Blobs and Azure virtual machine disks under a single account. This type of storage account has two performance tiers which include a premium storage performance tier which supports Azure Virtual Machine disks. This option could be used for Qlik Sense purposes. To learn more about Azure Storage, visit the following link:

https://azure.microsoft.com/en-us/documentation/articles/storage-introduction/

It is important that as part of the planning process, the expected needed storage is determine ahead of time so then the Azure environment can be configured accordingly from the beginning and avoid storage issues down the road.

Security and Authentication

After a standard Qlik Sense installation, all security and authentication settings are managed directly from the Qlik Sense Management Console. As far as the different components of Qlik Sense Serve, it is the Qlik Sense Proxy Service (QPS) that includes a module that handles authentication of Microsoft Windows users. Other authentication methods are available, and it is also possible to implement customized solutions for different authentication scenarios. Here are some of the things that can be managed directly from the Qlik Sense Management Console (QMC):

- Admin roles to grant QMC users administrator level access to various sections
- Proxy certificate for communication between the web browser and the proxy component
- Virtual proxies to allow different modules based on the URI to be used to access Qlik Sense environment
- Custom properties to allow using your own values in security rules
- Access control and security rules to grant users access to Qlik Sense resources

In order to learn more about the different security configuration options, visit the following link:
http://help.qlik.com/en-US/sense/2.2/Subsystems/ManagementConsole/Content/configure-security.htm

There are different authentication options that can be leveraged to grant users access to different applications and functionality based on the user's identity as well as preventing unauthorized users to have access to such applications and functionality. Authentication is often used in conjunction with a single sign-on (SSO) system that supplies a reverse proxy or filter for authentication of the user. In order to learn more about the different configuration settings involved in the different authentication mechanisms, visit the following link:

http://help.qlik.com/en-US/sense/2.2/Subsystems/ManagementConsole/Content/authentication-methods.htm

Here are some of the most frequently options used for authentication:

a. **NTLM/Kerberos:** After a default installation of Qlik Sense Server, the Qlik Sense Proxy Service (QPS) includes a module that handles authentication of Microsoft Windows users. The module supports the use of Kerberos, NTLM, and basic authentication. The default authentication module requires that the proxy that handles the authentication is part of the Microsoft Windows domain in question.

b. **Security Assertion Markup Language (SAML):** Qlik Sense Server can leverage SAML to enable single-sign-on (SSO) and thereby minimizes the number of times a user has to log on to cloud applications and websites. For example, the following series of videos explain how to set up Qlik Sense Server SAML with Active Directory Federation Services:

   https://www.youtube.com/playlist?list=PLRWKt1LuYZ7m-zu9x-sQXnQP94UgyEi2G

This approach leverages the following entities in the authentication process:

- The user

- The identity provider (IdP) which is used for authentication. When the identity provider has asserted the user identity, the service provider can give the user access to their services. Because the IdP has enabled SSO, the user can access several service provider sites and applications without having to log into each site. In the authentication process, Qlik Sense Server plays the role of a service provider. When a user logs into Qlik Sense Server, the login is transferred to the identity provider that handles the actual SSO authentication.

- The service provider (SP)

c. **Anonymous authentication:** You can allow users to access Qlik Sense without supplying the user identity and credentials. This is done by editing the virtual proxy property **Anonymous access mode**. There are various levels of anonymous use; see the descriptions in the procedure below. To learn more about how to set up this type of authentication, visit the following link:

   http://help.qlik.com/en-US/sense/2.2/Subsystems/ManagementConsole/Content/anonymous-authentication.htm

d. **Session/Ticket API:** Once the user's credentials have been identified, it is necessary to transfer the user to Qlik Sense. In order to do so, Qlik Sense Enterprise can leverage its own Session and Ticket API. For example, one option is to bring the user and user's properties with a one-time ticket. Another option is to use an external module that can transfer web sessions that identify the user and the user's properties to Qlik Sense. In order to learn more about these options, visit the following links:

   - http://help.qlik.com/en-US/sense/3.0/Subsystems/PlanningQlikSenseDeployments/Content/Server/Server-Security-Authentication.htm

   - http://help.qlik.com/en-US/sense/3.0/Subsystems/PlanningQlikSenseDeployments/Content/Server/Server-Security-

Authentication-Solutions-Session-Solution.htm

- http://help.qlik.com/en-US/sense/3.0/Subsystems/PlanningQlikSenseDeployments/Content/Server/Server-Security-Authentication-Solutions-Ticket-Solution.htm

While setting up security and authentication on the Qlik Sense Server side, Azure Active Directory can be leveraged. Azure Active Directory (Azure AD) is Microsoft's multi-tenant cloud based directory and identity management service. For IT Admins, Azure AD provides an affordable, easy to use solution to give employees and business partners single sign-on (SSO) access to thousands of cloud SaaS Applications like Office365, Salesforce.com, Concur, etc. Azure AD also includes a full suite of identity management capabilities including multi-factor authentication, device registration, self-service password management, self-service group management, privileged account management, role based access control, application usage monitoring, rich auditing and security monitoring and alerting. These capabilities can help secure cloud based applications, streamline IT processes, cut costs and help assure corporate compliance goals are met. Additionally, with just four clicks, Azure AD can be integrated with an existing Windows Server Active Directory, giving organizations the ability to leverage their existing on-premises identity investments to manage access to cloud based SaaS applications.

Azure AD makes it easy to setup and run Microsoft Active Directory (AD) in the Azure cloud, or connect your Azure resources with an existing on-premises Microsoft Active Directory. Azure AD makes connecting to an on-premises directory and leveraging the connection while using existing corporate credentials to grant access to Azure resources easy. This option leverages Microsoft Security Support Provider Interface (SSPI) to read the Windows user name and password. Upon successful reading of the user's credentials, the experience and process to gain access is similar to a single-sign-on experience. In the case that there are multiple nodes in the Qlik Sense Server environment, all nodes need to be part of the same domain.

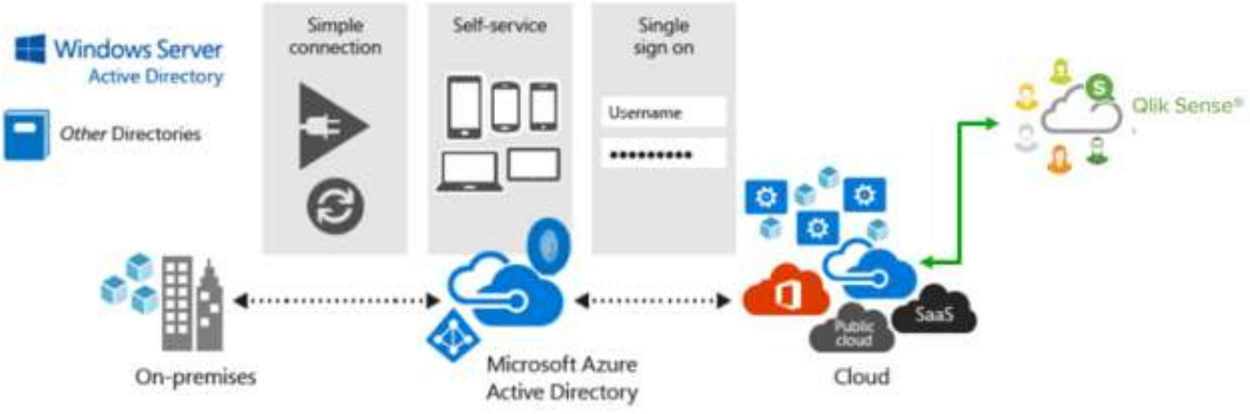

To learn more about Azure AD, visit the following the site:
https://azure.microsoft.com/en-us/documentation/articles/active-directory-whatis/

Operations

Once a Qlik Sense Server environment has been configured, there are a number of tasks and activities that typically take place right away as soon as the users start leveraging the environment. Aside from creating new applications and establish connectivity to various data sources (which has been described earlier in the document), this section provides and overview and description of such activities:

- Using the Qlik Management Console
- Importing/exporting applications
- Creating new Streams
- Scheduling tasks
- Publishing applications

Using the Qlik Management Console

Before expanding on such topics, having knowledge about the Qlik Sense Management Console is highly recommended. In order to learn more about some of the basic concepts, visit the following link that covers: (a) an introduction and licensing of a Qlik Sense site, (b) User Directory Connectors, (c) Token Allocation, (d) Streams, Apps and setting up tasks on the server, (e) Monitoring and Auditing the server, (f) Custom properties and security rules and (g) Virtual proxies, visit the following link:

https://www.youtube.com/playlist?list=PLW1uf5CQ_gSpUIEWu0-0TzzEaNVQo346i

Additionally, there are other great sources of reference on various topics from overviews, to Qlik Connectors to advanced topics for developers. For users that are new to Qlik, it is highly recommendable to go through some of the links below to learn more about Qlik's technology:

- https://www.youtube.com/channel/UCFxZPr8pHfZS0n3jxx74rpA
- https://community.qlik.com/docs/DOC-6932
- https://community.qlik.com/docs/DOC-7144

Importing and exporting applications

When a new Qlik Sense Server environment is first set up or throughout its life, it is quite common to bring applications from other environments through an import process. Given that Qlik Sense applications are self-contained, it is quite easy to import and export applications from one environment to another one. The following links lists the steps that need to be followed in order to import and export applications from one environment to another one. Please notice that if the app contains ODBC data connections, you must create new connections in the target environment, or use the ones that already exist at the new site. You also need to make sure that the related ODBC data sources exist on the target environment. The ODBC data sources need to be named and configured identically, and point to the same databases or files.

- **Import**: http://help.qlik.com/en-US/sense/2.2/Subsystems/ManagementConsole/Content/import-apps.htm?q=import
- **Export**: http://help.qlik.com/en-US/sense/2.2/Subsystems/ManagementConsole/Content/export-apps.htm?q=export

Creating New Streams

A stream enables users to read and/or publish apps, sheets, and stories. Users who have publish access to a stream, create the content for that specific stream. The stream access pattern in a Qlik Sense site is determined by the security rules for each stream. By default, Qlik Sense includes two streams: Everyone and Monitoring apps. An app can be published to only one stream. To publish an app to another stream, the app must first be duplicated and then published to the other stream. All authenticated users have read and publish rights to the **Everyone** stream and all anonymous users read-only rights. Three of the predefined admin roles (RootAdmin, ContentAdmin, and SecurityAdmin), have read and publish rights to the Monitoring apps stream. To learn how to create a stream, visit the following link: http://help.qlik.com/en-US/sense/2.2/Subsystems/ManagementConsole/Content/create-streams.htm?q=streams

Once a stream has been created, it is easy to edit and modify its configuration. To learn how to edit a stream, visit the following link: http://help.qlik.com/en-US/sense/2.2/Subsystems/ManagementConsole/Content/edit-streams.htm

Deleting streams is also possible, to learn how to do so, visit the following link: http://help.qlik.com/en-US/sense/2.2/Subsystems/ManagementConsole/Content/delete-streams.htm

Once a stream has been created, the access rights to such stream can be managed through the Qlik Sense Management Console. To learn how to do so, visit the following link: http://help.qlik.com/en-US/sense/2.2/Subsystems/ManagementConsole/Content/create-access-rights-for-streams.htm

The following video summarizes the concepts about streams described above: https://www.youtube.com/watch?v=fM85UttVzrM&index=4&list=PLW1uf5CQ_gSpUIEWu0-0TzzEaNVQo346i

Schedule Tasks

Qlik Sense comes with a services called Qlik Sense Scheduler Service (QSS) that can be used to schedule tasks incorporating time and event based triggers. Typically, tasks are used to perform a wide variety of operations and can be chained together in just any pattern. The tasks are handled by the Qlik Sense Scheduler Service (QSS). There are two types of tasks:

- **Reload**: it fully reloads the data in an app from the source. Any old data is discarded.
- **User synchronization:** A user synchronization task imports the users and the users' information from a user directory. When you create a new instance of a user directory connector (UDC) a synchronization task with a scheduled trigger is created by the system.

In order to learn how to create, view, edit, delete, enable/disable, start and stop tasks, visit the following link: http://help.qlik.com/en-US/sense/2.2/Subsystems/ManagementConsole/Content/manage-tasks-and-triggers.htm

The execution of a task is initiated by a trigger or manually from the tasks overview page. You can create additional triggers to execute the task and there are two types of triggers:

- **Scheduled**: they can be applied to both reload tasks and user synchronization tasks
- **Task Event**: they can only be applied to reload tasks.

The following video (second half) exemplifies how a task can be created and set up in the Qlik Sense Management Console: https://www.youtube.com/watch?v=fM85UttVzrM&index=4&list=PLW1uf5CQ_gSpUIEWu0-0TzzEaNVQo346i

Publishing Applications

Qlik Sense applications that include sheets and stories can be published to streams to make content available to other users that have access to the stream in question. When an application is first published, the sheets and stories in it will be accessible by everybody that has access to the application itself. Additional sheets and stories can be published as a part of an app that is already published. Access Rights control which users are allowed to publish different kinds of content. To learn more about how to publish new sheets and stories and what you can do within an application that is already published vs. an application that has not been published yet, visit the following link:

http://help.qlik.com/en-US/sense/2.2/Subsystems/Hub/Content/Publishing/Publish.htm?q=publishing

You can create an application in Qlik Sense Server Hub, if you have the appropriate access rights. Alternatively, in order to publish an application that was created with the Qlik Sense Desktop client, the application must be first imported by using the Qlik Sense Management Console. The security rules applied to the application, stream, or user, determine who can access the content and what the user is allowed to do. The application is locked when published. Content can be added to a published app through the Qlik Sense hub in a server deployment, but content that was published with the original app cannot be edited. The following link lists the steps on how to publish and app that has not been published yet: http://help.qlik.com/en-US/sense/2.2/Subsystems/ManagementConsole/Content/publish-apps.htm?q=publishing

Designs for Deploying Qlik Sense on Azure

Qlik Sense Server is architected to scale up (vertically) by adding more resources (RAM and CPU cores) to the server node as well as to scale out (horizontally) by adding more nodes to an existing cluster. In the case of a deployment on Azure, this topic becomes especially relevant given how easy it is to add more resources to an existing deployment in order to scale massively. In the simplest scenario with a single node environment, the different Qlik Sense services run within the same Azure Virtual Machine instance and/or Resource Group:

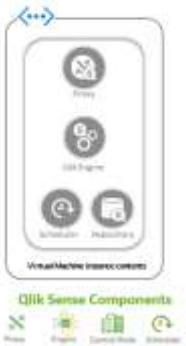

Multi-node deployments

Qlik Sense architecture is extremely flexible. Hence, when a multi-node environment is being considered, different configurations are possible. The Qlik Sense Server components can be split in different ways depending on what the overall goal is. One common way to look at a multi-node environment is by considering the following layers: data access, application and web front-end. In such case, then the different Qlik Sense Server components are assigned to each one of the different layers so resources can be assigned to each layer independently depending on their overall usage.

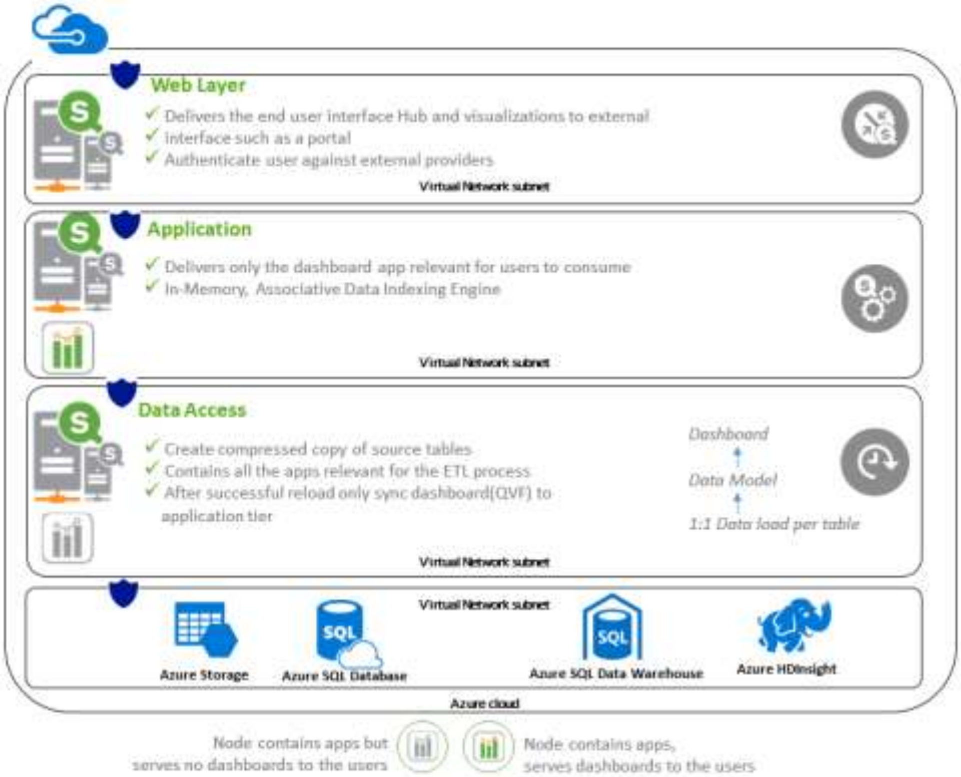
Web Layer
✓ Delivers the end user interface Hub and visualizations to external
✓ interface such as a portal
✓ Authenticate user against external providers
Virtual Network subnet
Application
✓ Delivers only the dashboard app relevant for users to consume
✓ In-Memory, Associative Data Indexing Engine
Virtual Network subnet
Data Access
✓ Create compressed copy of source tables
✓ Contains all the apps relevant for the ETL process
✓ After successful reload only sync dashboard(QVF) to application tier
Dashboard
Data Model
1:1 Data load per table
Virtual Network subnet
Virtual Network subnet
Azure Storage
Azure SQL Database
Azure SQL Data Warehouse
Azure HDInsight
Azure cloud
Node contains apps but serves no dashboards to the users
Node contains apps, serves dashboards to the users

When considering then the different layers (front-end, application, management and back-end), there would be various possible configurations. For example on the image below, (1) represents a single node installation where all the services are running on the same node; (2) represents a two node installation where the front end is on one node and the backend services on another node; (5) represents a three layer environment where the front end runs on one node, the application layer runs on a two node cluster for the two Qlik Sense Servers and the backend layer runs on a two node cluster running the backend Qlik Sense services; (8) represents a shared file layer for both sites which requires the applications to be stored as many times as there are sites; etc.

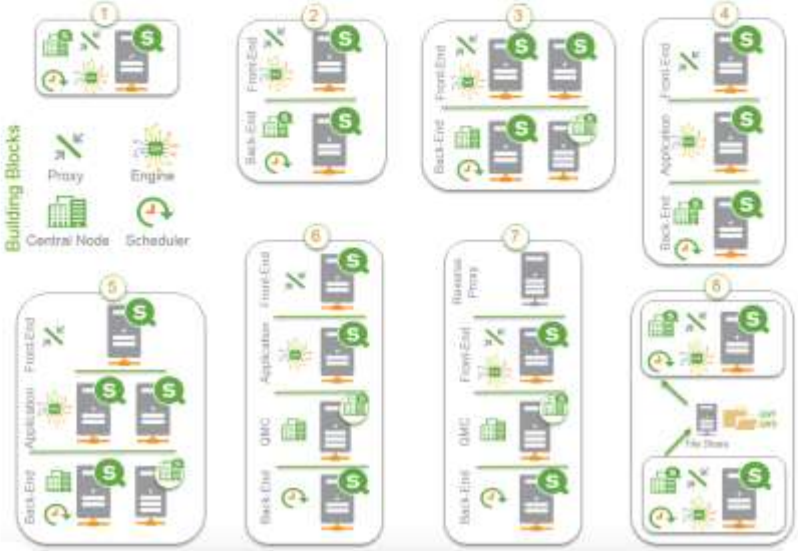

In a multi-node environment, there would be a node that would be considered the central node. All additional nodes, would be called RIM nodes. In order to learn more about some of the basic concepts of a Qlik Sense multimode environment, visit the following links:

- Guidelines for deploying multi-node sites
  http://help.qlik.com/en-US/sense/2.2/Subsystems/PlanningQlikSenseDeployments/Content/Server/Server-Deployment-Scenarios-Multi-Node-Site-Guidance.htm?q=multi node

- Deploying multi-node sites:
  http://help.qlik.com/en-US/sense/2.2/Subsystems/PlanningQlikSenseDeployments/Content/Server/Server-Deployment-Scenarios-Multi-Node-Site.htm?q=multi node

- Multi-node scenario: Production deployment
  http://help.qlik.com/en-US/sense/2.2/Subsystems/PlanningQlikSenseDeployments/Content/Server/Server-Deployment-Scenarios-Multi-Node-Site-Examples-Prod.htm?q=multi node

Load Balancing

When multiple nodes exist in the Qlik Sense Server environment, then a Load Balancer will be needed to provide the role of distributing users to servers (note that the Load Balancer must support WebSocket communication):

- **Load Balancing Qlik Sense Proxy Services:** The Qlik Sense Proxy Service (QPS) manages site authentication, session handling, and load balancing of Qlik Sense Engines. In this scenario multiple QPS can be deployed and configured to send all user connections to share engine servers. A Network Load Balancer ensures distribution is performed to both QPS based upon the configuration.

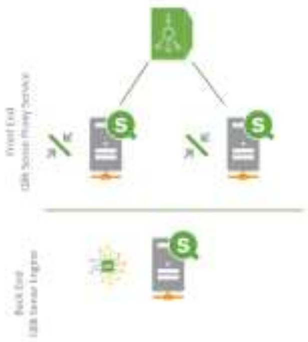

- **Load Balancing Qlik Sense Engines:** By default the Qlik Sense Proxy (QPS) provides Load Balancing for the first node created within a site (the central node). This action ensures the first QPS is able to send sessions to the first engine. When multiple engines and virtual proxies are installed, you must add load balancing to the new nodes and virtual proxies. If you create a node without configuring the virtual proxy, the node will never actually be used. If you create a new virtual proxy, you must configure it by adding load balancing and selecting which nodes that the virtual proxy can forward work to. The configuration of a Virtual Proxy provides the capability to add additional engines that are Load Balanced using round robin. It is the Qlik Sense Proxy Service that performs this function.

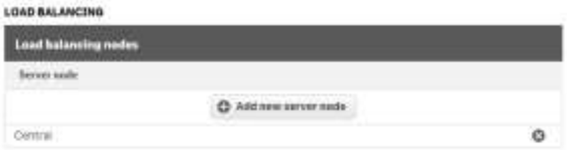

LOAD BALANCING

Load balancing nodes
Server node
Add new server node
Central

In addition, it is possible to utilize a third party Load Balancer by configuring in the Qlik Sense Management Console the Virtual Proxy to use a “Load balancing module base URI” under the Integration settings (see Appendix A for more information).

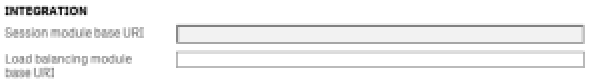

INTEGRATION

Session module base URI

Load balancing module base URI

High Availability and Disaster Recovery Strategy

Having a multi-node environment, allows to have a more robust strategy for High Availability scenarios as well as a Disaster Recovery strategy. Qlik Sense Server environments can be backed up and recovered easily. To learn more about how to back up and recover a Qlik Sense Server environment, visit the following link: http://help.qlik.com/en-US/sense/2.2/Subsystems/PlanningQlikSenseDeployments/Content/Server/Server-Backup-Restore.htm

As mentioned earlier, the scope of this deployment guide is a single node, single region deployment scenario. Under such scenario, High Availability and Disaster Recovery can be addressed by leveraging some of the services that Azure offers. All the information about the Qlik Sense environment is stored in the Virtual Machine Storage (under the Storage account). This includes the Qlik Sense Repository which snapshots are taken via Qlik Repository Snapshot Manager used as an alternative to or conjunction with Azure Recovery Services Vault. To learn more about the Qlik Repository Snapshot Manager, visit the following link: http://help.qlik.com/en-US/sense/2.2/Subsystems/PlanningQlikSenseDeployments/Content/Server/Server-Backup-Restore-Site-Using-RSM.htm?q=repository%20snapshot%20manager

The storage associated with the Virtual Machine can be replicated either by using geo-redundant storage and/or locally redundant storage. By default, Azure Recovery Services Vault has geo-redundant storage. Leave the option set to geo-redundant storage if this is your primary backup. Choose locally redundant storage if you want a cheaper option that isn't quite as durable. To learn more about Storage replication, visit the following link: https://azure.microsoft.com/en-us/documentation/articles/backup-azure-vms-first-look-arm/

Notice that a backup policy can be defined as part of this process. A backup policy defines a matrix of when the data snapshots are taken, and how long those snapshots are retained. When defining a policy for backing up a VM, you can trigger a backup job *once a day*. When you create a new policy, it is applied to the vault. Additionally, Virtual Machines can be backed up to a Backup Vault. To learn more about this topic, visit the following site: https://azure.microsoft.com/en-us/documentation/articles/backup-azure-vms/

When the volumes are backed up into Azure Recovery Services on scheduled basis, in parallel, a mirror passive environment of Qlik Sense Server (Disaster Recovery Environment or DR Environment) under a different Azure Resource Group is always ready to go. The DR environment has access to the latest Azure Recovery Services snapshots via a replication process within Azure. All users are redirected to the DR Environment via Virtual Network Gateway.

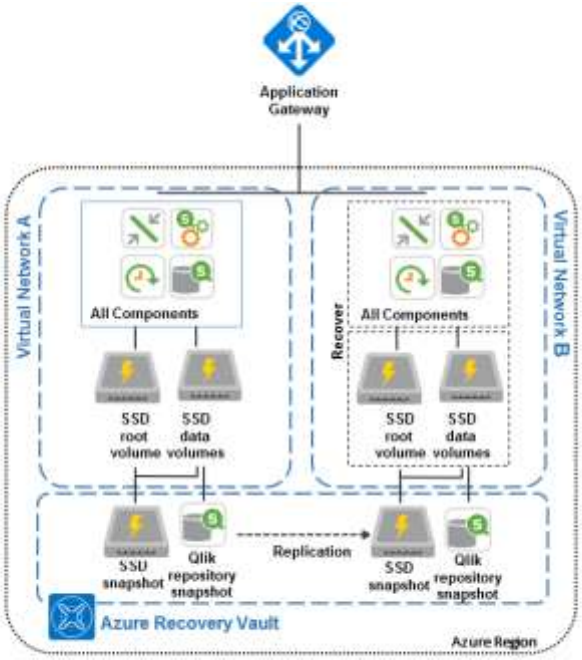

In the case of having a multi-node deployment of Qlik Sense Server, some changes are needed in order to have a High Availability and DR environment. As mentioned in the prior section, a Load Balancer will be needed and Azure Load Balancer does not support WebSockets and stickiness simultaneously. Hence, a third party Load Balancer would be needed. Alternatively, instead of a Load Balancer, an Application Gateway associated with one Qlik Sense Proxy could be used instead. In the event of Virtual Network failure, then the Application Gateway would have to be manually associated to the Qlik Sense Proxy on a different Virtual Network. In such case, Qlik Sense Proxy needs to be configured to load balance to engines in both Virtual Networks. Another change is the distribution of the Qlik Sense applications across the different Virtual Networks. It is recommended to distribute the Qlik Sense applications evenly for High Availability purposes. In the case of uneven distribution, Qlik Sense Proxies should load balance to engines in both Virtual Networks. In all cases, it is strongly recommended to use SSD storage, Azure Recovery Vault and the Qlik Repository Snapshot Manager to take the Repository snapshots. The following diagram depicts a multi-node Qlik Sense environment in Azure with High Availability and Disaster Recovery

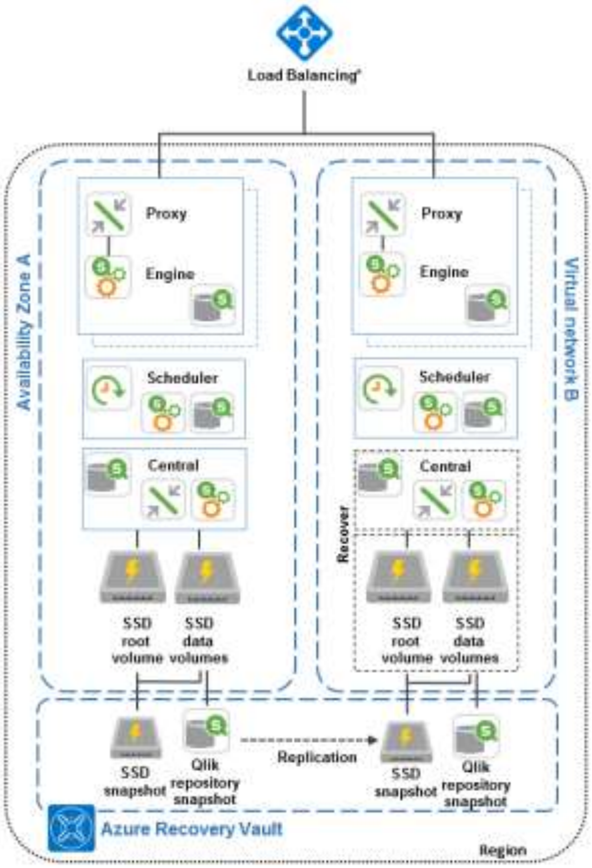

High-Performance Storage

Qlik Sense Server provides and Operations Monitor that helps understand the consumption and usage of the resources in the server node. Within the Operations Monitor, the Performance sheet displays the history of hardware utilization, active users, and active documents on the current node over a period selected by the user. In a multi-node environment, data comes from all nodes, unless specific nodes have been selected. The average and maximum usage is for all nodes combined or all selected nodes. The user can select on months, weeks, dates, and days of the week. Selections can also be made by hour and by ten-minute time period. In order to learn more about the Operations Monitor, visit the following link:

http://help.qlik.com/en-US/sense/2.2/Subsystems/Monitoring/Content/OperationsMonitor/Operations-monitor-app.htm?q=operations

Within the Operations Monitor, the performance tab contains charts and summary table can highlight periods of peak CPU and RAM usage and help identify concurrent events that might be contributing to the high usage. They can also help diagnose trends for concurrent users and apps that could contribute to periods of high activity that cause maximum utilization of resources reflected in RAM or CPU usage.

Qlik Sense Server is primarily RAM based, so high-performance storage does not play a critical role within the overall performance of the environment based on internal benchmarking tests. Hence, there is no specific recommendation for faster storage above and beyond the General Purpose SSD.

Conclusion

Deploying Qlik Sense Server on Azure Cloud platform enables Organizations to provide cloud based analytical solutions to their various stakeholders. The mechanics to deploy Qlik Sense Server on Azure are not too different from a traditional on premise deployment. This Deployment Guide summarizes considerations pertaining to architecture, configuration, security, storage, performance, management, monitoring that need to be considered. Following such guidelines will help to achieve optimal performance, availability and reliability. Additionally, this guide highlights how other Azure Services (Azure SQL databases, Azure SQL Data Warehouse, Azure HDInsight, Azure Storage, Azure Virtual Machines, etc.) can be leverage as part of the overall solution and ecosystem.

Contributors

The following individuals and organizations contributed to this document:

- Marcus Spitzmiller, Master Principal Enterprise Architect, Qlik
- Jeffrey Goldberg, Principal Enterprise Architect, Qlik
- Clint Carr, Principal Enterprise Architect, Qlik
- Raymond Neves, Senior Enterprise Architect, Qlik
- Thomas Hopp, Product Manager - Deployment and Performance, Qlik
- Jesús Centeno, Global Technology Alliances Manager, Qlik
- Kristofer Liljeblad, MS Senior Technical Evangelist & Cloud Solutions Architect, Microsoft

- Olivia Klose, MS Senior Technical Evangelist & Cloud Solutions Architect, Microsoft
- Rob Bagby, MS Senior Technical Evangelist & Cloud Solutions Architect, Microsoft
- Nicolas Caudron, MS Global Partner Business Evangelist & Alliance Manager, Microsoft
- Wesley Backelant, Cloud Data Solution Architect, Microsoft
- Bert Vandebroek, Cloud Solution Architect, Microsoft

Appendix A – Load Balancing Module

By adding a Load Balancing module base URI additional load balancing options are available (depending on the Load Balancer being used), for example the distribution of load could be based on RAM utilization or CPU utilization but the idea is to have load balancing capabilities on the front end for the proxy service as well as in the back end for the engine service as well.

There are a number of options in the market when it comes down to Load Balancers. For example, Azure Load Balancer automatically distributes incoming traffic across multiple Virtual Machines by creating a load balancer and registering instances with the load balancer in one or more Azure Resource Groups. The load balancer serves as a single point of contact for clients which increases the availability of Qlik Sense application. Virtual Machines can be added or removed from the load balancer as needed, without disrupting the overall flow of information. If a Virtual Machine fails, Azure Load Balancer automatically reroutes the traffic to the remaining running Virtual Machines. However, there is one major consideration to have in mind. Qlik Sense platform leverages WebSockets technology (a complete list of Qlik Sense system requirements is available here: http://help.qlik.com/en-US/sense/2.2/Subsystems/Installation/Content/Introduction/system-requirements.htm?q=websockets). Currently, Azure Load Balancer does not support WebSockets and stickiness simultaneously. There are some third party sources that propose how to work around such limitation without compromising the WebSockets by using TCP as a Load Balancer and Proxy protocol while creating a Load Balancer.

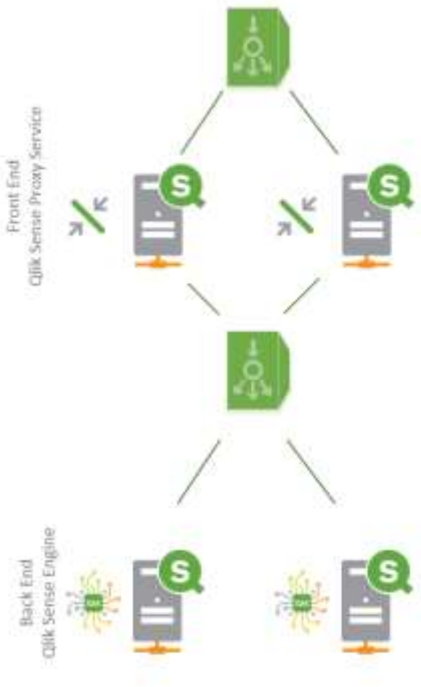

It is recommended to use a third party Load Balancer such as NGINIX or HAProxy. When considering using such Load Balancers, it is critical to ensure that they support and are configured to enable “sticky sessions”. In order to learn more about NGINIX, visit the following link: https://www.nginx.com/resources/admin-guide/load-balancer/

150 N. Radnor Chester Road
Suite E120
Radnor, PA 19087
Phone: +1 (888) 828-9768
Fax: +1 (610) 975-5987

qlik.com