



Qlik Sense on AWS

Deployment Guide 9.19.2016

Contents

1. [Abstract](#)
2. [Introduction](#)
 - a. [Qlik Sense: A Primer](#)
3. [Qlik on AWS - high level Architecture](#)
4. [Installation](#)
5. [Scalability and Sizing](#)
6. [Connectivity](#)
 - a. [How to get started](#)
 - b. [How to connect](#)
 - c. [How to build an app](#)
7. [Reference Architecture](#)
8. [Storage and Access Patterns](#)
9. [Security and Authentication](#)
10. [Operations](#)
 - a. [Using the Qlik Management Console](#)
 - b. [Importing/exporting applications](#)
 - c. [Creating new Streams](#)
 - d. [Scheduling tasks](#)
 - e. [Publishing applications](#)
11. [Designs for Deploying Qlik Sense on AWS](#)
 - a. [Multi-node deployments](#)
 - b. [Load Balancing](#)
 - c. [High Availability and Disaster Recovery Strategy](#)
12. [High Performance Storage](#)
13. [Conclusion](#)
14. [Contributors](#)
15. [Notes](#)

Abstract

Amazon Web Services (AWS) is a flexible, cost-effective, easy-to-use cloud computing platform. Qlik Sense is a platform that delivers visual analytics to where they deliver matter the most – the point at which decisions are made. Qlik's platform is widely deployed on AWS Cloud infrastructure to support analytics and visual analysis at various levels within organizations. Running Qlik Sense solutions on Amazon Elastic Compute Cloud (Amazon EC2) is a great option for organizations of any size that want to deliver SaaS (software as a service) solutions to internal and/or external users with cost effective, high-performing operations on any number of data sources. SaaS solutions based on Qlik's platform and delivered through AWS Cloud infrastructure enable individuals and teams to easily create and share applications with a wide number of users across multiple devices.

This Deployment Guide provides an overview on how to deploy and implement the Qlik Sense platform on the AWS Cloud infrastructure including architectural considerations and configuration steps that explain how Qlik Sense can leverage various AWS services such as Amazon Redshift, Amazon Elastic Compute Cloud (Amazon EC2), Amazon Virtual Private Cloud (Amazon VPC), AWS Directory Service, and Amazon Relational Database Services (RDS) to run scalable, secure and reliable analytical solutions.

Introduction

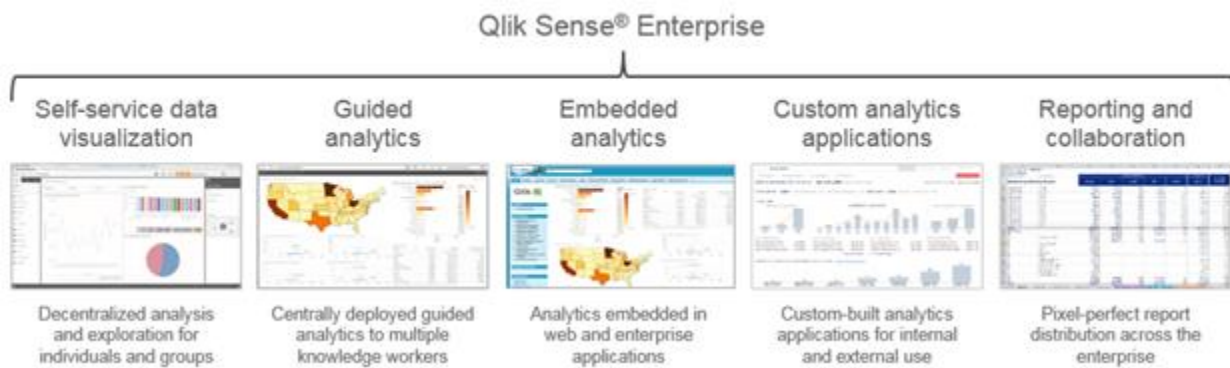
AWS provides an infrastructure with all the services and computing power needed to provide global solutions that can reliably perform as needed, regardless of unexpected changes in demand and concurrency across various levels of any organization. Qlik® Sense is a next-generation visual analytics platform 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. This gives 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, including:

- 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
- AWS Install

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 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 AWS – 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 AWS infrastructure so they can rapidly deploy new applications in a 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 AWS infrastructure on demand in minutes. For those customers, the AWS based approach is ideal for a SaaS solutions. This guide will help customers pursuing this approach so that they understand the different aspects of AWS 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 AWS implementations are:

I. Pre-requisites:

- Tight timeline to develop and deploy Qlik Sense applications
- Appropriate Licensing from Qlik is available
- Proactive support and stewardship from IT (i.e. Active Directory extract which can be used by an LDIF, Direct Connect, etc.)

II. Competencies

- Good understanding of AWS architecture and services (i.e. AWS EC2, AWS VPC, AWS Elastic IP, AWS S3, AWS RDS, AWS EMR, AWS Redshift, AWS Cloud formation, etc.) within the team
- Internal resources that can maintain and manage the Qlik and AWS environments are available
- Ability to setup a QA environment within AWS and appropriate time for testing is available

III. Constraints

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

AWS is an excellent platform for running many advanced data systems in the cloud. Characteristics of cloud computing solutions built on AWS, such as scalability, security and tight integration with IT resources 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. Amazon EC2 and Amazon VPC
- II. Qlik Sense Server
- III. Qlik Sense Configuration
- IV. Windows Server 2012
- V. Windows Server Active Directory and DNS
- VI. Amazon Redshift

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 AWS 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 AWS, see the 'Getting Started' section of the AWS documentation (which provides high level references and links to specifics areas of documentation) at: <http://docs.aws.amazon.com/gettingstarted/latest/awsgsg-intro/intro.html>

Installation

In order to install Qlik Sense on AWS, you should download Qlik Sense Enterprise from <http://www.qlik.com/download> (a Qlik account will be needed in order to access the download site). Before you launch an AWS EC2 instance (to learn more visit: <http://aws.amazon.com/ec2/>) where you can install Qlik Sense Enterprise, follow the next steps:

1. Create a security group for Qlik Sense (the name could be something generic such as Qlik Sense) with the rules in the following screenshot. Using the left navigation panel within the AWS Management Console, go to EC2 > Security Groups. Click "Create Security Group".

Name	Qlik Sense
Inbound Rules	HTTP, HTTPS, RDP, Custom TCP Rule (4244), Custom TCP Rule (4248)

If you are using VPC, the VPC ID can be entered here at this time, or it updated later once it is available. In the diagram, both HTTP and HTTPS are selected, but in most installations, you would choose one of them (typically only HTTPS).

Create Security Group

Security group name: Qlik Sense

Description: Qlik Sense

VPC: vpc-f032b495 (172.30.0.0/16)

Security group rules:

Inbound | Outbound

Type	Protocol	Port Range	Source
HTTP	TCP	80	Anywhere 0.0.0.0/0
HTTPS	TCP	443	Anywhere 0.0.0.0/0
RDP	TCP	3389	Anywhere 0.0.0.0/0
Custom TCP Rule	TCP	4244	Anywhere 0.0.0.0/0
Custom TCP Rule	TCP	4248	Anywhere 0.0.0.0/0

Add Rule

Cancel Create

Make sure to modify the Source Inbound Rule (Source) drop down menu to provide the correct source IPs. This can be a single IP or a range of addresses in CIDR notation. In essence, the security group requires necessary ports to be opened depending on the setup. For example, if the data source is on premise and requires the connection to, for example ODBC ports, those ports will need to be opened too.

2. **Create a key pair called “Qlik Sense”.** In the case that an existing key pair is available, it can be used instead. Using the left navigation panel within the AWS Management Console, go to EC2 > Key Pairs and click Create Key Pair. Save this file (Qlik Sense.pem) somewhere safe and make sure not to lose it. You will use this file to access the instance.


Create Key Pair

Key pair name: Qlik Sense

Cancel Create

3. **Launch a Windows Server 2012 R2 Base instance.** Using the AWS Management Console, go to EC2 > Instance and click “Launch Instance”. This will start a seven step process to create the instance.
 - a. **Choose AMI:** Using the Quick Start option use Windows Server 2012 R2 Base. Click ‘Next’ to continue. If for whatever reason you have already an AMI available that will be used for this purpose, typically, the following characteristics are desired as a minimum:
 - Recommended 4 cores / 16 GB minimum (this depends on the data volumes)
 - Clean O/S
 - .Net 4.5.2 installed and updates applied (default is 4.5 for Windows 2012)
 - IPv4 or IPv6 (dual stack)
 - Internet access from server (for license registration)

- Administrative rights and Remote Desktop access to the instance
- Service account with administrative rights on the server is available



- b. **Choose the Instance Type.** Qlik runs best on memory optimized instances, such as the R3 series. Choose one of the following 3 instances which are a recommended starting point:

- R3.xlarge is a good choice with 4 cores / 60 GB RAM
- R3.2xlarge is a good choice with 8 cores / 60 GB RAM
- R3.4xlarge is a good choice with 16 cores / 122 GB RAM


1. Choose AMI 2. Choose Instance Type 3. Configure Instance 4. Add Storage 5. Tag Instance 6. Configure Security Group 7. Review

Step 2: Choose an Instance Type

<input type="radio"/>	Memory optimized	r3.xlarge	4	30.5	1 x 80 (SSD)	Yes	Moderate
<input checked="" type="radio"/>	Memory optimized	r3.2xlarge	8	61	1 x 160 (SSD)	Yes	High
<input type="radio"/>	Memory optimized	r3.4xlarge	16	122	1 x 320 (SSD)	Yes	High
<input type="radio"/>	Memory optimized	r3.8xlarge	32	244	2 x 320 (SSD)	-	10 Gigabit

Notice that other larger instances are available and could be used. Some of the larger servers may require an email to AWS support first, depending on your organization's current usage and limits. If a larger server may be needed, please refer to the next section: "Sizing and Scalability".

Note: while it is better to go through all the steps, if the 'Review and Launch' button is clicked at this point, a warning about security groups will come up. Change from the default security group to "Qlik Sense" security group that was previously created.



- c. **Configure Instance:** this section will show the default configuration options. On this screen, the new instance could be assigned to the default VPC network or to a new one if needed. Set your VPIC, or if there are no changes, click 'Next' to continue.
- d. **Add Storage:** this section will show the default storage assigned to the new instance. Typically, the default disk type assigned (general purpose SSD) will be sufficient. For any scenario other than a simple proof of concept, we recommend you increase the size of the root volume to at least 100GB+. If

the instance is going to be used for a Production environment where large volumes of data will be handled with frequent reload of Qlik Sense applications, the speed of the disk may become a bottleneck. In this case you may need to increase the size that was allocated to the Qlik Sense system.

1. Choose AMI 2. Choose Instance Type 3. Configure Instance 4. Add Storage 5. Tag Instance 6. Configure Security Group 7. Review

Step 4: Add Storage

Your instance will be launched with the following storage device settings. You can attach additional EBS volumes and instance store volumes to your instance, or edit the settings of the root volume. You can also attach additional EBS volumes after launching an instance, but not instance store volumes. [Learn more](#) about storage options in Amazon EC2.

Volume Type	Device	Snapshot	Size (GiB)	Volume Type	IOPS	Throughput (MiB/s)	Delete on Termination	Encrypted
Root	/dev/sda1	snap-deb0000f	100	General Purpose SSD (GP2)	300 / 3000	N/A	<input checked="" type="checkbox"/>	Not Encrypted
EBS	/dev/sda2	Search (case-insensitive)	100	General Purpose SSD (GP2)	300 / 3000	N/A	<input type="checkbox"/>	<input type="checkbox"/>

[Add New Volume](#)

Free tier eligible customers can get up to 30 GiB of EBS General Purpose (SSD) or Magnetic storage. [Learn more](#) about free usage tier eligibility and usage restrictions.

Notice that the root volume can be increased accordingly and in some cases, there will be an attached Z: volume (usually referred to as 'instance store' or 'instance storage') to your instance where the extra space has been allocated. The Z: volume is a temporary volume that will be cleaned every time the instance is stopped. Without modification, it is then not recommended to install anything on this drive as the installation would be lost upon the first restart. The data in an instance store persists only during the lifetime of its associated instance. Hence, if an instance reboots (intentionally or unintentionally), the instance store (Z: drive) persists. However, the data that was stored on the Z: drive will be gone. It is a good practice to increase the C: drive as needed to store the data there instead.

While Qlik's engine is primarily memory intensive, larger configurations require appropriate disk I/O resources as well. Should you need to increase disk performance, the following link provides some guidance for selecting the appropriate disk configuration (including attaching additional EBS volumes for extra storage with the similar configuration).

<http://aws.amazon.com/ebs/details/#piops>

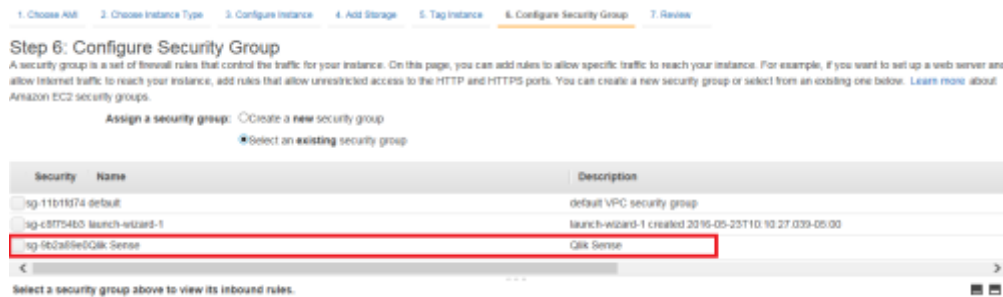
The following link describes the process of adding an EBS volume to an existing instance:

https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/ebs-attaching-volume.html?icmpid=docs_ec2_console

- e. **Tag instance:** this section will allow you to specify a tag name for the instance. Tags enable you to categorize your AWS resources in different ways, for example, by purpose, owner, or environment. The following link provides more information on how to set up tags. A blank tag is the default value and while a tag name doesn't have to necessarily be provided, it is recommended.

https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/Using_Tags.html?icmpid=docs_ec2_console

- f. **Configure Security Group:** this section will allow to specify the Security Group that should be used for the new instance. Select the option to use an existing Security Group and select the 'Qlik Sense' security group that was previously created. Click 'Review and Launch' when ready.

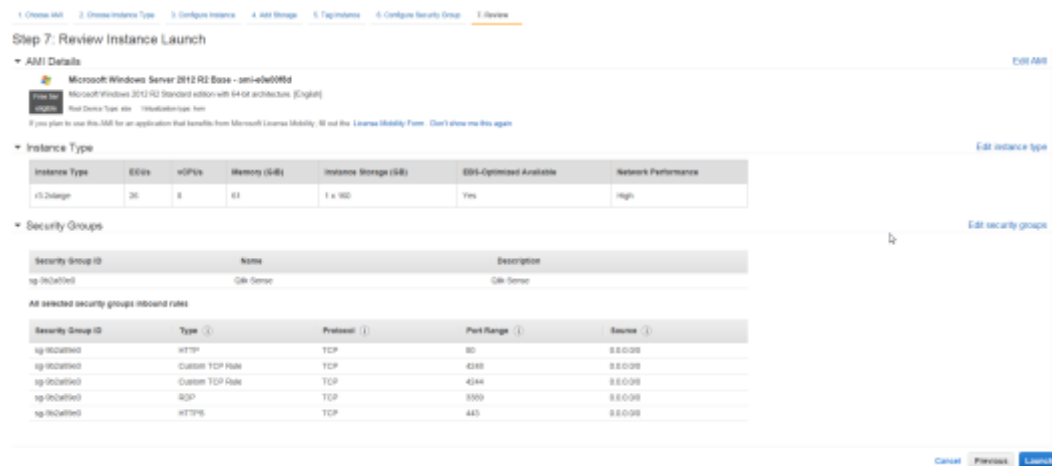


After selecting the 'Qlik Sense' security group, the Inbound rules configured will be displayed:

Inbound rules for sg-8b2af8e0 (Selected security groups: sg-8b2af8e0)

Type	Protocol	Port Range	Source
HTTP	TCP	80	0.0.0.0/0
Custom TCP Rule	TCP	4248	0.0.0.0/0
Custom TCP Rule	TCP	4244	0.0.0.0/0

- g. **Review:** in this last section, the configuration of the instance can be reviewed before being launched. Any modification needed can be made by either using the Back button or the links on the top. Click 'Launch' when ready.



4. When launched, a warning about key pairs will show up. Choose either an existing key pair or the "Qlik Sense" key pair that was previously created and saved.

Select an existing key pair or create a new key pair

A key pair consists of a **public key** that AWS stores, and a **private key file** that you store. Together, they allow you to connect to your instance securely. For Windows AMIs, the private key file is required to obtain the password used to log into your instance. For Linux AMIs, the private key file allows you to securely SSH into your instance.

Note: The selected key pair will be added to the set of keys authorized for this instance. Learn more about [removing existing key pairs from a public AMI](#).

Choose an existing key pair

Select a key pair

☐ I acknowledge that I have access to the selected private key file (Qlik Sense.pem), and that without this file, I won't be able to log into my instance.

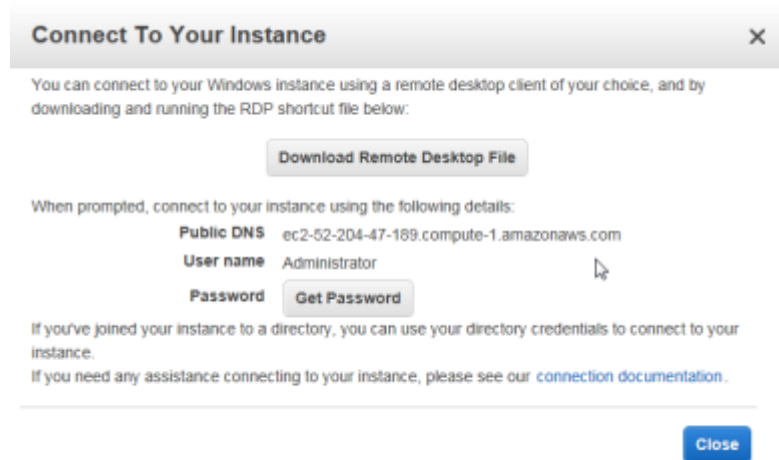
Cancel

Launch Instances

5. **(optional)** Create an Elastic IP so that the server IP address doesn't change even when the instance is stopped and started again. Please notice that there is a cost associated with this step, hence it is marked as **optional**.
 - Create an elastic IP via EC2 > Elastic IPs > Allocate New Address > Yes, Allocate
 - Select the new Elastic IP address, and click Associate Address
 - Choose the running instance and click Associate
6. Go back to the EC2 console and wait for the new instance to say "running" and "2/2 checks passed"

<input type="checkbox"/>	Name	Instance ID	Instance Type	Availability Zone	Instance State	Status Checks
<input checked="" type="checkbox"/>		i-1a506af1	m3.2xlarge	us-east-1c	running	2/2 checks passed

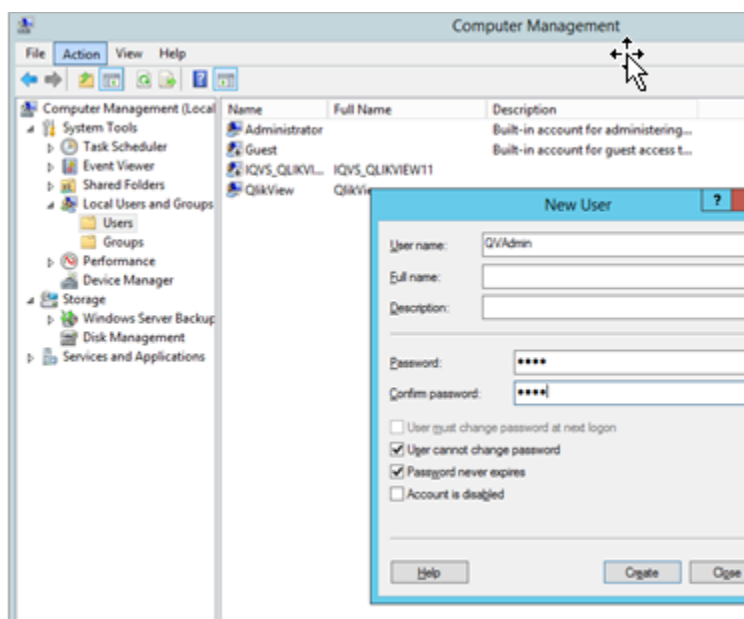
7. **RDP to the new instance**
 - Using the left hand side navigation panel, Click EC2 > Instances > Choose your instance > Connect
 - Click Download Remote Desktop File > Saves a RDP link you can use to connect



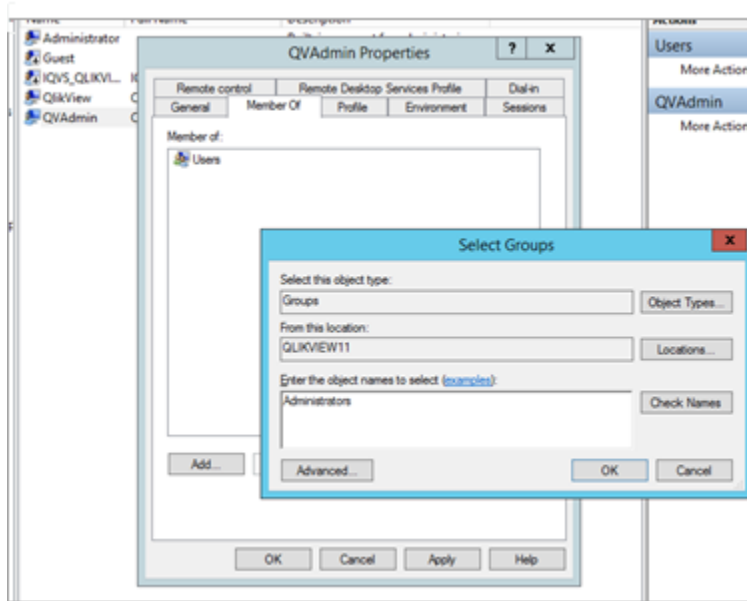
- Click Get Password > Choose your QlikSense.pem file > Click Decrypt password
- If applicable, securely store this string for future reference

Clicking on the Remote Desktop file that was downloaded will open an RDP session on the server. The decrypted password is needed in order to access the instance. After the initial login, the local Administrator password on the instance can be changed, if desired.

8. 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. Note 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 username (i.e. 'QSAdmin') and a password (i.e. 'QlikSense!')



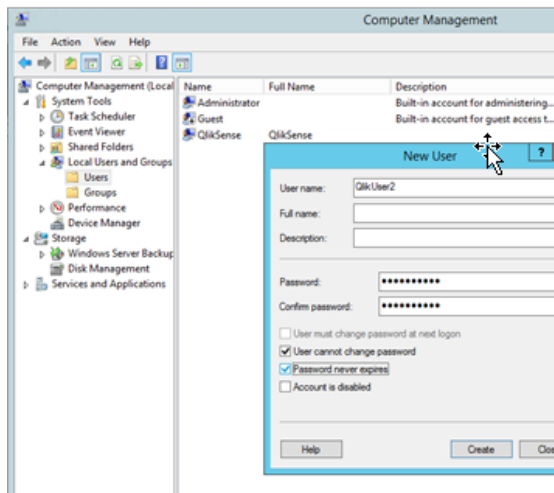
- d. Uncheck 'User must change password at next logon'
- e. Check 'User cannot change password' and 'Password never expires'
- f. Click Create
- g. Double click on 'QSAdmin'
- h. Click 'Member Of'
- i. Click Add
- j. Type 'Administrators'
- k. Click 'Check Names'



- 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'

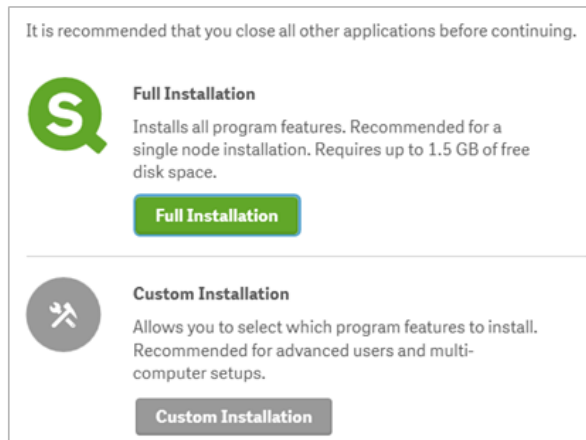


- f. If applicable, check 'User cannot change password'
- g. If applicable, check 'Password never expires'
- h. Click 'Create'

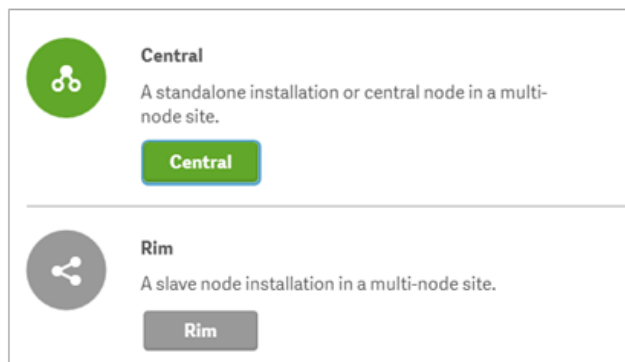
9. It is typically recommended to document the setup of all accounts in a table format to make it easy for administrators to keep control of such accounts. For example, the table below is a way in which such information can be documented. All security rules for password creation that are specific to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

User Name	Password	Account Settings	Windows Role Description	Function Description
QlikService	Q1kSense!	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	Q1kSense!	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

10. Once all users have been created, install Qlik Sense Enterprise.
 - a. Using a local administrator account on the instance, 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)



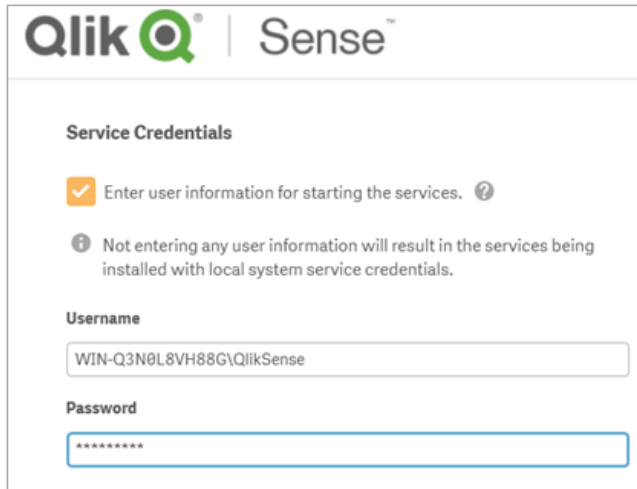
- d. Accept the license agreement
- e. Choose “**Central**”



- f. Enter a repository password. It can be anything, but be sure to document it, you will need it later! For example, the password could be the same as the service account user's password (i.e. 'Q1kSense!')

A screenshot of a form titled "Repository Database Superuser Password". It contains two input fields. The first field is labeled "Choose a password" with a question mark icon. The second field is labeled "Confirm the password". A mouse cursor is pointing at the first input field.

- 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.



Qlik Sense

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

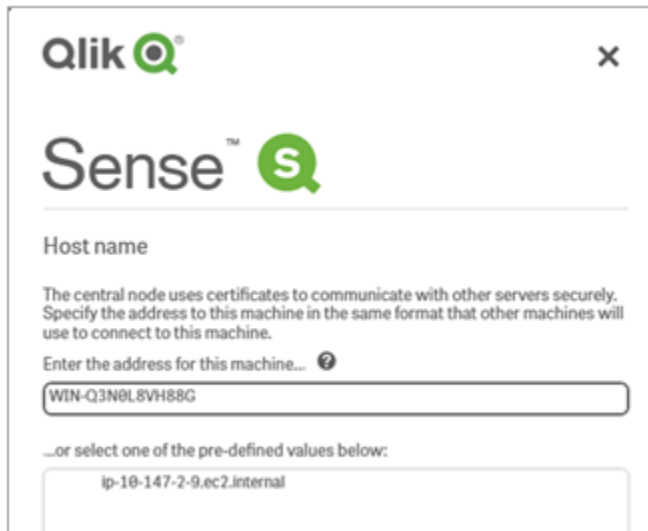
WIN-Q3N0L8VH88G\QlikSense

Password

11. The following is a critical step, **so follow the instructions carefully**. Enter the machine name without the domain name as shown below (replace the machine name with the one below). Avoid using the fully qualified name.

For example:

- IP address - 10.1.123.234
- machine name – WIN-Q3N0L8VH88G
- fully qualified machine name – WIN-Q3N0L8VH88G.ExampleDomain.COM



Qlik Sense

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.

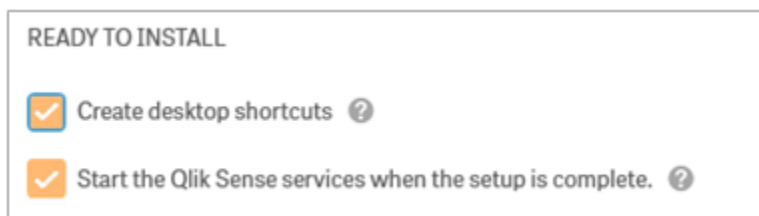
Enter the address for this machine... ?

WIN-Q3N0L8VH88G

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

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

- h. On the next screen, click 'Install'

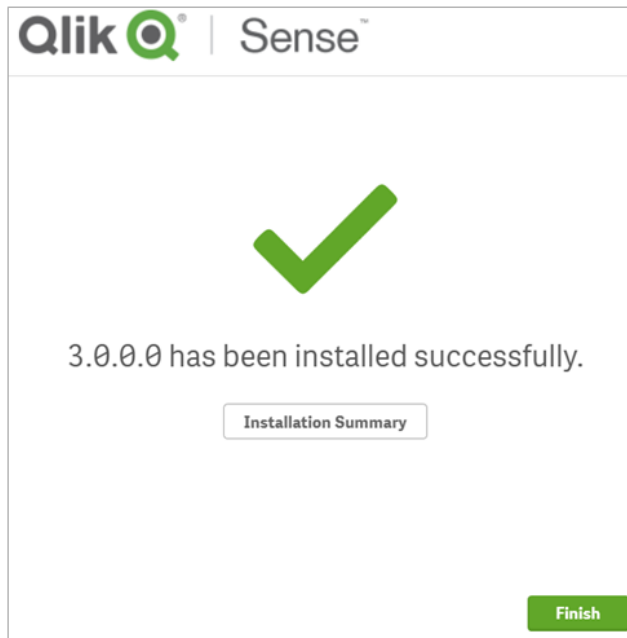


READY TO INSTALL

☒ Create desktop shortcuts ?

☒ Start the Qlik Sense services when the setup is complete. ?


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



- 12. Once the installation has been completed, you will need to set up the Qlik Sense Management Console (QMC). 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, wait at least 30 seconds, as the services may still be 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 RDP into the Qlik Sense Server. If you cannot connect, make sure that the right security group is being used.



- 13. 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, contact a Qlik representative to acquire the LEF.

SITE LICENSE

Owner name

Owner organization

QT

Serial number

Control number

▼ LEF access

Get LEF and preview the license

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

✓

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.

View the guide for Setting up your Qlik Sense site

14. Now that the license has been validated, the tokens need to be assigned to users. 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 allocate tokens, follow these 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. From a 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 that was created earlier (i.e. QSAdmin).

User access allocations Total: 1 Showing: 1 Selected: 1				Actions
Name	User directory	Status	Last used	
administrator	WIN-RFJRP6RR3Q	Allocated		

License and tokens

License usage summary

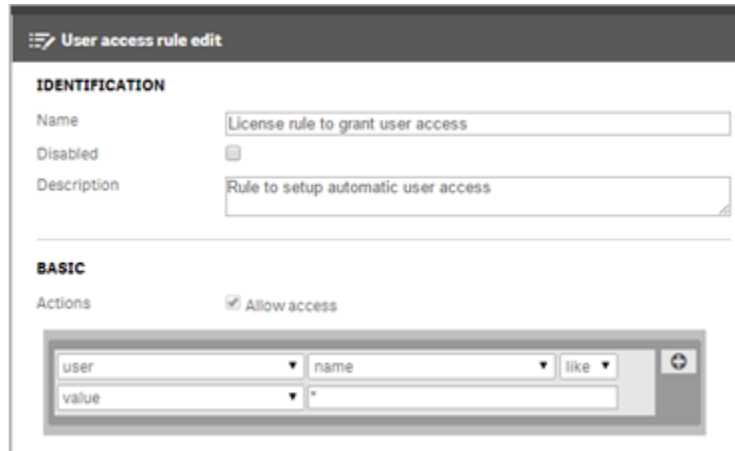
✓ User access allocations

- d. Typically in development and testing environments, it comes handy to have anyone connecting to Qlik Sense Server being able to get a license. This approach may not be adequate for production environments. To learn more on how to manage a production environment, visit the following link:

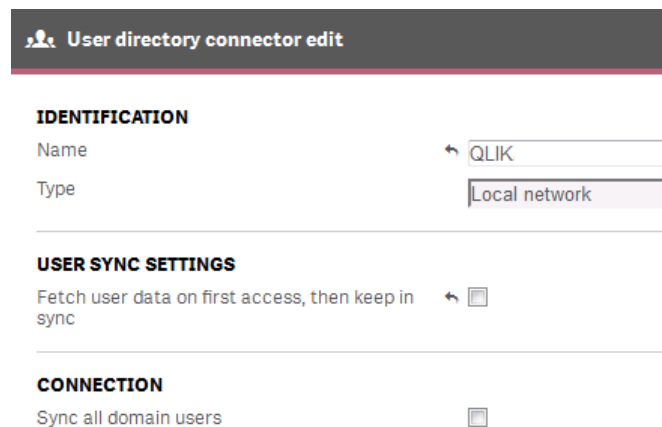
<http://help.qlik.com/en-US/sense/2.2/Subsystems/ManagementConsole/Content/manage-license-and-tokens.htm>

In order to grant access for users, either a User Access Rule or a User Directory could be leveraged. In order to create an access rule, click on the Start button on the top left corner of the QMC > License and Tokens > User Access Rules >

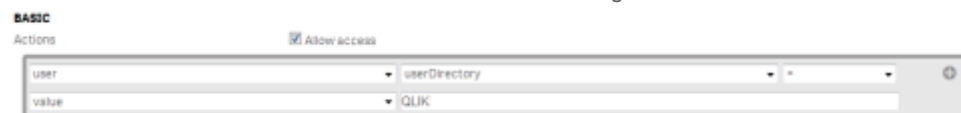
- e. Click “Create New”
- f. Click “Basic”
- g. Create the rule as ‘user name like value *’



- h. Click Apply.
- i. Creating a User Directory (rather than adding individual rules) allows to have the rules on how to assign tokens based on such User Directory. In the simplest scenario, a User DirectoryConnector would be created. The local Directory could be called anything (i.e. QLIK):



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



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>

15. 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. If an Elastic IP address was created (optional), add the elastic IP address of the Qlik Sense Server (external IP address found in EC2 > Instances > Machine > Public IP/Public DNS/Public Domain fields)
- i. Click Add New Value
- j. Add the AWS machine name of the Qlik Sense Server
- k. Click Add New Value
- l. Add the Public DNS (found in EC2 > Instances > Machine > Public IP/Public DNS/Public Domain fields)
- m. Click Add New Value
- n. Add the Public Domain (found in EC2 > Instances > Machine > Public IP/Public DNS/Public Domain fields)



- o. Click Apply. A warning saying that the proxy will have to be restarted will appear. Click OK. If any DNS entries has been created (e.g., qlikbi.company.com), add these to the whitelist. If using HTTPS, import a SSL certificate. Adding the values to the Websocket Origin Whitelist allows Qlik to accept URLs of these formats.
16. 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 AWS 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:

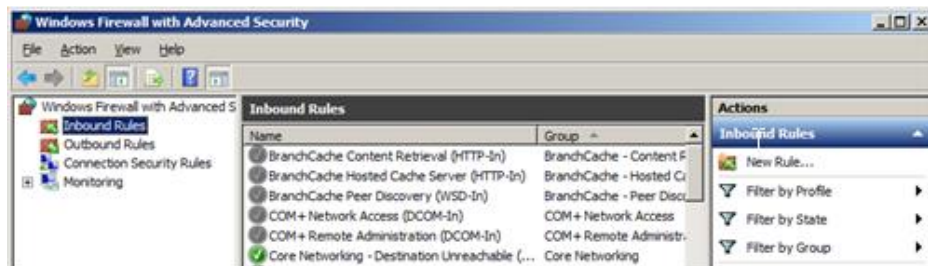
<http://help.qlik.com/en-US/sense/2.2/Subsystems/Hub/Content/Hub/hub.htm?q=hub>

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:

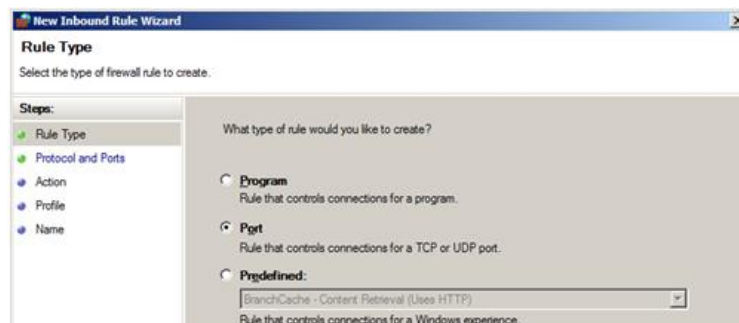
`http://MachineName/hub`
`http://MachineName/qmc`

These addresses will only work locally within the AWS VPC 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 (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.



- c. Select 'Port' as the rule type.



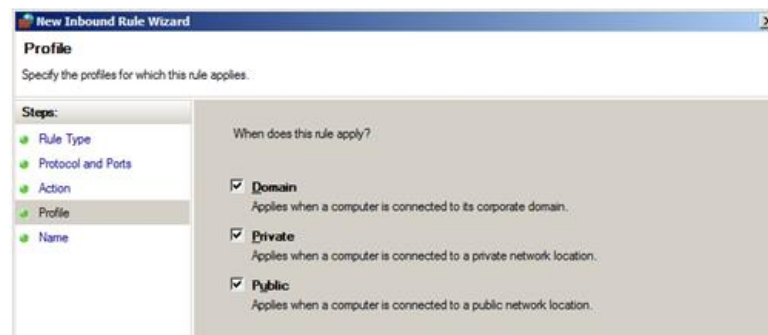
- d. In the 'Specific local ports' field, enter '80,443,4244,4248'. This will allow you to access the Qlik Sense Hub and QMC. As a reminder, probably HTTP and HTTP were not both enabled, so only the appropriate port out of these two would have to be used here.



- e. Ensure that 'Allow the connection' is selected and click 'Next'.



- f. Ensure that all options are ticked and click 'Next'.



- 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.

17. In order to test access to the Qlik Sense HUB and QMC from outside the AWS 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://AWSPublicDNS /hub
https://AWSInstancePublicIP/hub
https://AWS PublicDNS/qmc
https://AWSInstancePublicIP/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://AWSPublicDNS /hub
http://AWSInstancePublicIP/hub
http://AWS PublicDNS/qmc
http://AWSInstancePublicIP/qmc
```

18. 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 AWS EC2 requires additional configuration. If a Client Provided Browser Certificate is available, the certificate could be imported into the Qlik Sense Server 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 can run 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 better understand how Qlik Sense running on AWS can easily scale, a benchmarking analysis has been documented in a technical brief that is publicly available (see link below for the Qlik® Sense Performance Benchmark Technical Brief). 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/en/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=8H0JJYUft8
- https://www.youtube.com/watch?feature=player_embedded&v=GiSF1yW9D9g
- 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/en/resource-library/qlik-sense-performance-benchmark>

Connectivity

As mentioned earlier on this document, Amazon Web Services 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 Amazon's cloud platform are:

- Ease of use – a platform can be constructed in hours, unlike acquiring a physical server which may take weeks
- Elasticity – capacity can be grown or shrunk on demand
- Cost optimization – the cost of a platform can be easily matched to the benefits gained.

Under the AWS banner, Amazon offers a number of web services that can be used as data stores where Qlik Sense applications can pull data from in order to build applications. Such services include:

- Amazon DynamoDB – NoSQL database
<http://aws.amazon.com/dynamodb/getting-started/>
- Amazon RDS – managed relational database service
http://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/CHAP_GettingStarted.html
- Amazon Redshift – data warehouse as a service
<https://aws.amazon.com/redshift/getting-started/>
- Amazon Simple Storage Service (S3) – scalable, object-based cloud storage
<https://aws.amazon.com/s3/>
- AWS Elastic Map Reduce (EMR) – managed Hadoop service
<https://console.aws.amazon.com/elasticmapreduce>

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 Amazon RDS and Amazon Redshift
- b. an application that uses data coming from an AWS data source as well as a combination between flat files and web based data sources (i.e. a web service data feed)
- c. an application that uses data stored in AWS 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 AWS infrastructure, it is very common to see Qlik Sense applications leveraging native AWS data sources such as Amazon Redshift. Amazon Redshift has been one of the fastest growing products within AWS, and for a good reason. It is a fast, simple and highly scalable relational data warehousing service that massively processes workloads in parallel to provide superior performance in a fully managed environment, all with zero administration overhead. By using such features as columnar storage, data compression, zone maps and high-speed direct attached storage, Amazon Redshift dramatically reduces I/O, which is often the bottleneck in the data warehousing systems. Amazon Redshift parallelizes and distributes queries, load jobs, backup/restore tasks and resizing activities. For organizations that are using Amazon Redshift for their traditional enterprise data warehouses, their Big Data scenarios or SaaS offerings, Qlik Sense makes it very easy to consume the data that is stored and loaded in Amazon Redshift and use it in concert with other data sources. To learn more about the integration of Qlik's platform with Amazon Redshift, please visit the following link: <http://aws.amazon.com/redshift/partners-detail/qlik/>



Qlik Sense is great for data sources outside AWS as well. Qlik Sense can connect to these data sources from the cloud by leveraging a VPN connection (i.e. using a virtual private gateway attached to the Amazon VPC that talks to a gateway located where the data source(s) in question reside). AWS Direct Connect is a service that allows enterprises to establish a dedicated network connection from their datacenter, office or colocation environment to AWS. AWS Direct Connect allows you to bypass the public Internet when connecting to AWS, which in many cases can reduce your network costs, improve bandwidth throughput and provide a more consistent and secure network experience. With each AWS Direct Connect connection, customers can enable one or more logical connections that allow access to public AWS resources (such as S3 buckets) and private VPC networks. AWS Direct Connect is compatible with all AWS services, including Amazon Elastic Compute Cloud (EC2), Amazon Virtual Private Cloud (VPC), Amazon Simple Storage Service (S3), and Amazon Relational Database Service (RDS) as well as private resources (such as Amazon EC2 instance running within an Amazon VPC using private IP space), while maintaining public and private environments separated. In order to learn more about AWS Direct Connect, visit the following link:

<http://docs.aws.amazon.com/directconnect/latest/UserGuide/Welcome.html>

As mentioned before in this section, Qlik Sense leverages native ODBC drivers when available. In the case of Amazon Redshift, Qlik Sense establishes an ODBC connection in order to connect. We recommend users to follow Amazon Redshift best practices to determine how to structure data in an optimal way for the fastest query performance. The following diagram represents, at a high level, the architecture of an Amazon Redshift and Qlik Sense environment. The Amazon Redshift Leader Node acts as the SQL endpoint storing metadata information and coordinating query execution. The compute nodes act on their local columnar storage to enable the parallel execution of queries, as well as resources for load, backup and restore activities via Amazon S3.




In order to establish ODBC connectivity from Qlik Sense Enterprise to AWS Redshift, first download and install the Amazon Redshift ODBC driver onto the Qlik Sense Server. In order to learn more about the Amazon Redshift ODBC client and to download, visit the following links:

<http://docs.aws.amazon.com/redshift/latest/mgmt/configuring-connections.html#connecting-drivers>
<http://docs.aws.amazon.com/redshift/latest/mgmt/install-odbc-driver-windows.html>

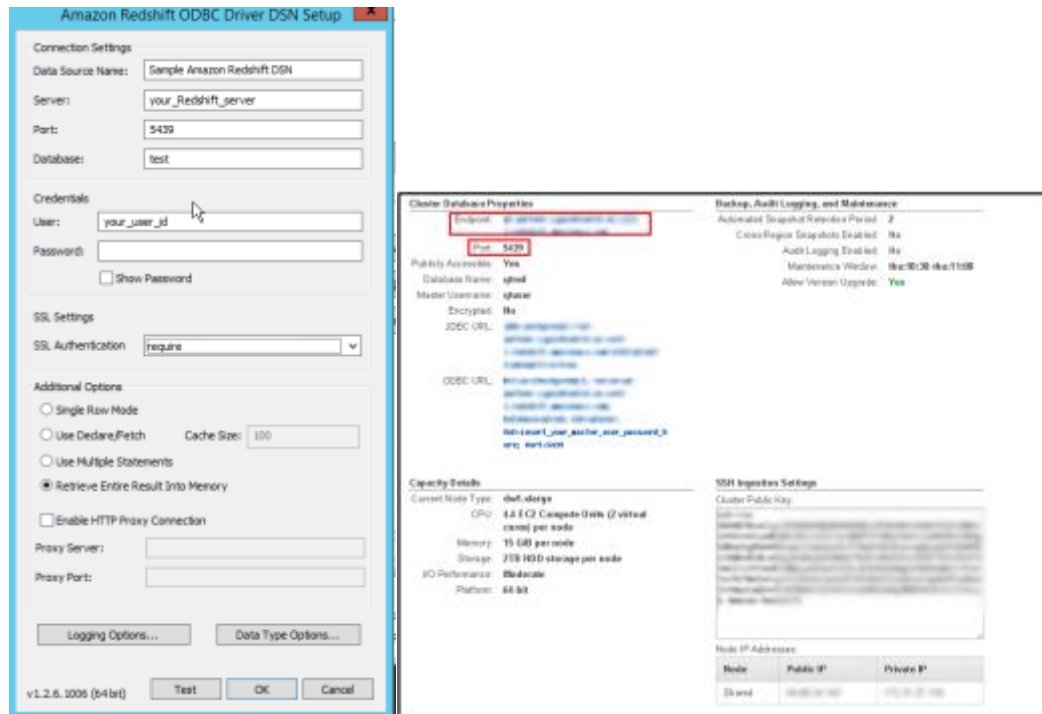
Once the ODBC driver, has been installed, follow the next steps to create the ODBC DSN that will be used by Qlik Sense Server:

1. Start the “Data Sources (ODBC) 64-bit” program from the Control Panel in Windows
2. You should see a dialog box that looks like this:



3. In the System DSN tab, highlight the Sample Amazon Redshift DSN that has been automatically created during the installation and click the “Configure” button
4. Enter the information requested. In the Server box, enter the endpoint string from the Amazon Redshift cluster definition page (from the AWS Management Console → Databases → Redshift → Clusters and click on the cluster in question). For the Port number, use the same port number that is defined in the Amazon Redshift Cluster definition page. For name, use the name of the database within the cluster where you want to connect. The DSN can be configured to enable SSL authentication if desired. The configuration is transparent to the Qlik Sense Server, it is

managed entirely through ODBC Administrator panel. Depending on the location of the Amazon Redshift database, the recommended option for the SSL authentication setting may vary.



5. Once all the settings have been set accordingly, click on the 'Test' button to make sure that the connection has been created correctly.

While this document does not cover every single AWS data source, each AWS data source provides native connectivity that would be leveraged in a similar way by Qlik Sense Server. Depending on the source in question, AWS provides a number of very good references 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 Amazon Redshift. Typically, Qlik recommends to leverage the processing power of AWS Redshift 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 AWS Redshift Data Warehouse to get optimal performance to work well with Qlik Sense Server:

- Understanding and analyzing explain plans
<http://docs.aws.amazon.com/redshift/latest/dg/c-optimizing-query-performance.html>
- Selecting sort keys
http://docs.aws.amazon.com/redshift/latest/dg/c_best-practices-sort-key.html
- Selecting best distribution keys
http://docs.aws.amazon.com/redshift/latest/dg/c_best-practices-best-dist-key.html
- Identifying smallest column size and data set.
http://docs.aws.amazon.com/redshift/latest/dg/c_best-practices-smallest-column-size.html
- Compression
http://docs.aws.amazon.com/redshift/latest/dg/t_Compressing_data_on_disk.html

- Understanding data distribution
http://docs.aws.amazon.com/redshift/latest/dg/t_Distributing_data.html


Creating a new Qlik Sense Application

Creating new applications in Qlik Sense is an uncomplicated process. Qlik Sense technology is based on a notion of applications. An application in the context of Qlik Sense is a file with a QVF extension. A QVF file (or application) is the container that captures contains everything required for users to be able to analyze, drill anywhere, visualize, explore and discover their data. In essence, think of a QVF as the unit of interaction for users within the Qlik Sense environment to analyze and discover the data. The application knows where the data sources are located and how they are being used.



The QVF file contains three elements:

- Visualizations: this layer of the application is composed of the following elements: sheets that contain visualizations, bookmarks for specific set of selections and interactive stories with live data.




- Data: when a Qlik Sense application connects to a specific data source, it only captures each one of the distinct data points once within each field, table, and data source, which allows for extremely efficient processing. Based on the data structure of the different data sources, the application automatically builds an associative data model that enables the discovery experience within 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](http://help.qlik.com/en-US/sense/2.2/Subsystems/Hub/Content/Scripting/introduction-data-modeling.htm?q=data%20loading)

- 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 logic automatically as part of the application's script. The script can be manually modified if the Use Case requires it.



There are many public sources that can be leveraged to learn more about how to build a new application in Qlik Sense from scratch. These great sources for guidance are: Qlik Community site (registration on the site is required), the Qlik Help site and Qlik's YouTube channel. We strongly recommend the first three links below for those that are getting started with Qlik.

<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_gSrO7bIgP1xF0HowDXnG1S6I&index=3

For more thorough training resources, visit the training section on the Qlik website (under the Training section):

<http://www.qlik.com/services/training/qlik-sense-training>

Reference Architecture

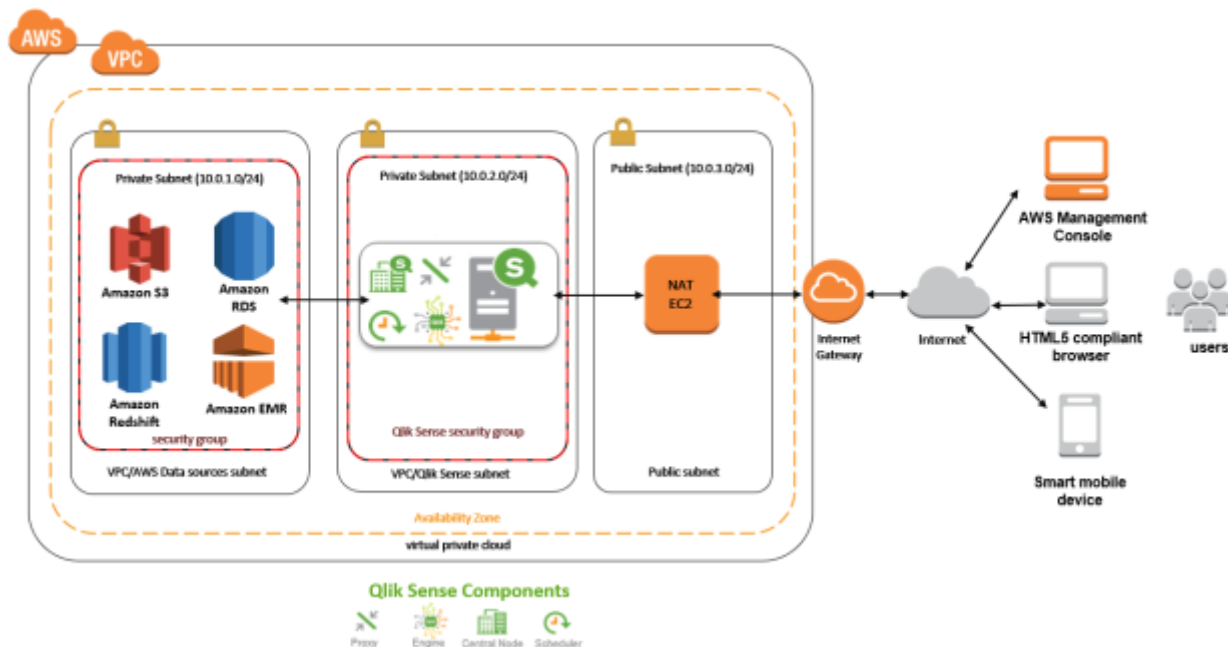
Whether Qlik Sense is deployed to AWS manually or by leveraging AWS Marketplace, the following components should be deployed and configured as part of the reference architecture:

- An Amazon VPC: Amazon Virtual Private Cloud (Amazon VPC) lets you provision a logically isolated section of the Amazon Web Services (AWS) cloud where you can launch AWS 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 AWS VPC, visit the following link:
<http://aws.amazon.com/vpc/>
- Subnet: at least one subnet (either public or private) within the VPC. In the case of just one, it would be

considered the default one

- NAT instance (optional): although not required, in the case of having a public subnet within the VPC, a NAT instance deployed into the public subnet and configured with an Elastic IP address (EIP) for outbound Internet connectivity. The NAT instance is used for Internet access if any Amazon EC2 instances are launched within the private network.
- Instance: A Windows 2012 Server instance deployed in the default subnet for Qlik Sense Server.
- Security groups: one for each instance or node within the environment with restricted access. Information on the protocols and ports opened will be part of the definition. At the very least, one group that goes with the VPC
 - AWS Identity and Access Management (IAM): depending on the expected administration of the environment, an AWS IAM may be needed in order to manage fine-grained permissions for access to various AWS services involved in the deployment process. To learn more about AWS IAM, visit the following site: http://docs.aws.amazon.com/IAM/latest/UserGuide/introduction.html?icmpid=docs_iam_console
- 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 AWS. Hence, one option is to launch the Qlik Sense node within a public subnet providing direct Internet access. Alternatively, another option is to deploy the Qlik Sense node in a private subnet. In the case that a configuration file or cloud formation template is being use to manage the deployment, the configuration file can manage whether the Qlik Sense node should be deployed to a public or private subnet. The decision on public vs. private deployment is typically based on the overall solution requirements. The diagram below represents a reference architecture where Qlik Sense is within its own private subnet and various AWS 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 about the installation process, part of the process of creating your AWS configuration 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 and logging user sessions. There are three main ways to architect the storage for this application:

- **Ephemeral storage:** part of the configuration process of the EC2 instance that is used for the Qlik Sense server is to determine the amount of storage use the storage within the EC2 instance before without having to add extra volumes for further storage. Once it is determined the expected storage space needed, then the storage can be allocated accordingly. This option provides plenty of storage space, (up to 16 TB - <https://aws.amazon.com/about-aws/whats-new/2015/03/amazon-efs-increases-maximum-performance-size/>). This is usually the lowest cost option.

Step 4: Add Storage

Your instance will be launched with the following storage device settings. You can attach additional EBS volumes and instance store volumes to your instance, or edit the settings of the root volume. You can also attach additional EBS volumes after launching an instance, but not instance store volumes. [Learn more](#) about storage options in Amazon EC2.

Volume Type	Device	Snapshot	Size (GB)	Volume Type	IOPS	Throughput (MB/s)	Delete on Termination	Encrypted
Root	/dev/sda1	snp-d8b600f	100	General Purpose SSD (GP2)	100 / 3000	N/A	<input checked="" type="checkbox"/>	Not Encrypted
EBS	/dev/sdb	Search case internal	100	General Purpose SSD (GP2)	300 / 3000	N/A	<input type="checkbox"/>	<input type="checkbox"/>

[Add New Volume](#)

Free tier eligible customers can get up to 30 GB of EBS General Purpose (SSD) or Magnetic storage. [Learn more](#) about free usage tier eligibility and usage restrictions.

- **EBS (Elastic Block Store):** EBS provides persistent block level storage volumes for use with Amazon EC2 instances in the AWS Cloud. Each EBS volume is automatically replicated within its Availability Zone to protect you from component failure, offering high availability and durability. EBS volumes offer the consistent and low-latency performance needed to run your workloads. With EBS, you can scale your usage up or down within minutes, all while paying a low price for only what you provision. During the creation of the EC2 instance for Qlik Sense, more volumes can be attached to the instance as needed in order to increase the storage space within the Qlik Sense Server. This is the recommended solution for most Use Cases over S3. For more information, visit the following link: <http://aws.amazon.com/ebs/>
- **Leverage Amazon S3 (Simple Storage Service):** provides a secure, durable, highly-scalable cloud storage for source data and backups. Amazon S3 is easy to use object storage, with a simple web service interface to store and retrieve any amount of data from anywhere on the web. Amazon S3 offers a range of storage classes designed for different use cases including Amazon S3 Standard for general-purpose storage of frequently accessed data, Amazon S3 Standard - Infrequent Access (Standard - IA) for long-lived, but less frequently accessed data, and Amazon Glacier for long-term archive. Given that Qlik Sense Server runs on a Windows operation system, leveraging S3 as a storage mechanism would imply mounting the files on Linux and then attach the volume to the Qlik Sense Server as a NAS drive. As mentioned above, EBS represents a more straight forward option for Qlik Sense. In order to learn more about S3, visit the following link: <http://aws.amazon.com/s3/>

It is important that as part of the planning process, the expected needed storage is determine ahead of time so then the AWS 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 Server, it is a module within the Qlik Sense Proxy Service (QPS) 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 settings 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 used options for authentication:

- 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.
- Security Assertion Markup Language (SAML):** Qlik Sense Server can leverage SAML to enable single-sign-on (SSO) and 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, the AWS Directory Services can be leveraged. AWS Directory Services makes it easy to setup and run Microsoft Active Directory (AD) in the AWS cloud, or connect your AWS resources with an existing on-premises Microsoft Active Directory. AWS Directory Service provides you with three directory types to choose from including AWS Directory Service for Microsoft Active Directory (Enterprise Edition), also referred to as Microsoft AD, as well as Simple AD, and AD Connector. AWS Directory Services makes connecting to an on-premises directory and leveraging the connection while using existing corporate credentials to grant access to AWS 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 AWS Directory Services, visit the following site:

<http://aws.amazon.com/directoryservice/>

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 an overview and description of the following 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 recommended 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. 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 about 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) describes 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 an app that has not been published yet:

<http://help.qlik.com/en-US/sense/2.2/Subsystems/ManagementConsole/Content/publish-apps.htm?q=publishing>

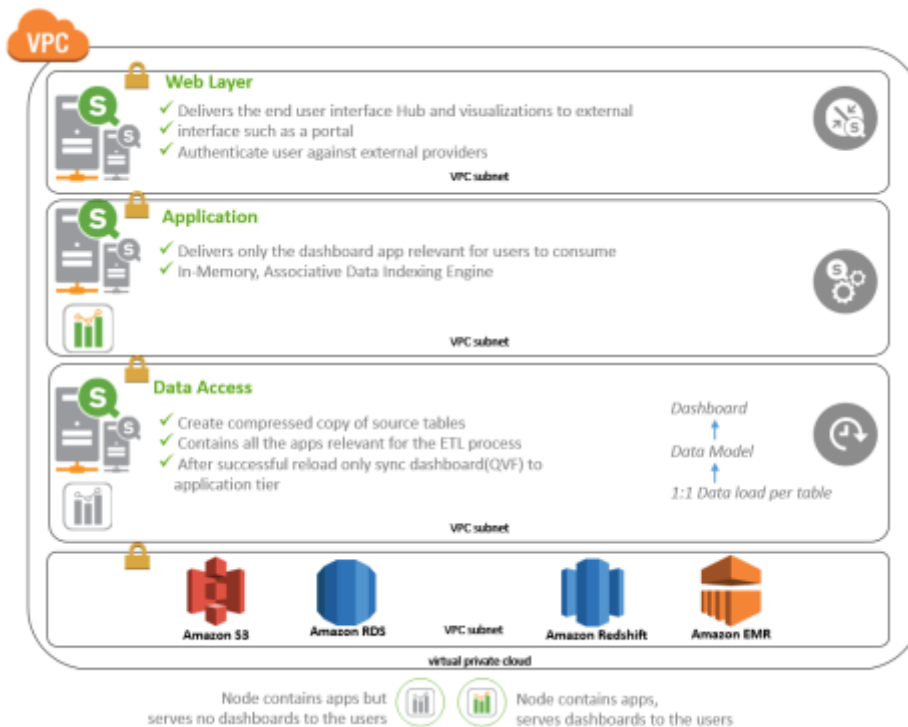
Designs for Deploying Qlik Sense on AWS

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 AWS, 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 AWS EC2 instance:



Multi-node deployments

While the scope of this document is for a single node deployment, the following section describes at a high level some of the considerations for a multi-node deployment. 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.



When considering then the different layers (front-end, application, management and back-end), there would be various possible configurations. For example, in one end of the spectrum, all services could be running off the same node (scenario #1 below). As some of the services are assigned to different nodes, there could be multiple layers. A classic example is having a front-end layer with the proxy service, a mid-tier layer running the engine and a back-end layer acting as the central node and running the scheduler service (scenario #4 below). Each layer can have one or more nodes (scenario #5 below). And whether it is a single node environment or a multi-node environment, multiple environments could leverage a centralized location of files that contain applications (QVF files) as well as data files in the form of QVD files (scenario #8 below). The are other scenarios depicted below:




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.glik.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.glik.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.glik.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 servers 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.



It is also possible to utilize a third party Load Balancer by configuring in the QMC the Virtual Proxy to use a “Load balancing module base URI” under the Integration settings. Further detail is available on **Appendix A** of this guide.



High Availability and Disaster Recovery (DR) Strategy

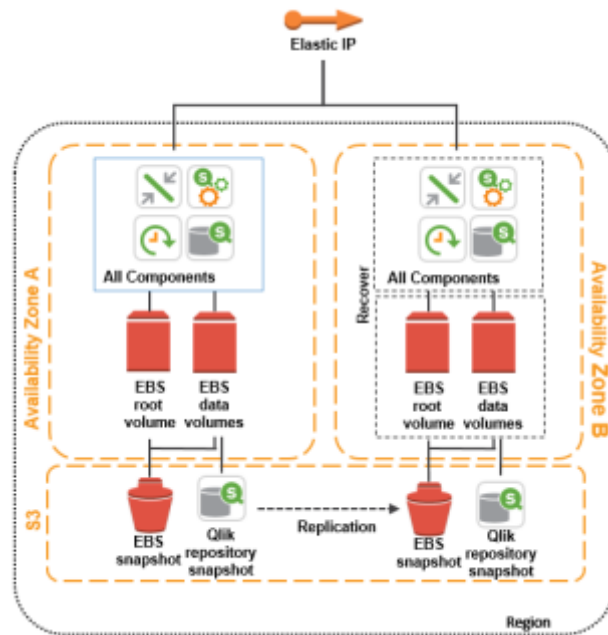
Having a multi-node environment, allows to have a more robust strategy for High Availability scenarios as well as a DR strategy. Qlik Sense environments can be backed up and recovered easily. To learn more, visit the following link:

<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 AWS offers. All the information about the Qlik Sense environment is stored in the EBS root and data volumes snapshots. This includes the Qlik Sense Repository which snapshots are taken via Qlik Repository Snapshot Manager used as an alternative to or in conjunction with EBS snapshots. 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>

Such volumes can be backed up into S3 on scheduled basis. In parallel, a mirror passive environment of Qlik Sense (DR Environment) under a different Availability Zone is always ready to go. The DR environment has access to the latest EBS snapshots via a replication process within S3. All users are redirected to the DR Environment via Elastic IP.



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. Use Amazon Application Load Balancer if using SAML authentication. **Qlik Sense with Windows Authentication through Amazon Application Load Balancer is not supported.** Do not use Amazon Classic Load Balancer, as it does not support websockets and stickiness simultaneously. Alternatively, instead of a Load Balancer, Elastic IP associated with one Qlik Sense Proxy could be used instead. In the event of Availability Zone failure, then the Elastic IP would have to be manually associated to the Qlik Sense Proxy on a different Availability Zone. In such case, Qlik Sense Proxy needs to be configured to load balance to engines in both Availability Zones. Another change is the distribution of the Qlik Sense applications across the different Availability Zones. 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 Availability Zones. In all cases, it is strongly recommended to use EBS storage and use the Qlik Repository Snapshot Manager to take the Repository snapshots. . The following diagram depicts a multi-node Qlik Sense environment in AWS with High Availability and Disaster Recovery.



High-Performance Storage

Qlik Sense Server provides an 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. It 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 AWS Cloud platform enables Organizations to provide cloud based analytical solutions to their various stakeholders. The mechanics to deploy Qlik Sense Server on AWS are not too different from a traditional on premises 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 AWS Services (Amazon Redshift, Amazon Elastic Compute Cloud, Amazon Relational Database Services, etc.) can be leveraged 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
- Karthik Krishnan, Solutions Architect, Amazon Web Services
- David Potes, Partner Solutions Architect, Amazon Web Services

Notes

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, AWS Elastic Load Balancing (ELB) automatically distributes incoming traffic across multiple EC2 instances by creating a load balancer and registering instances with the load balancer in one or more AWS Availability Zones. The load balancer serves as a single point of contact for clients which increases the availability of Qlik Sense application. EC2 instances can be added or removed from the load balancer as needed, without disrupting the overall flow of information. If an EC2 instance fails, Elastic Load Balancing automatically reroutes the traffic to the remaining running EC2 instances. 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, AWS ELB 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 (AWS ELB with TCP protocol supports the WebSockets): <http://blog.flux7.com/web-apps-websockets-with-aws-elastic-load-balancing>

However, 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

