

# Using Microsoft Power BI with the AWS Cloud

*January 2021*



## Notices

Customers are responsible for making their own independent assessment of the information in this document. This document: (a) is for informational purposes only, (b) represents current AWS product offerings and practices, which are subject to change without notice, and (c) does not create any commitments or assurances from AWS and its affiliates, suppliers or licensors. AWS products or services are provided “as is” without warranties, representations, or conditions of any kind, whether express or implied. The responsibilities and liabilities of AWS to its customers are controlled by AWS agreements, and this document is not part of, nor does it modify, any agreement between AWS and its customers.

© 2021 Amazon Web Services, Inc. or its affiliates. All rights reserved.

# Contents

- Introduction ..... 1
- The Microsoft Power BI Suite.....2
  - Power BI Desktop.....2
  - Power BI service .....2
  - Power BI Report Server .....3
  - On-premises data gateway .....3
- Connecting Microsoft Power BI Desktop to AWS data sources .....3
  - Using Power BI Desktop on premises .....3
  - Using Microsoft Power BI Desktop in the AWS Cloud ..... 10
  - Summary of Microsoft Power BI Desktop connectivity options ..... 14
- Connecting the Microsoft Power BI service to AWS data sources ..... 14
  - Recommended configuration ..... 14
  - Additional considerations ..... 15
- Using Microsoft Power BI Report Server in AWS..... 17
- Using Microsoft Power BI with Amazon QuickSight ..... 20
- Conclusion ..... 22
- Contributors ..... 22
- Further reading ..... 23
- Document Revisions..... 23
- Appendix: Microsoft Power BI supported AWS data sources ..... 23

# Abstract

This whitepaper discusses how to integrate and use Microsoft Power BI (Desktop, Report Server, Service, and on-premises data gateway) with the AWS Cloud. It presents options for customers looking to connect Microsoft Power BI products to AWS services such as Amazon Redshift, Amazon Athena, and Amazon RDS, with a focus on connectivity, security, performance, and cost optimization.

This whitepaper is for IT decision makers and architects looking to quickly understand Microsoft Power BI concepts and what options exist to make use of those technologies when using AWS services as data sources.

## Introduction

Customers of all sizes are using AWS products and services to store their data reliably, cost effectively, and securely. This is due in part to the broad ecosystem of mature data storage and analytics offerings that are available. Some of these offerings include the following services:

- Amazon S3 provides a simple, scalable, secure, and cost-effective data repository. It has become an industry standard for storing application data, as well as a first choice for customer data lakes.
- Amazon Athena is an interactive query service that makes it easy to analyze data in Amazon S3 using standard SQL.
- Amazon Relational Database Service (Amazon RDS) makes it easy to set up, operate, and scale a relational database in the cloud. It provides cost-efficient and resizable capacity while automating time-consuming administration tasks such as hardware provisioning, database setup, patching, and backups. SQL Server, Oracle Database, MySQL, MariaDB, and PostgreSQL engines are available.
- Amazon Redshift is fully managed, massively scalable data warehouse that makes it easy to analyze both structured and unstructured datasets.
- Amazon QuickSight is a fast, cloud-powered business intelligence service that makes it easy to deliver insights to everyone in your organization.

AWS gives customers full flexibility in mixing the technologies they prefer for their data needs. While many customers choose Amazon QuickSight for their business intelligence (BI) needs, other customers choose vendors such as Microsoft Power BI, Tableau, and Qlik.

This document focuses on the Microsoft Power BI suite of products and services, and how to use them in combination with AWS services.

To better understand how services relate to one another, we often label data services as either being *data sources* or *data consumers*. A data source allows customers and applications to store and retrieve data from the service. Frequently, data sources also have built-in compute and can provide computational analysis and filtering. But, ultimately, data is loaded into these data sources and eventually data is retrieved from them by data consumers. Amazon S3, Amazon Athena, and Amazon Redshift are good examples of data sources.

Data consumers, on the other hand, access the data from data sources and, typically, process it. They might optionally display it too. Amazon QuickSight and the Microsoft Power BI suite are good examples of data consumers. They read from data sources, and then assist in the analysis, visualization, and publication of information.

It's easy to see how a data consumer can quickly become a data source for another upstream service. However, for the purposes of this paper, we consider Microsoft Power BI a data consumer, and the AWS services such as Amazon RDS, Amazon Redshift, and Amazon Athena, as data sources.

## The Microsoft Power BI Suite

To reduce confusion due to product naming similarities, we present what each Microsoft Power BI product and service is.

### Power BI Desktop

Power BI Desktop is a free application you install on your local computer. It lets you connect to, transform, and visualize your data. With Power BI Desktop, you can connect to multiple different sources of data and combine them (often called *modeling*) into a data model. This data model lets you build visuals and collections of visuals you can share as reports with other people inside your organization.

Power BI Desktop can connect to any supported data source that is available locally or over the network. For supported data sources, see [Microsoft Power BI supported AWS data sources](#).

Most users who work on business intelligence projects use Power BI Desktop to create reports. Then they push content to either Power BI Report Server or the Power BI service in order to share their reports with others. The act of pushing content from Power BI Desktop to the Power BI Report Server or the Power BI service is known as *publishing*. For more information, see [What is Power BI Desktop?](#)

**Note:** Power BI Desktop is a Windows-only application. It is not available for Linux, macOS, or other operating systems.

### Power BI service

Power BI is a collection of software services, apps, and connectors that work together to help you create, share, and consume business insights in a way that serves you and your business most effectively. The Power BI service, sometimes referred to as Power BI online, is the *software as a service* (SaaS) part of Power BI.<sup>1</sup>

The Power BI service is a cloud-based service. It supports light report editing and collaboration for teams and organizations. You can connect to data sources in the Power BI service too, but modeling is limited.

Most report designers who work on business intelligence projects use Power BI Desktop to create reports, and then use the Power BI service to distribute their reports with others.<sup>2</sup>

## Power BI Report Server

Power BI Report Server is a private report server with a web portal in which you display and manage reports and KPIs. Customers use Power BI Report Server in cases where they do not want their reports published to the Power BI service. Although it was originally intended for on-premises environments, the Power BI Report Server can run on AWS as well.

## On-premises data gateway

The Microsoft On-premises data gateway acts as a bridge between privatized data sources that are stored outside of Power BI service, and the Power BI service. It allows the Power BI service to access privatized data sources, which are located in another facility and accessible by internal network connectivity between the data source and the data gateway. Although it is typically installed as a server component, you can also install a personal mode on your local computer as an application. This provides connectivity to its data sources as well. For the purposes of this whitepaper, we focus only on the standard mode.

## Connecting Microsoft Power BI Desktop to AWS data sources

Most often customers who start with Microsoft Power BI Desktop are interested in how they can connect to AWS data sources from their on-premises computers and network. The desktop application is typically running on their local Windows laptop and physical and logical connectivity to AWS data sources are the biggest perceived barriers to entry.

However, another option exists, which is to run the Microsoft Power BI Desktop in the AWS Cloud. This option significantly reduces connectivity barriers to AWS data sources, but also requires some additional considerations. Both models are discussed in this chapter. We examine the implications of each in relation to connectivity, security, performance, and costs so that you can decide which option is best for you.

## Using Power BI Desktop on premises

If you plan on using Power BI desktop on premises with data sources that are stored within the AWS Cloud, Power BI can access these sources in one of three ways:

1. Connecting to data sources using the internet
2. Connecting to data sources using AWS VPN
3. Connecting to data sources using Direct Connect

Each method is detailed in the following sections. For alternate strategies, see [Using Microsoft Power BI Desktop within the AWS Cloud](#).

## Connecting to data sources through the internet

In this model, the Power BI Desktop application places an outbound connection that is routed over the internet to an IP address of an internet-accessible AWS data source. For example, Amazon RDS and Amazon Redshift, which are instantiated within a customer's Amazon VPC, support the *public accessibility* option to make instances accessible over the internet. Amazon Athena can be queried directly from the internet by using the service endpoint for your specific Region.




Figure 1 – Power BI connectivity to AWS data sources over the internet

Although this method of connectivity is technically possible, we don't recommend it for anything other than a small number of users. The following table lists important considerations.

Table 1 — Considerations for accessing AWS data sources over the internet

Criteria	Considerations for accessing AWS data sources over the internet
<b>Network connectivity</b>	Data sources must be accessible over the internet. Access must be either through a public IP address (Amazon RDS, Amazon Redshift, Amazon EC2 based data sources), or by a regional service endpoint (Amazon Athena).
<b>Security</b>	<p data-bbox="435 489 675 520"><b>IP access control</b></p> <p data-bbox="435 552 1425 682">A <i>security group</i> acts as a virtual firewall for your instance to control inbound and outbound traffic. In order to limit access to trusted entities, configure security groups to only allow inbound IP ranges associated with known CIDR ranges.</p> <p data-bbox="435 709 719 741"><b>Encryption in transit</b></p> <p data-bbox="435 772 1458 934">We recommend that you configure encryption for any data sources that use public IP addresses, such as Amazon RDS, Amazon Redshift, or any Amazon EC2-based data sources. This ensures that the risk of data or credentials being compromised while in transit, is reduced. Failure to configure encryption represents a significant risk. Do not overlook this aspect.</p> <p data-bbox="435 961 1425 1060">Regional service endpoints, such as Amazon Athena, are TLS encrypted. In addition, Amazon Athena query results that stream to JDBC or ODBC clients are encrypted using Transport Layer Security (TLS).</p> <p data-bbox="435 1087 889 1119"><b>Authentication and authorization</b></p> <p data-bbox="435 1150 1437 1213">We recommend that you use credentials that provide read-only access to datasets, and set up processes to rotate credentials per your company policy.</p>
<b>Performance</b>	<p data-bbox="435 1270 1393 1333">Some factors that might impact the overall Power BI Desktop performance when accessing AWS data sources over the internet include:</p> <ul data-bbox="483 1365 1458 1654" style="list-style-type: none"> <li data-bbox="483 1365 1458 1465">• The size of the dataset being accessed. Larger datasets take longer to retrieve. We recommend limiting queries and using filters to reduce the amount of data retrieved over the internet.</li> <li data-bbox="483 1491 1458 1654">• The quality of the internet connection, including bandwidth, latency, and packet loss. Where possible access data in AWS regions which you are geographically close to in order to reduce the effect of latency. If your internet is shared, consider loading data sources at off peak times and/or ensuring that enough bandwidth is available.</li> </ul> <p data-bbox="435 1680 1442 1743">In general, we recommend testing the experience at different times of the day, with different datasets, and with progressively larger number of users.</p>

Criteria	Considerations for accessing AWS data sources over the internet
<b>Cost</b>	Data sources that reside in a VPC and are queried using public IP address over the internet incur standard AWS VPC data egress charges. In order to reduce costs, we recommend limiting queries and using filters to reduce the amount of data retrieved over the internet.

## Connecting to data sources via AWS VPN

In this model, Power BI Desktop installations connect to data sources in the AWS network using one of two AWS VPN methods: AWS Site-to-Site VPN or AWS Client VPN. Each connection type delivers a highly available, managed, and elastic cloud VPN solution to protect your network traffic.

Site-to-Site VPN creates encrypted tunnels between your network and your Amazon Virtual Private Cloud (Amazon VPN) or AWS Transit Gateway. Client VPN connects your users to AWS or on-premises resources using a free VPN software client.

VPN traffic from both Site-to-Site VPN and Client VPN connections terminates in your VPC. As such, it can route to private IP addresses so your instances no longer need public-facing IP addresses. For services with a data path accessible from a publicly facing service endpoint, such as Athena, these service requests can either be routed over the internet, or over the VPN connection and through a VPC endpoint.




Figure 2 – Connecting Power BI Desktop to AWS data sources over Site-to-Site VPN and Client VPN

Site-to-Site VPN can also connect to AWS Transit Gateway, facilitating access to data sources spread across multiple VPCs.

Using AWS VPN provides the benefit of employing encryption when accessing data sources stored in AWS, without requiring that each data source to be explicitly configured. Once configured, VPN technology is largely seamless to end users.

Table 2 — Considerations for accessing AWS data sources using AWS VPN

Criteria	Considerations for accessing AWS data sources using AWS VPN
<b>Network connectivity</b>	<p>Data sources are available by connecting to private IP addresses in a VPC, or using a regional service endpoint. Power BI Desktop connects via VPN to the VPN Gateway and either access data sources directly (Amazon RDS, Amazon Redshift, Amazon EC2-based data sources), or by a VPC endpoint for services with a regional endpoint (Amazon Athena).</p>
<b>Security</b>	<p><b>IP access control</b></p> <p>You can use a combination of routing and security groups to control access to data sources stored in the AWS Cloud.</p> <p><b>Encryption in transit</b></p> <p>Both types of AWS VPN use <a href="#">IPsec</a> encryption, meaning that data transferred is encrypted as it travels between AWS and on premises. This ensures that even if data sources are not configured to use encrypted communications, that data is protected while traversing the internet.</p> <p><b>Authentication</b></p> <p>Site-to-Site VPN requires a one-time configuration and, once established, is seamless to users. End users are not required to authenticate to use the Site-to-Site VPN, but they require authentication to data sources.</p> <p>On the other hand, Client VPN does require authentication by the end users in order to establish the connection. Client VPN authentication can take place via Active Directory (user-based), mutual authentication (certificate-based), or SAML SSO (user-based). Once authenticated, the connection is seamless to the end user. AWS data sources added to Power BI Desktop require authentication.</p> <p>We recommend that you authenticate with AWS data sources using an identity that has read-only access only to the datasets required.</p>

Criteria	Considerations for accessing AWS data sources using AWS VPN
<b>Performance</b>	<p>The use of AWS VPN occurs over the internet. As such, its performance envelope is similar to the first scenario presented. Some factors can impact the overall Power BI Desktop performance when accessing AWS data sources over the internet. They include:</p> <ul style="list-style-type: none"> <li>• The size of the dataset being accessed. Larger datasets take longer to retrieve. We recommend limiting queries and using filters to reduce the amount of data retrieved over the internet.</li> <li>• The quality of the internet connection, including bandwidth, latency, and packet loss. Where possible, access data in AWS Regions that you are geographically close to. This reduces the effect of latency. If your internet is shared, consider loading data sources at off-peak times and/or ensuring that enough bandwidth is available.</li> </ul> <p>In general, we recommend testing the experience at different times of the day, with different datasets, and with progressively larger number of users.</p>
<b>Cost</b>	<p>Data sources that reside in a VPC and are queried using AWS VPN incur standard AWS VPN data transfer charges. To reduce costs, we recommend limiting queries and using filters to reduce the amount of data retrieved over the internet.</p>

## Connecting to data sources via AWS Direct Connect

AWS Direct Connect links your internal network to an AWS Direct Connect location over a standard Ethernet fiber-optic cable. One end of the cable is connected to your router, the other to an AWS Direct Connect router. With this connection, you can create virtual interfaces directly to public AWS services (for example, to Amazon S3) or to Amazon VPC, bypassing internet service providers in your network path. An AWS Direct Connect location provides access to AWS in the Region with which it is associated. You can use a single connection in a public Region or AWS GovCloud (US) to access public AWS services in all other public Regions.

In this model, the customer's on-premises network is connected through AWS Direct Connect directly to the AWS network. While there are multiple methods for configuring AWS Direct Connect, in its simplest mode, you are given access to IP ranges within a VPC using constructs known as a Private Virtual Interfaces (Private VIF). You access to the internet/public IP ranges using a Public Virtual Interface (Public VIF).




Figure 3 – Connecting Power BI Desktop to AWS data sources over AWS Direct Connect

When adding data sources in Power BI, you select the private IP address if it’s located in a VPC, or the public endpoint if it’s a regional service.

Table 3 — Considerations for accessing AWS data sources using AWS Direct Connect

Criteria	Considerations for accessing AWS data sources using AWS Direct Connect
<b>Network connectivity</b>	Once you configure Direct Connect, it can access data sources by connecting to private IPs in a VPC, or by using a regional service endpoint.
<b>Security</b>	<p><b>IP access control</b></p> <p>You can use a combination of routing and security groups to control access to data sources stored in the AWS Cloud.</p> <p><b>Encryption in transit</b></p> <p>Direct Connect does not provide line-level encryption of data. We recommend that you either enable TLS encryption at the data source level, or combine Site-to-Site VPN with Direct Connect to ensure that data and credentials are not compromised during transmission.</p> <p><b>Authentication</b></p>

Criteria	Considerations for accessing AWS data sources using AWS Direct Connect
	<p>Direct Connect does not require any additional authentication once configured.</p> <p>We recommend that you authenticate AWS data sources using an identity that has read-only access only to the datasets required.</p>
<b>Performance</b>	<p>With Direct Connect, 1 Gbps and 10 Gbps ports are available. You can order speeds of 50 Mbps, 100 Mbps, 200 Mbps, 300 Mbps, 400 Mbps, and 500 Mbps from any APN Partners that support Direct Connect.</p> <p>When compared to internet-based connections, Direct Connect connections typically provide increased bandwidth and reduced latency. Depending on the capacity of the connection, Power BI Desktop performance might no longer be constrained by the network when loading data sources. As such, when accessing datasets in the same Region that the Direct Connect is associated to, your users can expect good performance for queries.</p> <p>Be mindful of loading large datasets, and note that Power BI Desktop has a 10-GB dataset limit.</p>
<b>Cost</b>	<p>In addition to standard Direct Connect port charges, data sources that are accessed using a Direct Connect connection incur standard Direct Connect data transfer charges, as explained in <a href="#">AWS Direct Connect Pricing</a>. Data sent to the AWS Cloud over Direct Connect does not incur any charges.</p> <p>In order to reduce costs, we recommend limiting queries and using filters to reduce the amount of data retrieved.</p>

## Using Microsoft Power BI Desktop in the AWS Cloud

Using the Microsoft Power BI Desktop in the AWS Cloud is a popular solution for many of the challenges described in the previous section. In this model, customers host the Microsoft Power BI Desktop in the AWS Cloud, and then access it remotely on premises. The following diagram shows an example.




Figure 4 – Microsoft Power BI Desktop deployed in the AWS Cloud

Although the diagram depicts user connectivity to the desktop occurring through the internet, AWS VPN and Direct Connect are both valid connection types too. Also, because only graphical management traffic is transmitted, the bandwidth requirements are well suited for typical internet connections.

In this model the Microsoft Power BI Desktop is hosted within the Amazon VPC and has direct network connectivity to data sources with private IP addresses, such as Amazon RDS and Amazon Redshift. You can connect to Amazon Athena and other regional services by using a VPC endpoint connection as the destination (pictured in the diagram), or by using the regional public service endpoint.

There are a number of options for hosting Microsoft Power BI Desktop in the AWS Cloud:

### Option 1: Install Microsoft Power BI Desktop on an Amazon EC2 instance

In this option, you create one or more Amazon EC2 instances with Microsoft Windows Server installed. Your users connect to the instance using an RDP application and install, configure, and use Microsoft Power BI as they would as if it was installed on their laptop. By default, Windows Server installations are limited to two concurrent RDP client sessions. If you require more, you can buy Client Access Licenses from a Microsoft reseller. You might find that increasing the number of simultaneous users on a single

server causes resource contention. You may achieve a better user experience by increasing the number of instances, each with fewer users.

## Option 2: Install Microsoft Power BI in an Amazon WorkSpaces environment

Amazon WorkSpaces is a managed, secure Desktop-as-a-Service (DaaS) solution. You can use Amazon WorkSpaces to provision either Windows or Linux desktops in just a few minutes and quickly scale to provide thousands of desktops to workers across the globe.

In this option, you provision Microsoft Windows Desktops with Microsoft Power BI Desktop installed for your users. Users can access the Windows desktop using a light-weight client application that is available for Android, iOS, Fire, Mac, PC, Chromebook, and Linux devices.

Using Amazon WorkSpaces, Microsoft Power BI Desktop users have dedicated resources, and desktops can be automatically stopped to save costs. This option also has the advantage of allowing non-Windows devices to access Amazon WorkSpaces, as well as providing users a desktop environment for performing other tasks, such as accessing Microsoft Office or other Windows applications.

## Option 3: Install Microsoft Power BI in an Amazon AppStream 2.0 environment

Amazon AppStream 2.0 is a fully managed application streaming service. You centrally manage your desktop applications on AppStream 2.0 and securely deliver them to any computer. You can easily scale to any number of users across the globe without acquiring, provisioning, and operating hardware or infrastructure.

In this option, you provide your on-premises users just the Microsoft Power BI Desktop application, streamed to an HTML5-compliant web browser, with no plugins required. Like the Amazon WorkSpaces option, non-Windows devices can connect and users have access to dedicated resources.

*Table 4 — Considerations for Power BI Desktop running in the AWS Cloud*

Criteria	Considerations for Power BI Desktop running in the AWS Cloud
<b>Network connectivity</b>	Network connectivity to data sources is straight forward because both the data consumer and the data sources reside in the AWS Cloud. Data sources that live in an Amazon VPC, such as Amazon RDS and Amazon Redshift, can be accessed directly. Data sources that use regional endpoints can be accessed through the Amazon VPC internet gateway, or an Amazon VPC endpoint.

Criteria	Considerations for Power BI Desktop running in the AWS Cloud
	<p>Connectivity to the Microsoft Power BI Desktop occurs over the internet, AWS VPN, or AWS Direct Connect. All three options have modest requirements that most internet connections can meet.</p>
<b>Security</b>	<p><b>IP access control</b></p> <p>Customers can use a combination of routing and security groups to control access to data sources stored in the AWS Cloud.</p> <p>For the first option (using Amazon EC2), you can additionally use a combination of routing and security groups to allow only specific on-premises CIDR ranges.</p> <p><b>Encryption in transit</b></p> <p>We recommend that data sources within an Amazon VPC are configured to use encryption for transmission of data. Regional services already make use of TLS encryption.</p> <p>Management access is encrypted for all three options.</p> <p><b>Authentication</b></p> <p>We recommend that you authenticate with AWS data sources using an identity that has read-only access only to the datasets required.</p> <p>All three options require the use of an authentication mechanism in order for the remote user to prove their identity before accessing the environment that hosts Microsoft Power BI Desktop. This mechanism is typically a user name and password, but MFA options are available as well.</p>
<b>Performance</b>	<p>When running in the AWS Cloud, Microsoft Power BI Desktop performance is typically better than when run on premises. Both networking and compute instances are likely to be modern and have higher specifications, and there is likely to be low latency between the application and the data sources.</p>
<b>Cost</b>	<p>By placing Microsoft Power BI Desktop in the AWS Cloud, data transfer charges can be significantly reduced or removed entirely. Use care when accessing data sources across Availability Zones, in different VPC's, or in different Regions, because data transfer charges can apply.</p> <p>For each option, there are additional cost considerations that need to be factored in. Amazon EC2, Amazon WorkSpaces, and Amazon AppStream 2.0 are priced based on usage. For more information, see the pricing page for each service.</p>

## Summary of Microsoft Power BI Desktop connectivity options

For a small number of users with light dataset requirements, running Microsoft Power BI Desktop on premises and connecting securely over the internet, or using AWS VPN, might be an adequate solution. Make sure that security is configured and maintained in this model. We also recommend testing this configuration to determine if it meets users performance expectations

As the number of users increase, we recommend that you consider connectivity through AWS Direct Connect. Direct Connect provides a better user experience when loading larger datasets. Make sure that users are aware of the cost implications of transferring large datasets.

We recommend that you evaluate running Microsoft Power BI Desktop in the AWS Cloud. This is likely to provide both the best performance experience for the end user and the best management experience for cloud administrators. Solutions using Amazon WorkSpaces and Amazon AppStream 2.0 in particular can scale from a small number of users to thousands of users. These services also provide significant security and management benefits.

## Connecting the Microsoft Power BI service to AWS data sources

Connecting Microsoft Power BI service (SaaS) to private data sources in an Amazon VPC, requires an application component called Microsoft On-premises data gateway. The Microsoft On-premises data gateway is downloaded and installed on an Amazon EC2 instance in the VPC and configured with Microsoft Power BI credentials. The gateway establishes an outbound connection to the Azure Service Bus over the internet, and is configured in Microsoft Power BI to connect to data sources that it can access.

The Microsoft On-premises data gateway does not accept inbound connections from the Azure cloud, and only initiates outbound connections to the Azure Service Bus.<sup>3</sup>

## Recommended configuration

We recommend that you install the Microsoft On-premises data gateway on an Amazon EC2 instance in the private subnet that contains your data sources. This subnet is configured to route requests to the internet via an [Amazon VPC NAT gateway](#) installed in a public subnet. You can use a network address translation (NAT) gateway to enable instances in a private subnet to connect to the internet or to other AWS services, but prevent the internet from connecting to those instances. If you require a highly available

data gateway implementation, we recommend using a cluster of on-premise data gateways installed across multiple EC2 instances that span different AWS Availability Zones. For information, see [Add another gateway to create a cluster](#).




Figure 5: Connecting AWS data sources to Microsoft Power BI service

## Additional considerations

Table 5 — Considerations for Microsoft Power BI service with data sources in the AWS Cloud

Criteria	Considerations for Microsoft Power BI service with data sources in the AWS Cloud
<b>Network connectivity</b>	<p>Microsoft On-premises data gateway connectivity to data sources is straight forward because both the data consumer and the data sources reside within the AWS Cloud. Data sources that live in an Amazon VPC, such as Amazon RDS and Amazon Redshift, can be accessed directly. Data sources that use regional endpoints can be accessed through the Amazon VPC internet gateway, or by an Amazon VPC endpoint.</p> <p>Microsoft On-premises data gateway connectivity to the Microsoft Power BI service occurs over the internet and is an outbound connection only.</p>

Criteria	Considerations for Microsoft Power BI service with data sources in the AWS Cloud
<b>Security</b>	<p data-bbox="456 310 699 342"><b>IP access control</b></p> <p data-bbox="456 373 1464 436">You can use a combination of routing and security groups to control access to data sources stored within the AWS Cloud.</p> <p data-bbox="456 468 1464 636">Because Microsoft On-premises data gateway is installed on an Amazon EC2 instance, it will have an associated security group that can be used to limit inbound access to the operating system. The gateway does not accept inbound requests. The instance does not need a public IP address, and should not be configured with one.</p> <p data-bbox="456 667 743 699"><b>Encryption in transit</b></p> <p data-bbox="456 730 1448 825">We recommend that data sources within an Amazon VPC are configured to use encryption for transmission of data. Regional services already make use of TLS encryption.</p> <p data-bbox="456 856 1464 1014">Microsoft On-premises data gateway connectivity can be configured to connect to the Microsoft Azure Service Bus using HTTPS instead of TCP. We recommend using the HTTPS mode for communication. This is also the default for new gateway installations since the June 2019 gateway software version release.</p> <p data-bbox="456 1045 662 1077"><b>Authentication</b></p> <p data-bbox="456 1108 1425 1266">We recommend that you authenticate with AWS data sources using an identity that has read-only access only to the datasets required. The credentials that you enter for a data source are encrypted and stored in the gateway cloud service. The credentials are decrypted at the gateway on premises. <sup>4</sup></p> <p data-bbox="456 1297 1464 1392">Make sure that Microsoft Power BI credentials are securely controlled. Access to the service permits access to AWS data sources and potentially sensitive information they might contain.</p>
<b>Performance</b>	<p data-bbox="456 1455 1464 1549">Microsoft On-premises data gateway in the AWS Cloud typically performs well due to the ability to size and scale up the Amazon EC2 instance. It also performs fast in Region networking and connectivity to the internet.</p>
<b>Cost</b>	<p data-bbox="456 1612 1417 1675">Three factors need to be considered: Amazon EC2 instance charges, data transfer charges, and Amazon NAT gateway charges.</p> <p data-bbox="456 1707 1433 1801">Size your Amazon EC2 instances according to Microsoft's <a href="#">requirements</a>. To reduce costs, you can purchase Amazon EC2 Reserved Instances or <a href="#">AWS Savings Plans</a>.</p>

Criteria	Considerations for Microsoft Power BI service with data sources in the AWS Cloud
	<p>Data transferred from the Microsoft On-premises data gateway to the Microsoft BI service incurs VPC egress charges. We recommend that you limit queries and use filters to ensure that only relevant data is transferred.</p> <p>If the Microsoft On-premises data gateway connects to data sources in different Availability Zones or different AWS Regions, data transfer charges also apply.</p> <p>If the Microsoft On-premises data gateways are located in private subnets and make use of an AWS NAT gateway, hourly and data processing charges apply. For more information, see the <a href="#">Amazon VPC pricing page</a>.</p>

## Using Microsoft Power BI Report Server in AWS

Microsoft Power BI Report Server provides a private report server that Microsoft Power BI Desktop users can publish reports to and then share with a wider audience. Although it has traditionally been deployed on premises, Microsoft Power BI Report Server can reside within the AWS Cloud as well. This allows you to avoid hosting it in your data center while still making it accessible to both Microsoft Power BI Desktop users and the audience that needs to view published reports.

Microsoft Power BI Report Server can also connect to data sources that reside within the AWS Cloud. For a list of supported data sources, see [Power BI report data sources in Power BI Report Server](#).

There are two licensing options available for Power BI Report Server: Microsoft Power BI Premium and SQL Server Enterprise Edition with Software Assurance. Because of changes Microsoft made to their licensing terms in October 2019, Power BI Report Server can be deployed on EC2 Dedicated Hosts if the eligible SQL Server Enterprise Edition license with active Software Assurance was purchased prior to 10/1/2019 or added as a true-up under an active Enterprise Enrollment that was effective prior to 10/1/2019. In these specific BYOL scenarios, the licenses can only be upgraded to versions that were available prior to 10/1/2019. Note that this type of license is not included as part of any Amazon Machine Images (AMIs).

For more information about licensing, see [Licensing Power BI Report Server](#). You can deploy Microsoft Power BI Report Server in an Amazon EC2 instance that resides within a VPC subnet. The following diagram shows Microsoft Power BI Report Server deployed in a public subnet so that users can connect to it over the internet. If you have

AWS Site-to-Site VPN, or AWS Direct Connect, you can place Microsoft Power BI Report Server in a private subnet instead.




Figure 6 – Microsoft Power BI Report Server deployed into a public subnet

The following table outlines some common considerations when running Microsoft Power BI Report Server in the AWS Cloud.

Table 6 — Considerations for Microsoft Power BI Report Server in AWS

Criteria	Considerations for Microsoft Power BI Report Server in AWS
<b>Network connectivity</b>	<p>Before deploying the instance, you must choose how Microsoft Power BI Desktop users and the report audience will connect to the instance.</p> <p><b>For the report audience:</b></p> <ul style="list-style-type: none"> <li>If your organization has an AWS Direct Connect connection, or a Site-to-Site VPN, then we recommend that you deploy the instance in a private subnet and the report audience accesses it over those networks.</li> </ul>


Criteria	Considerations for Microsoft Power BI Report Server in AWS
	<ul style="list-style-type: none"> <li>If your organization does not have private connectivity, then the instance is accessed over the public internet. Deploy the instance in a public subnet with an Elastic IP address.</li> </ul> <p><b>For Microsoft Power BI Desktop users:</b></p> <ul style="list-style-type: none"> <li>If Microsoft Power BI Desktop is hosted on premises, it can connect to Power BI Report Server using the internet, Direct Connect, Site-to-Site VPN, or Client VPN.</li> <li>If Microsoft Power BI desktop is hosted in the AWS Cloud, it can connect to a report server in either a public or a private subnet using native AWS networking, such as the VPC local route, VPC peering, or AWS Transit Gateway.</li> </ul> <p><b>For AWS data sources:</b></p> <ul style="list-style-type: none"> <li>Because Microsoft Power BI Report Server lives within an Amazon VPC it can access AWS data sources using native AWS networking, such as local routes within an Amazon VPC, VPC Peering, or Amazon VPC endpoints.</li> </ul>
<p><b>Security</b></p>	<p><b>IP access control</b></p> <p>In order to limit access to trusted entities, you must configure the security group associated with the instance hosting Microsoft Power BI Report Server so that it only allows inbound IP ranges associated with trusted CIDR ranges.</p> <p><b>Encryption in transit</b></p> <p>We recommend that you configure Microsoft Power BI Report Server Web Service and web portal to use HTTPS. This ensures that connections from Microsoft Power BI Desktop and report viewers are encrypted in transit. For more information, see <a href="#">Configure TLS connections on a native mode report server</a>.</p> <p><b>Authentication and authorization</b></p> <p>Authentication and authorization mechanisms for Microsoft Power BI publishing and reporting viewing are contained in the application layer for Microsoft Power BI Report Server. This can be configured to use Microsoft Active Directory services, which can also live in the AWS Cloud.</p>
<p><b>Performance</b></p>	<p>Before deploying Microsoft Power BI Report Server in AWS, consider the following:</p>

Criteria	Considerations for Microsoft Power BI Report Server in AWS
	<ul style="list-style-type: none"> <li>The Amazon EC2 instance types should be sized in accordance with Microsoft's <a href="#">recommended hardware configuration</a>. For information about available Amazon EC2 instance types, see <a href="#">Amazon EC2 Instance Types</a>. To get started quickly we recommend that you narrow your focus to T3, T3a, C5, or C5a instance types, because these are likely to make good candidates.</li> <li>The experience of report users is also dependent on the quality of the network connection to the Microsoft Power BI Report Server. Report users are typically external to AWS and are using public or private networks. In cases where you have a large number of reporting users, make sure that there's sufficient bandwidth to handle peak requirements.</li> <li>The experience of report authors publishing from Microsoft Power BI Desktop is dependent on whether the software is running on premises or in the AWS Cloud. For on-premises installations, calculate how much bandwidth is required by estimating the number of authors and how much data is being published at peak times. Microsoft Power BI Desktop running in AWS (either in Amazon EC2, Amazon WorkSpaces, or Amazon AppStream 2.0) has high bandwidth and low latency connections, so network sizing is unlikely to be a factor.</li> </ul>
<b>Cost</b>	<p>Consider the following cost factors when deploying Microsoft Power BI Report Server in AWS:</p> <ul style="list-style-type: none"> <li>Amazon EC2 instance charges for the deployment of Microsoft Power BI Report Server. We encourage you to properly size your instances to get the best price and performance mix. Additionally, because Amazon EC2 instances associated with Microsoft Power BI Report server are likely to be running continuously, you should consider reducing costs further by purchasing <a href="#">Amazon EC2 Reserved Instances</a> or <a href="#">AWS Savings Plans</a>.</li> <li>Content viewed by reporting users over the internet is subject to standard Amazon VPC egress charges. If content is viewed over Site-to-Site VPN, or Direct Connect, those charges will apply instead.</li> </ul>

## Using Microsoft Power BI with Amazon QuickSight

Customers considering using the Microsoft Power BI Suite with AWS are encouraged to evaluate [Amazon QuickSight](#) as an alternative. This fully managed cloud service

natively connects to data sources in AWS, reducing the complexity and cost when compared to other BI solutions.



When compared to other BI solutions, QuickSight has the following benefits:

- With QuickSight, there's no need to download and install a client application. All functionality, including authoring and reporting, can be accessed from any platform (Windows, Mac, Linux, etc.) by a web browser.
- QuickSight is delivered as a fully managed, cloud-native SaaS application and is simple to build and deploy dashboards to production. The service is serverless, which means that you do not need to calculate how many nodes/servers you need to support your users. QuickSight also takes full advantage of high availability features provided by AWS for resiliency.
- It's easy to get started in small or large settings, with the ability to add users from a point-and-click interface within QuickSight. No external administrator intervention needed.
- QuickSight is powered by Super-fast, Parallel, In-memory Calculation Engine (SPICE) for a fast response time (in the milliseconds) and interactive visualizations. Datasets can currently scale up to 200 GB.
- QuickSight pricing is simple, inexpensive, and has two components; report authors and report readers. Report authors, who create and publish interactive dashboards, are priced per user. If users do not log in during a given month, there are no charges for those users. Report readers are charged per 30-minute session, with a maximum of \$5.00 per reader per month. A free trial allows you to evaluate QuickSight without any charges. For more information, see [Amazon QuickSight Pricing](#).

## Conclusion

If you're looking to use Microsoft Power BI Desktop, we generally find that customers start experimenting with the software on premises, connecting to data sources over the internet. While private connectivity options exist for using AWS VPN and Direct Connect, many customers have concluded that running Microsoft Power BI Desktop in the AWS Cloud provides a better performing experience. You can experiment between using Amazon EC2 instances, Amazon WorkSpaces, and Amazon AppStream 2.0 – and a mix of options is often employed within the same customer.

If you want to connect data sources in AWS to Microsoft Power BI Service, you should feel comfortable knowing that this is an established architectural pattern. You can install the Microsoft On-premises data gateway within an Amazon VPC and connect data sources such as Amazon RDS, Amazon Redshift, and Amazon Athena seamlessly to the service.

If you plan on using Microsoft Power BI Report Server in AWS, there is also an established path forward. You can install the Microsoft Power BI Report Server within an Amazon VPC, close to AWS data sources, and there are connectivity options for both authors and report users.

If you want a solution that provides the same business outcomes, without the added complexity of installing, configuring, patching, and scaling self-managed BI solutions, we recommend Amazon QuickSight. This fully managed service combines all the required functionality in a simple web browser experience with pay-per-user pricing. There is nothing to install and no additional components are required.

Hopefully, this is just the start of your business intelligence journey with AWS. For additional resources to help you get started, see [Appendix: Microsoft Power BI supported AWS data sources](#). If you need hands-on help or have additional questions, [contact us](#).

## Contributors

Contributors to this document include:

- Ralph Holm, Sr. Solutions Architect, Amazon Web Services
- Barret Newman, Sr. Solutions Architect, Amazon Web Services
- Fabrizio Napolitano, Specialist Solutions Architect for Database and Analytics, Amazon Web Services

## Further reading

- [Integrate Power BI with Amazon Redshift for insights and analytics](#)
- [Connect to an Amazon Redshift database in Power BI Desktop](#)

## Document Revisions

Date	Description
November 2020	First publication
January 2021	Updated Report Server Licensing

## Appendix: Microsoft Power BI supported AWS data sources

When considering the list of supported data sources keep the following points in mind:

- Amazon Redshift is natively supported as a Power BI data source.
- Amazon Relational Database Service provides multiple database engines including SQL Server, MariaDB, MySQL, Oracle Database and PostgreSQL. Note that the database engines are listed, not the Amazon RDS service.
- Amazon Athena uses standard SQL and can be queried using an ODBC or JDBC connection, so it is not specifically listed. Drivers are provided by AWS for common operating systems. See [Connecting to Amazon Athena with ODBC and JDBC Drivers](#) to get started.

The full list of supported data sources is provided by Microsoft, see [Power BI data sources](#).

## Notes

<sup>1</sup> [What is the Power BI service?](#)

<sup>2</sup> [Comparing Power BI Desktop and the Power BI service](#)

<sup>3</sup> [On-premises data gateway FAQ](#)

<sup>4</sup> The credentials that you enter for a data source are encrypted and stored in the gateway cloud service. The credentials are decrypted at the gateway on premises.