Deploy Remote Desktop Gateway on the AWS Cloud

Santiago Cardenas

April 2014
Last updated: July 2016 (revisions)

This guide is also available in HTML format at http://docs.aws.amazon.com/quickstart/latest/rd-gateway/.

Table of Contents

Abstract 3
Before You Get Started 3
Three Ways to Use this Guide 3
Considerations When Deploying RD Gateway 4
The Principle of Least Privilege 4
Amazon Virtual Private Cloud (Amazon VPC) 4
Network Access Control Lists 5
Security Groups 5
RD Gateway Setup 7
Initial Remote Administration Architecture 7
RD Gateway Installation 7
SSL Certificates 8
Connection and Resource Authorization Policies 10
RD Gateway Architecture on the AWS Cloud 13
Client Configuration 14
Installing the Root Certificate 14
Name Resolution 15
Configuring the Remote Desktop Connection Client 15
Automated Deployment 18
Scenario 1: Deploy RD Gateway into a new Amazon VPC 18
Scenario 2: Deploy Standalone RD Gateway into an Existing VPC 20
Scenario 3: Deploy Domain-Joined RD Gateway into an Existing VPC 20
RD Gateway Deployment Checklist 21
Further Reading 22
Appendix 23
Send Us Your Feedback 24
Document Revisions 24

Abstract

This reference deployment guide includes architectural considerations and configuration steps for deploying Remote Desktop Gateway (RD Gateway) on the Amazon Web Services (AWS) cloud. We'll discuss best practices for securely accessing your Windows-based instances using the Remote Desktop Protocol (RDP) for remote administration. We also provide links to automated AWS CloudFormation templates that you can leverage for your deployment or launch directly into your AWS account.

Amazon Web Services provides a comprehensive set of services and tools for deploying Microsoft Windows-based workloads on its highly reliable and secure cloud infrastructure. The RD Gateway uses the Remote Desktop Protocol (RDP) over HTTPS to establish a secure, encrypted connection between remote users on the Internet and Windows-based, Amazon EC2 instances, without needing to configure a virtual private network (VPN) connection. This allows you to reduce the attack surface on your Windows-based instances while providing a remote administration solution for administrators.

This guide is aimed at organizations running workloads in the AWS cloud that require secure remote administrative access to Windows-based, Amazon EC2 instances over the Internet. After reading this guide, IT infrastructure personnel should have a good understanding of how to design and deploy an RD Gateway infrastructure on the AWS cloud.

Before You Get Started

Implementing the RD Gateway on the AWS cloud is an advanced topic. If you are new to AWS, see the Getting Started section of the AWS documentation. In addition, familiarity with the following technologies is recommended:

- Amazon EC2
- Amazon VPC
- Windows Server 2012 or 2008 R2
- Remote Windows Administration using the Remote Desktop Protocol (RDP)

This guide focuses on infrastructure configuration topics that require careful consideration when you are planning and deploying an RD Gateway infrastructure on the AWS cloud. We don't cover general Windows Server installation and software configuration tasks. For general software configuration guidance and best practices, consult the Microsoft product documentation.

Three Ways to Use this Guide

- This guide provides considerations and deployment steps for running the RD Gateway role in the AWS cloud. The examples and steps provided in this document can be used to deploy this infrastructure from scratch, allowing you to securely administer your Windows-based, Amazon EC2 fleet using RDP over HTTPS. You can follow these steps in your own existing environment to deploy a fully configured RD Gateway infrastructure.

- You can use the provided AWS CloudFormation templates in this guide to help automate the deployment of your RD Gateway infrastructure.
- We have published a set of Quick Starts that provide solutions for deploying common Microsoft workloads on AWS, including Microsoft Active Directory, Microsoft SharePoint, Microsoft Exchange, and Microsoft SQL Server. These reference deployments include RD Gateways that you can deploy using AWS CloudFormation templates. This guide describes the architecture used in those Quick Starts for remote administration, and the implementation steps used to deploy the solution.

Considerations When Deploying RD Gateway

The Principle of Least Privilege

When considering remote administrative access to your environment, it is important to follow the principle of "least privilege." This principle refers to users having the least possible privilege necessary to perform their job functions. This helps reduce the attack surface of your environment, making it much harder for an adversary to exploit. An attack surface can be defined as the set of exploitable vulnerabilities in your environment, including the network, software, and users who are involved in the ongoing operation of the system.

Following the principle of least privilege, we recommend reducing the attack surface of your environment by exposing the absolute minimal set of ports to the network while also restricting the source network or IP address that will have access to your Amazon EC2 instances.

There are several AWS capabilities, alongside the functionality that exists in the Microsoft platform, to help you implement the principle of least privilege.

Amazon Virtual Private Cloud (Amazon VPC)

Amazon VPC lets you provision a private, isolated section of the AWS cloud where you can launch AWS resources in a virtual network that you define. With Amazon VPC, you can define a virtual network topology closely resembling a traditional network that you might operate on your own premises. 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.

When deploying a Windows-based architecture on the AWS cloud, we recommend an Amazon VPC configuration that supports the following requirements:

- Critical workloads should be placed in a minimum of two Availability Zones to provide high availability.
- Instances should be placed into individual tiers. For example, in a Microsoft SharePoint deployment, you should have separate tiers for web servers, application servers, database servers, and Domain Controllers. Traffic between these groups can be controlled to adhere to the principle of least privilege.
- Internal application servers and other non-Internet facing servers should be placed in private subnets to prevent direct access to these instances from the Internet.
- RD Gateways should be deployed into public subnets in each Availability Zone for remote administration. Other components, such as reverse proxy servers, can also be placed into these public subnets if needed.

For details on the Amazon VPC design used in this reference, see the Quick Start for building a modular and scalable virtual network architecture with Amazon VPC.

Network Access Control Lists

A network access control list (ACL) is a set of permissions that can be attached to any network subnet in an Amazon VPC to provide stateless filtering of traffic. Network ACLs can be used for inbound or outbound traffic and provide an effective way to blacklist a CIDR block or individual IP addresses. These ACLs can contain ordered rules to allow or deny traffic based upon IP protocol, service port, or source or destination IP address. Figure 1 shows the default ACL configuration for an Amazon VPC subnet.

Network ACL: Default (replace)
Inbound:

Rule #	Port (Service)	Protocol	Source	Allow/Deny
100	ALL	ALL	0.0.0.0/0	ALLOW
*	ALL	ALL	0.0.0.0/0	DENY

Outbound:

Rule #	Port (Service)	Protocol	Destination	Allow/Deny
100	ALL	ALL	0.0.0.0/0	ALLOW
*	ALL	ALL	0.0.0.0/0	DENY

Figure 1: Default Network ACL Configuration for an Amazon VPC Subnet

You may choose to keep the default network ACL configuration, or you may choose to lock it down with more specific rules to restrict traffic between subnets at the network level. For example, you could set a rule that would allow inbound administrative traffic on TCP port 3389 from a specific set of IP addresses. In either case, you'll also need to implement Security Group rules to permit access from users connecting to RD Gateways and between tiered groups of Amazon EC2 instances.

Security Groups

All Amazon EC2 instances are required to belong to one or more security groups. Security groups allow you to set policies to control open ports and provide isolation between application tiers. In an Amazon VPC, every instance runs behind a stateful firewall with all ports closed by default. The Security Group contains rules responsible for opening inbound and outbound ports on that firewall. While Security Groups act as an instance-level firewall, they can also be associated with multiple instances, providing isolation between application tiers in your environment. For example, you can create a Security Group for all of your web servers that will allow traffic on TCP port 3389, but only from members of the Security Group containing your RD Gateway servers.

Figure 2 shows an RD Gateway architecture running on the AWS cloud that is aligned with AWS best practices.

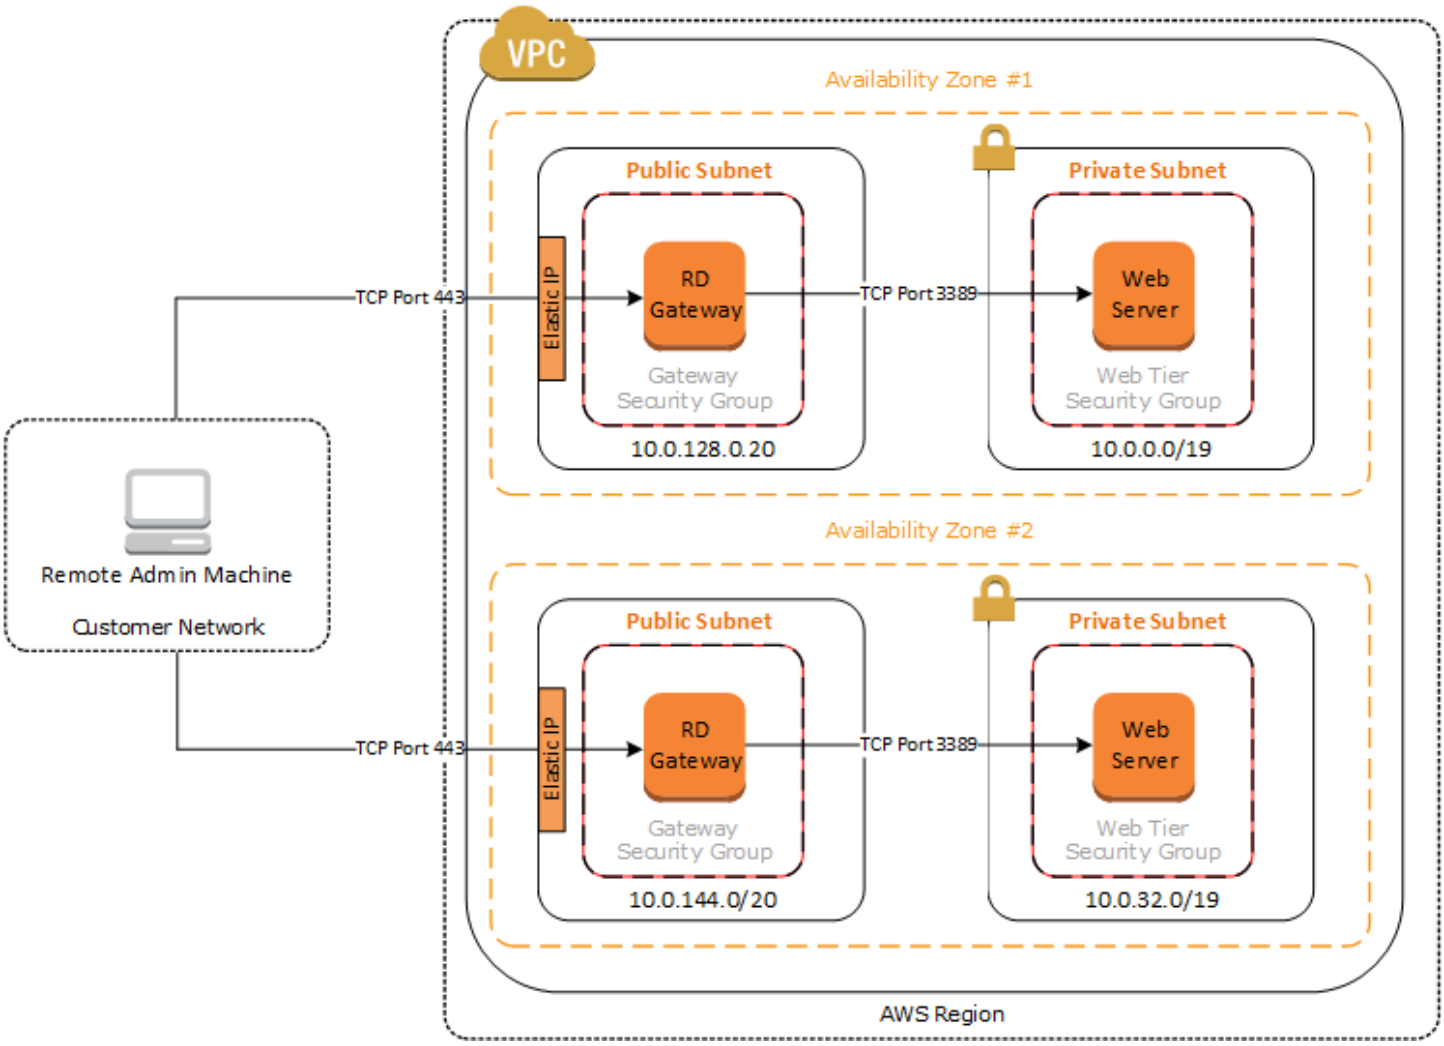

Figure 2: RD Gateway Architecture on the AWS Cloud

Notice that inbound connections from the Internet are only permitted over TCP port 443 to the RD Gateways. The RD Gateways have an Elastic IP address (EIP) assigned and have direct access to the Internet. The remaining Windows instances are deployed into private subnets and are assigned private IP addresses only. Security Group rules allow only the RD Gateways to initiate inbound connections for remote administration to TCP port 3389 for instances in the private subnets.

In this architecture, RDP connections are established over HTTPS to the RD Gateway and proxied to backend instances on the standard RDP TCP port 3389. This configuration allows you to reduce the attack surface on your Windows-based instances while allowing administrators to establish connections to all of your instances through a single gateway.

While it is possible to provide remote administrative access to all of your Windows-based instances through one RD Gateway, we recommend placing gateways in each Availability Zone for redundancy.

RD Gateway Setup

Initial Remote Administration Architecture

When you initially configure your RD Gateways, the servers in the public subnet will need an inbound Security Group rule permitting TCP port 3389 from the administrator's source IP address or subnet. Windows instances sitting behind the RD Gateway in a private subnet should be in their own isolated tier. For example, a group of web server instances in a private subnet may be associated with their own "web tier" security group. This security group will need an inbound rule allowing connections from the RD Gateway on TCP port 3389.

Using this architecture, an administrator can use a traditional RDP connection to an RD Gateway to configure the local server. The RD Gateway can also be used as a "jump box"; once an RDP connection is established to the desktop of the RD Gateway, an administrator can start a new RDP client session to initiate a connection to an instance in a private subnet.

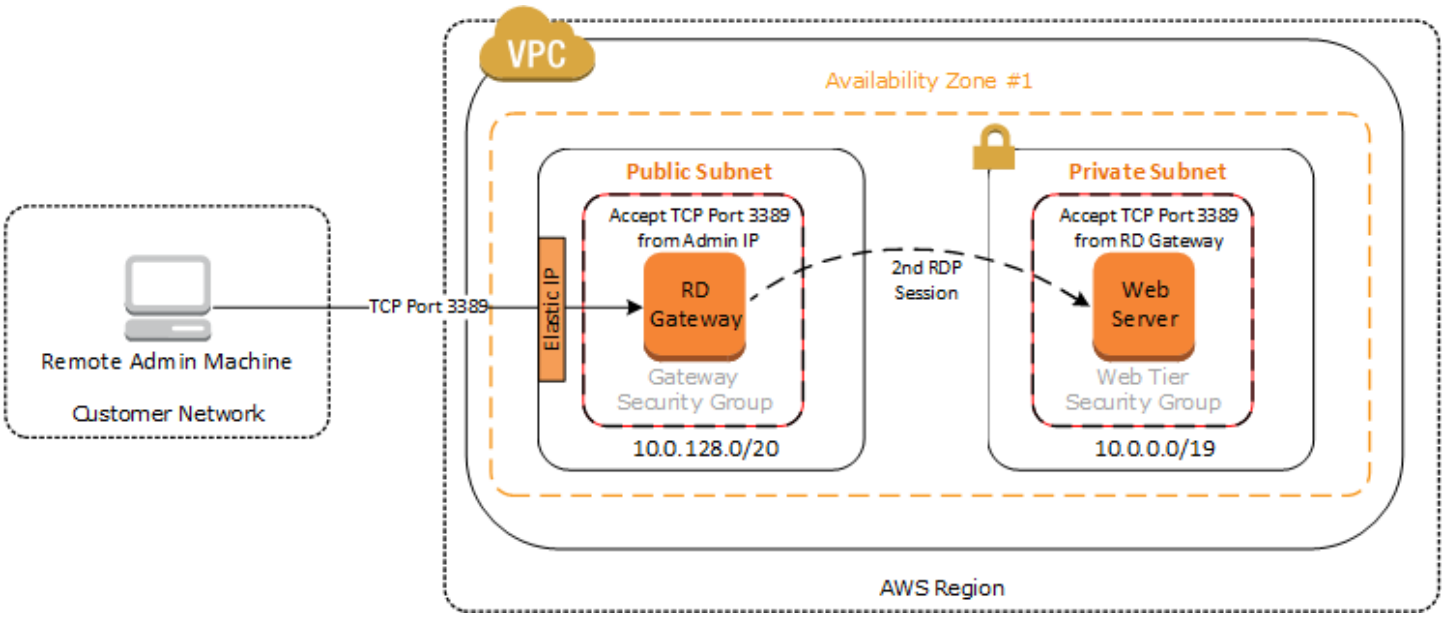

Figure 3: Initial Architecture for Remote Administration

While this architecture works well for initial administration, it is not recommended for the long term. To further secure connections and reduce the number of RDP sessions required to administer the servers in the private subnets, the RD Gateway service should be installed and configured with an SSL certificate and Connection and Authorization policies.

RD Gateway Installation

The installation of the RD Gateway role is very straightforward. This can be performed from the Server Manager or with a single PowerShell command on Windows Server 2012.

```
Install-WindowsFeature RDS-Gateway -IncludeManagementTools
```

The above command should be run from a PowerShell instance started with administrative privileges. Once complete, the RD Gateway role, along with all pre-requisite software and administration tools, will be installed on your Windows Server 2012, Amazon EC2 instance.

For Windows Server 2008 R2 based installations, we recommend following the detailed installation instructions in the Remote Desktop Services documentation (Microsoft TechNet Library).

SSL Certificates

The RD Gateway role uses Transport Layer Security (TLS) to encrypt communications over the Internet between administrators and gateway servers. To support TLS, a valid X.509 SSL certificate must be installed on each RD Gateway. Certificates can be acquired in a number of ways, including the following common options:

- Your own PKI infrastructure, such as a Microsoft Enterprise Certificate Authority (CA)
- Certificates issued by a public CA, such as Verisign or Digicert
- Self-Signed Certificates

For smaller test environments, implementing a self-signed certificate is a straightforward process that allows you to get up and running quickly. However, if you have a large number of varying administrative devices that need to establish a connection to your gateways, we recommend using a public certificate.

In order for an RDP client to establish a secure connection with an RD Gateway, the following certificate and DNS requirements must be met:

- The issuing CA of the certificate installed on the gateway must be trusted by the RDP client. For example, the root CA certificate must be installed in the client machine's "Trusted Root Certification Authorities" store.
- The subject name used on the certificate installed on the gateway must match the DNS name used by the client to connect to the server: for example, rdgw1.example.com.
- The client must be able to resolve the host name (for example, rdgw1.example.com) to the EIP of the RD Gateway. This will require a Host (A) record in DNS.

There are various considerations when choosing the right CA to obtain an SSL certificate. For example, a public certificate may be ideal since the issuing CA will be widely trusted by the majority of client devices that need to connect to your gateways. On the other hand, you may choose to utilize your own PKI infrastructure to ensure that only the machines that are part of your organization will trust the issuing CA.

Implementing a Self-Signed Certificate

If you choose a self-signed certificate, you will need to install the root CA certificate on every client device. Keep in mind that in order to provide an automated solution, the CloudFormation templates provided in this guide utilize a self-signed certificate for the RD Gateway service. If you are not using the automated deployment, you can follow the steps below to generate a self-signed certificate.

The RD Gateway management tools provide a mechanism for generating a self-signed certificate.

1. To install a self-signed certificate, launch the **RD Gateway Manager.**
2. Right-click on the local server name, and select **Properties.**

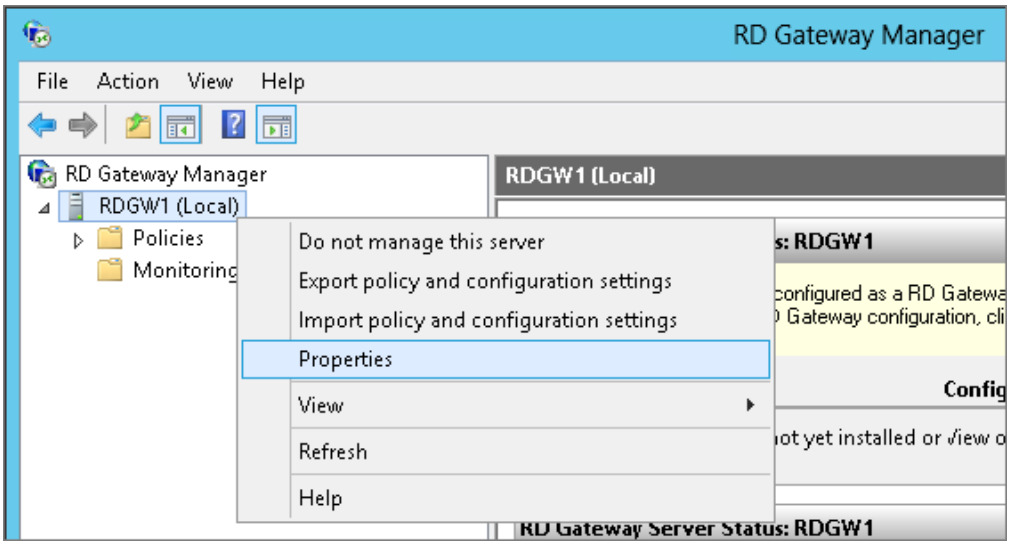

Figure 4: Navigating the RD Gateway Manager

3. On the SSL Certificate tab, ensure that **Create a self-signed certificate** is selected and click **Create and Import a Certificate**.

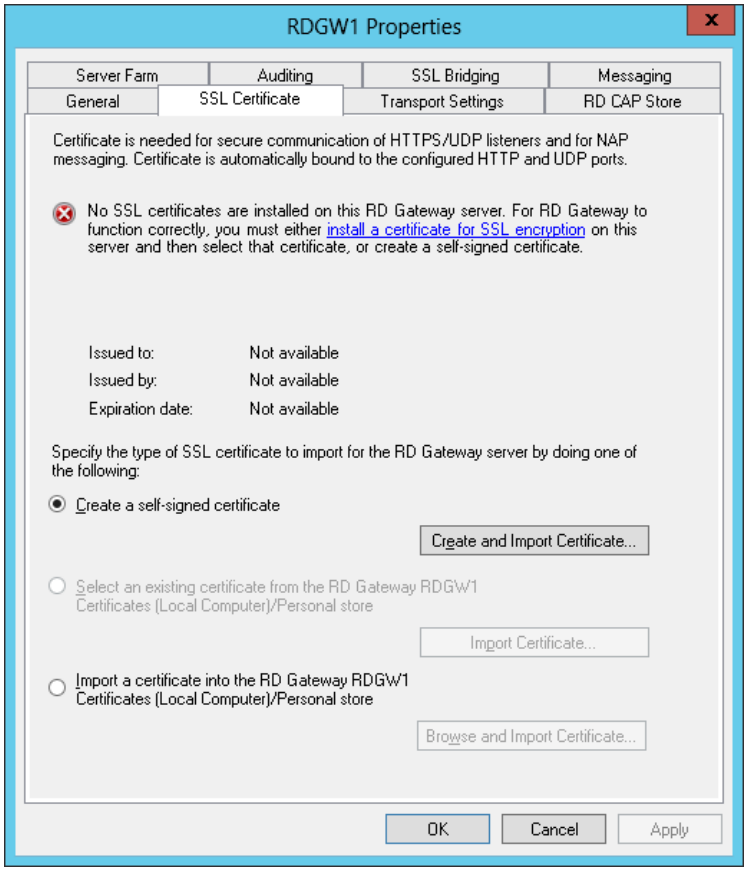

Figure 5: SSL Certificate Settings on the RD Gateway

4. Ensure that the correct fully-qualified domain name (FQDN) is listed for the **Certificate name**. Make note of the root certificate location and Click **OK.**

Figure 6: Creating a Self-Signed Certificate

5. After installing the certificate, closing and reopening the server's **Properties** dialog box will show the new self-signed certificate successfully installed.

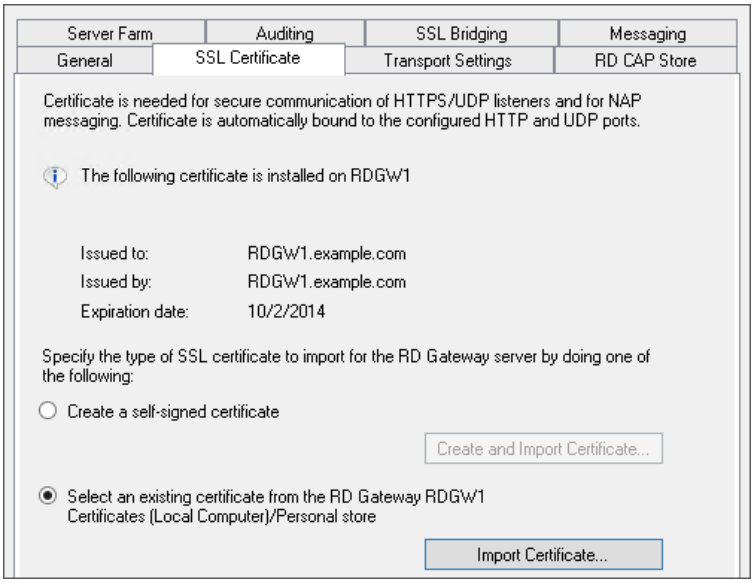

Figure 7: Viewing the SSL Certificate Settings After Creating a New Certificate

Connection and Resource Authorization Policies

Once you've installed the RD Gateway role and an SSL certificate, you are ready to configure Connection and Resource Authorization policies.

- **Connection Authorization Policies**—Remote Desktop connection authorization policies (RD CAPs) allow you to specify who can connect to an RD Gateway instance. For example, you can select a group of users from your domain, such as "Domain Admins."

- **Resource Authorization Policies**—Remote Desktop resource authorization policies (RD RAPs) allow you to specify the internal Windows-based instances that remote users can connect to through an RD Gateway instance. For example, you can choose specific domain-joined computers which administrators can connect to through the RD Gateway.

1. To configure the policies, launch the **RD Gateway Manager**.

2. Right-click the **Policies** branch and select **Create New Authorization Policies**.

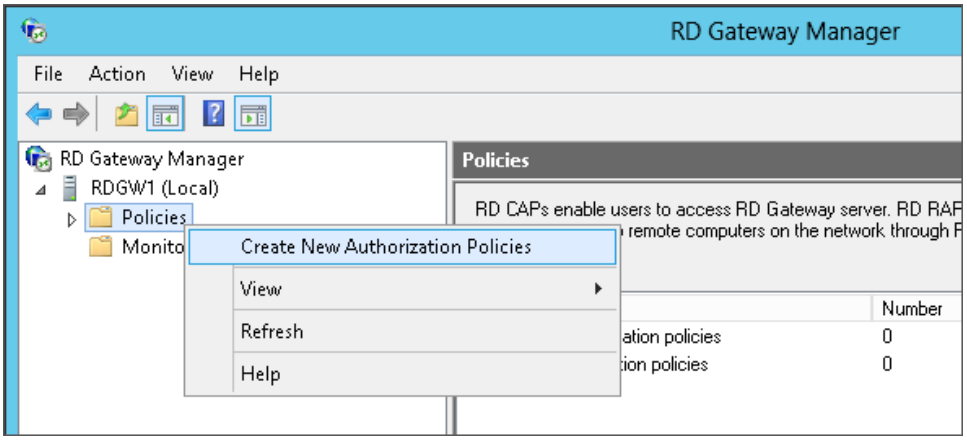

Figure 8: Navigating the RD Gateway Manager

3. In the **Create New Authorization Policies** Wizard, select **Create a RD CAP and a RD RAP (recommended)** and click **Next**.

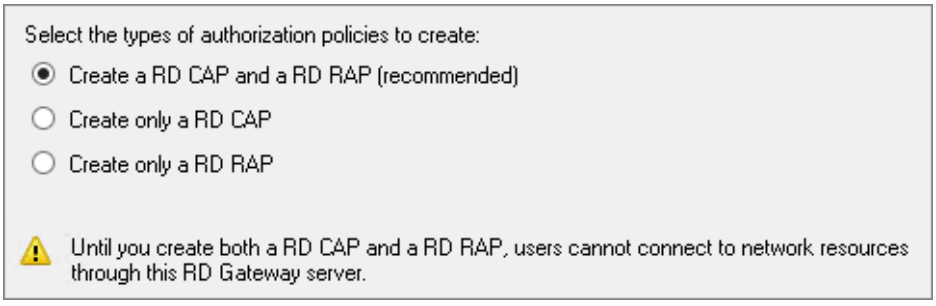

Figure 9: Select Authorization Policies

4. Provide a friendly name for your RD CAP and click **Next**.

5. On the Select Requirements screen, define the authentication method and groups that should be permitted to connect to the RD Gateway and click **Next**.

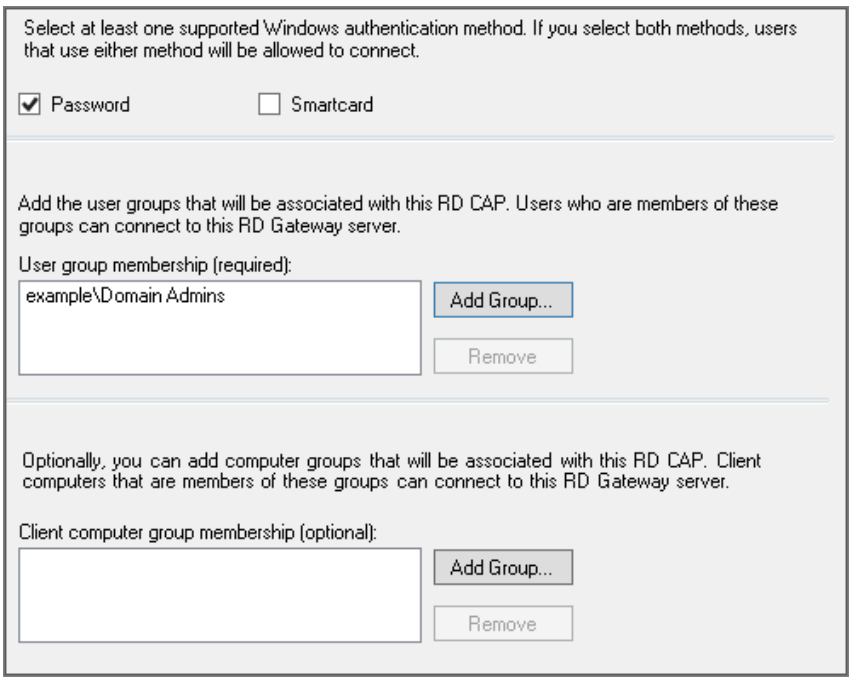

Figure 10: Configure Authentication Method and Groups for RD CAP

6. Choose whether to enable or disable Device Redirection and click **Next**.
7. Specify your timeout and reconnection settings and click **Next**.
8. On the RD CAP Settings Summary screen click **Next**.
9. Provide a friendly name for your RD RAP and click **Next**.
10. Select the user groups that will be associated with the RAP and click **Next**.

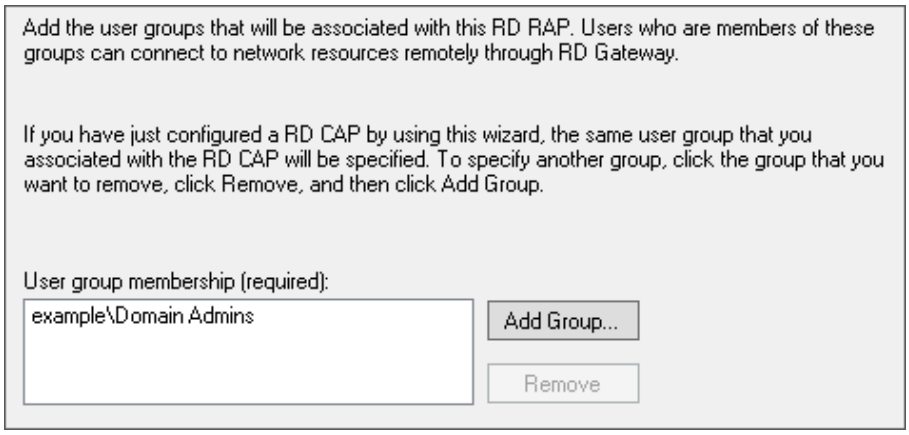

Figure 11: Select Group Memberships for RD RAP

11. Select the Windows-based instances (network resources) that administrators should be able to connect to through the RD Gateway. This can be a security group in AD containing specific computers. For this example, we'll allow administrators to connect to any computer. Click **Next**.

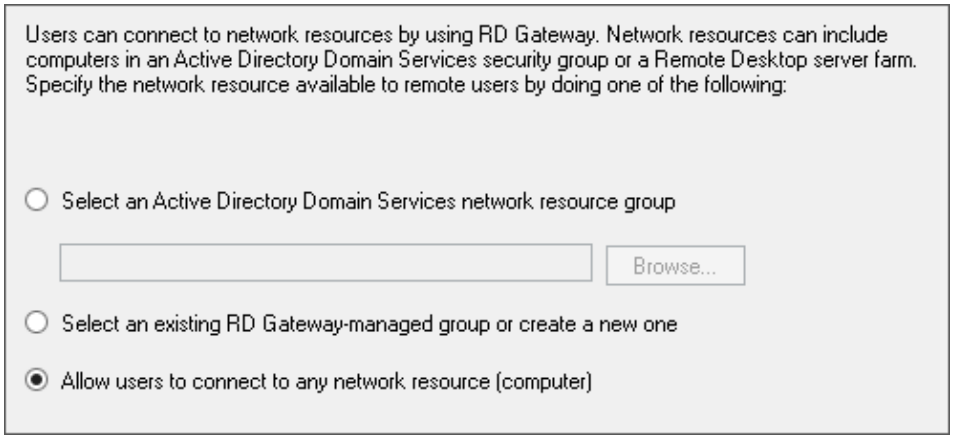

Figure 12: Select Network Resources

12. Allow connections to TCP port 3389 and click **Next**.

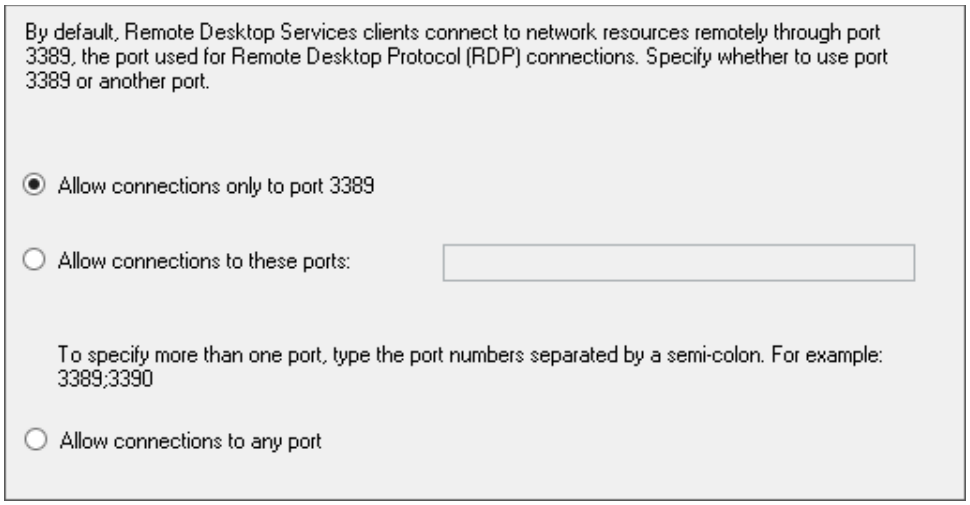

Figure 13: Select RDP Port

13. Click **Finish** and click **Close.**

RD Gateway Architecture on the AWS Cloud

After you configure Connection and Resource Authorization policies, you can modify the Security Group for RD Gateway to use a single inbound rule permitting TCP port 443. This modification will allow a Transport Layer Security (TLS) encrypted RDP connection to be proxied through the gateway over TCP port 443 directly to one or more Windows-based instances in private subnets on TCP port 3389. This configuration increases the security of the connection and also prevents the need to initiate an RDP session to the desktop of the RD Gateway.

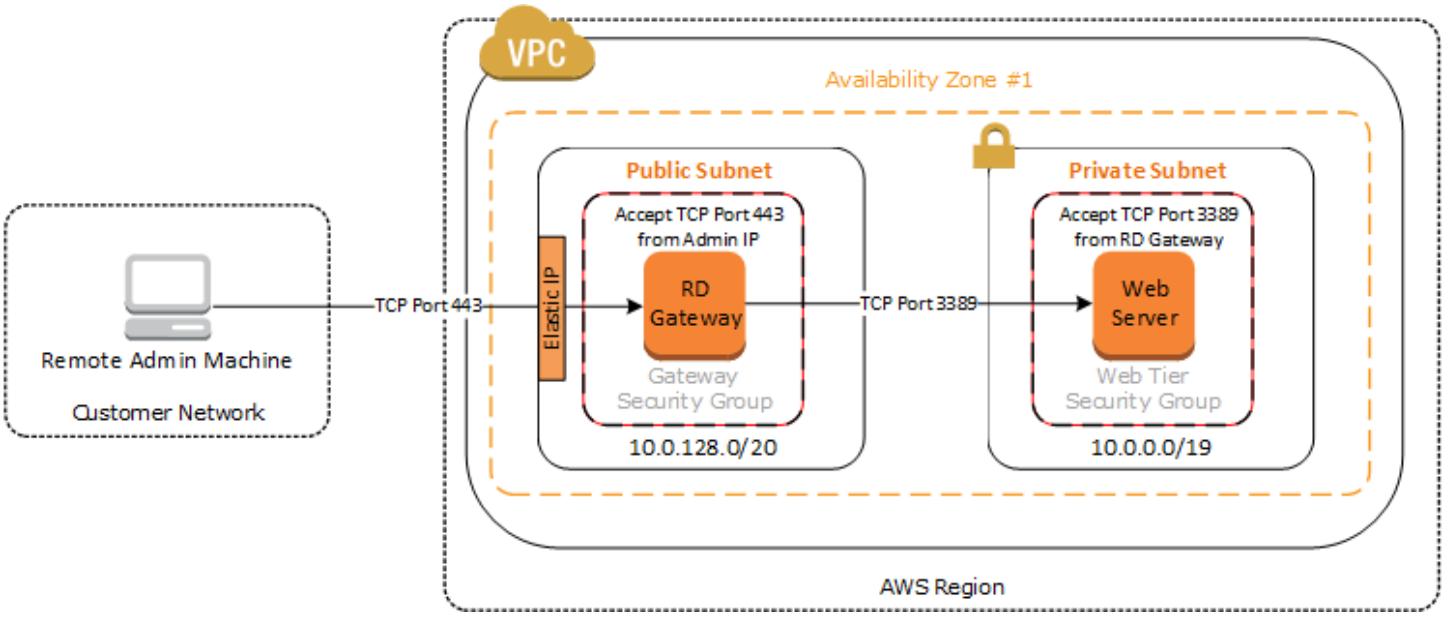

Figure 14: Architecture for RD Gateway Administrative Access

Client Configuration

Configuring your administrative clients requires the installation of any root certificates, name resolution for the RD Gateway FQDN, and proper configuration of the Remote Desktop Connection client.

Installing the Root Certificate

If you are using a self-signed certificate on the RD Gateway, you must install the root certificate on your administrative clients before you configure the RDP client to connect to your RD Gateway. As you'll recall from the section in this guide on Implementing a Self-Signed Certificate, the root certificate will automatically be stored as *c:\users\<username>\documents\<servername>.cer.*

This file can be distributed to administrator workstations and installed using the following steps:

1. Open a **Command Prompt** window using administrative credentials.
2. Type **mmc** and press **Enter**.
3. On the **File** menu, click **Add/Remove Snap In.**
4. Click **Add.**
5. In the **Add Standalone Snap-in** dialog box, select **Certificates.**
6. Click **Add.**
7. In the **Certificates** snap-in dialog box, select **Computer account** and click Next.
8. In the **Select Computer** dialog box, click **Finish.**

9. In the **Add Standalone Snap-in** dialog box, click **Close.**
10. On the **Add/Remove Snap-in** dialog box, click **OK.**
11. In the Console Root window, expand **Certificates (Local Computer)**.
12. Under **Certificates (Local Computer)**, expand **Trusted Root Certification Authorities.**
13. Right click on **Certificates**, and select **All Tasks** > **Import**.
14. Navigate to the root certificate (e.g. RDGW1.cer) to complete the installation.

Note: The root certificate will be stored as *c:\<servername>.cer* on each RD Gateway when deploying servers using the CloudFormation templates.

Name Resolution

Make sure that your administrative clients can resolve the FQDN of your RD Gateway. You can create an A (Host) record in DNS that maps the FQDN to the RD Gateways EIP or Public IP address. For testing purposes, you can configure this mapping in the local host's file on the machine.

Configuring the Remote Desktop Connection Client

Use the following steps to configure the Remote Desktop Connection on admin clients.

1. Start the Remote Desktop Connection client. In the computer name field, type the name or IP address of the Windows instance you want to connect to. Keep in mind that this instance only needs to be reachable from the RD Gateway, not from the client machine.

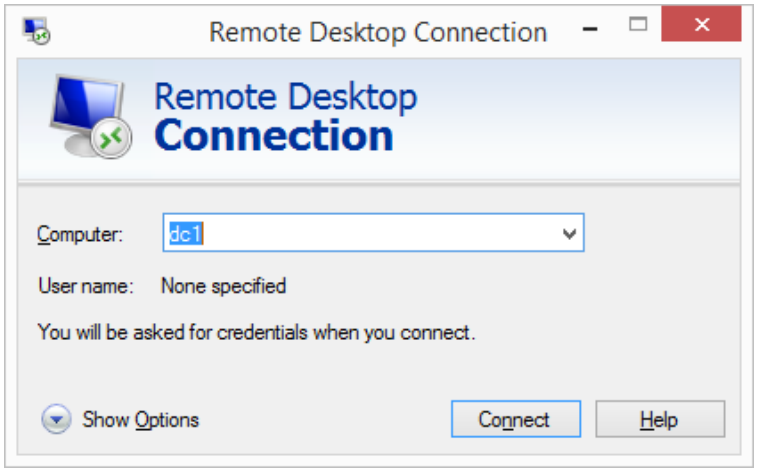

Figure 15: The Remote Desktop Connection Client

2. Click on **Show Options**, and select the **Advanced** tab.

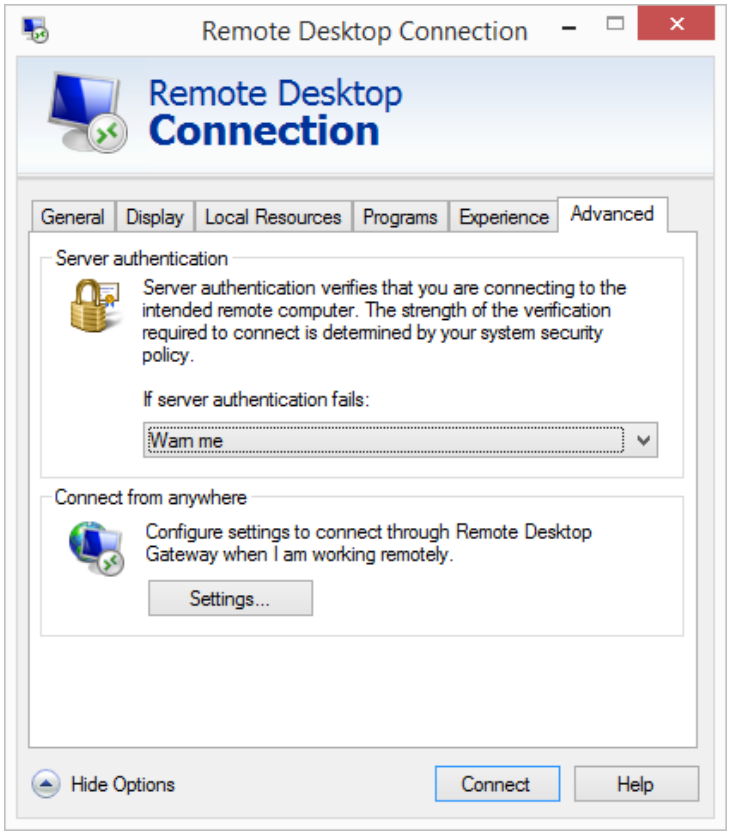

Figure 16: Advanced Properties for the Remote Desktop Connection Client

3. Click the **Settings** button, select **Use these RD Gateway server settings,** and specify the FQDN of the RD Gateway in the **Server name** field. If the RD Gateway and the server you want to connect to are in the same domain, select **Use My RD Gateway credentials for the remote computer** and click **OK.**

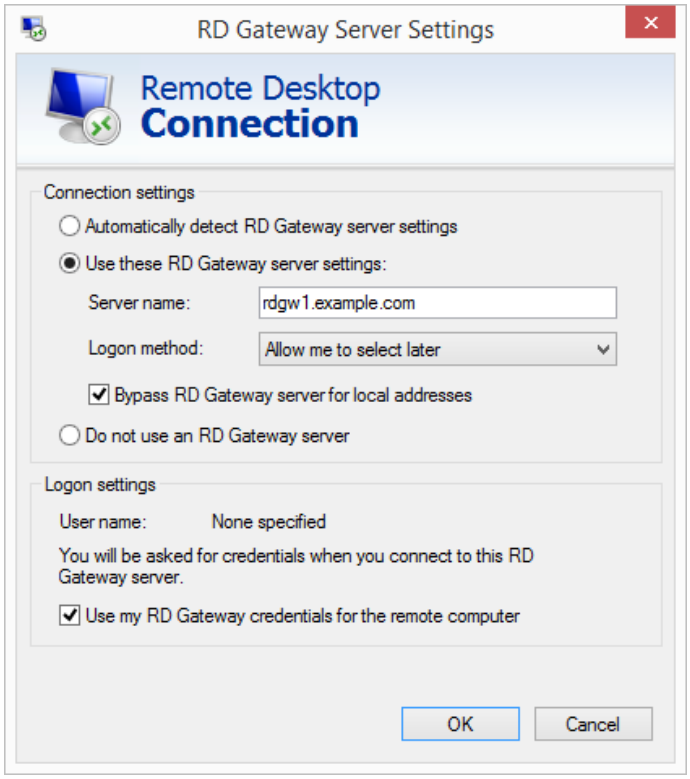

Figure 17: Advanced Settings of the Remote Desktop Connection Client

4. You can supply the same set of credentials for the RD Gateway and Destination server. Notice that in Figure 18 the dialog box shows that these credentials will be used for both the RD Gateway and for the destination server. When the RD Gateway and destination computer are in the same Active Directory domain, a single set of credentials should work fine. If your servers are not domain joined, you will need to authenticate twice: once for the RD Gateway and once for the destination server. When you click **OK**, the RDP connection to the server will be initiated.

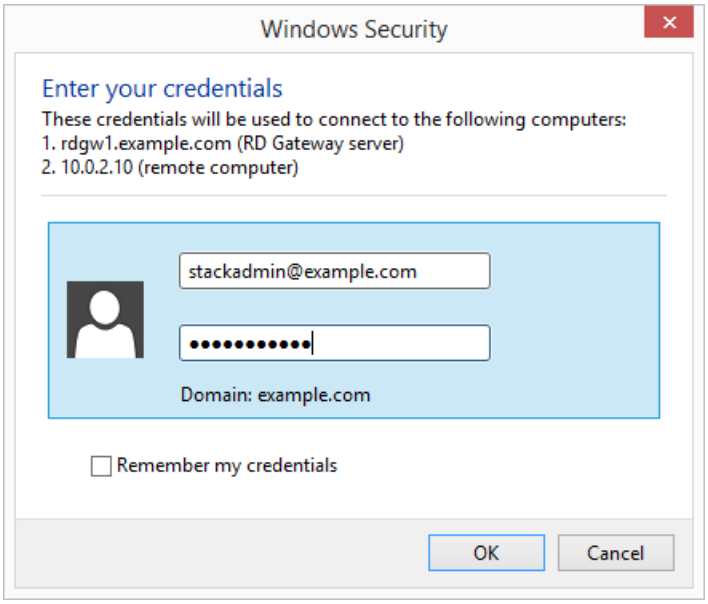

Figure 18: Providing Credentials to the RD Gateway and Destination Server

Automated Deployment

While AWS provides existing reference deployments that include RD Gateway servers, we realize that you may want to deploy your solution using another approach. For example, you may not want or need an Active Directory infrastructure, or you may want to deploy standalone RD Gateways into an existing Amazon VPC, or you may want to domain-join your RD Gateway instances. To accommodate these approaches, we've provided CloudFormation templates that will launch an RD Gateway deployment for each scenario.

For an automated deployment that includes Active Directory Domain Services and RD Gateways, see Active Directory Domain Services on the AWS Cloud: Quick Start Reference Deployment.

Scenario 1: Deploy RD Gateway into a new Amazon VPC

The CloudFormation template for this scenario creates a new Amazon VPC and launches Windows Server 2012 instances into public subnets in two Availability Zones that will host the RD Gateway Role. After launching this CloudFormation stack, you will have deployed the infrastructure illustrated in Figure 19.

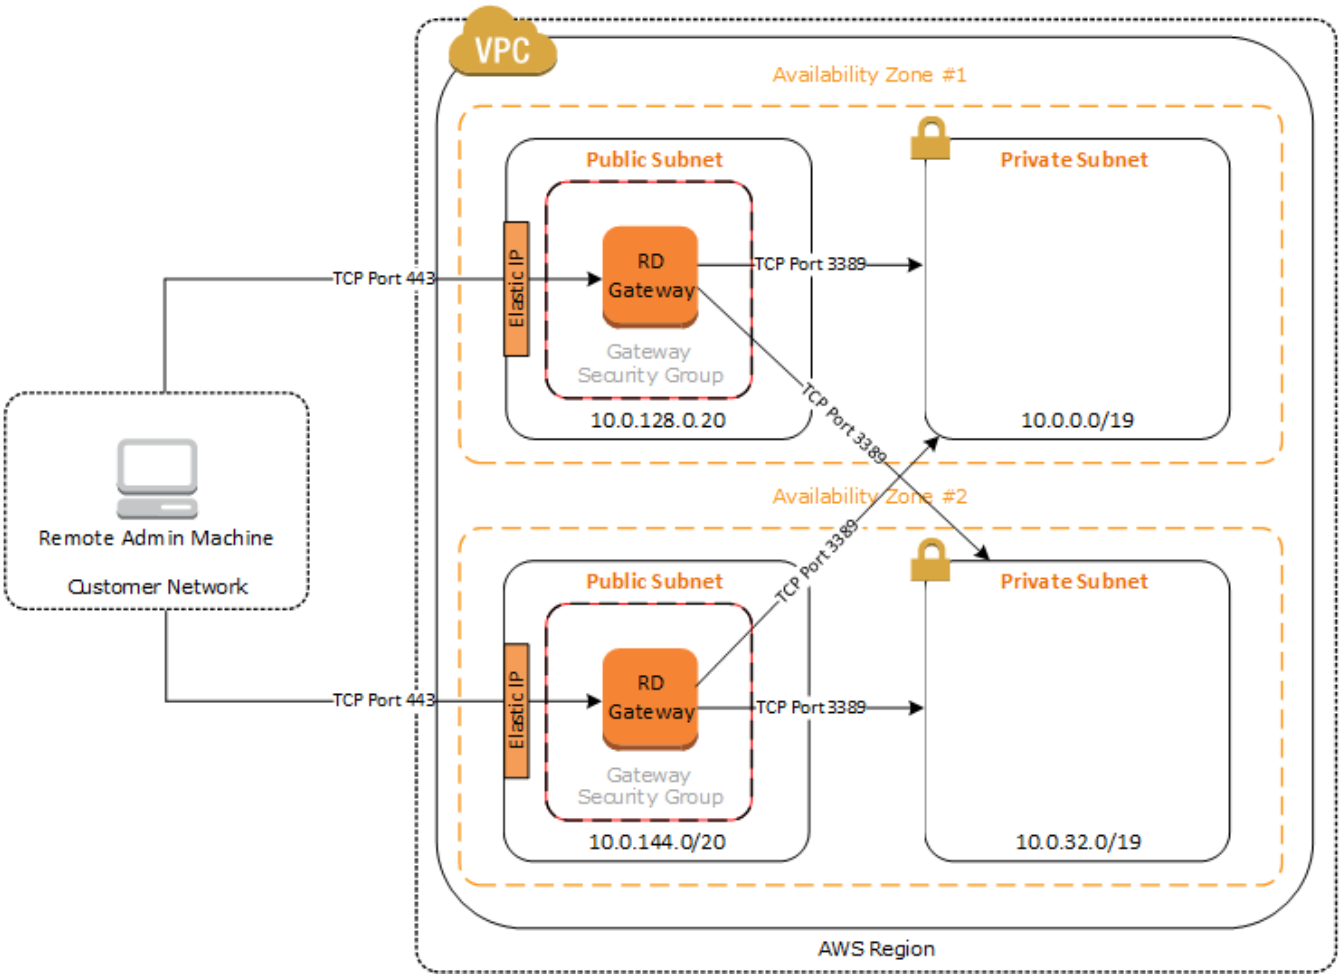

Figure 19: Architecture for Sample Deployment Scenario 1

Template Customization

The template for this scenario allows for rich customization of parameters at template launch. You can modify these parameters, change the default values, or, if you choose to edit the code of the template itself, create an entirely new set of parameters based on your specific deployment scenario. The template parameters include the following default values.

Network Configuration:

Parameter	Default	Description
Availability Zones	*Requires input*	The list of Availability Zones to use for the subnets in the VPC. The Quick Start uses two Availability Zones from your list and preserves the logical order you specify.
VPC CIDR	10.0.0.0/16	CIDR block for the VPC.
Private Subnet 1 CIDR	10.0.0.0/19	CIDR block for private subnet 1 located in Availability Zone 1.
Private Subnet 2 CIDR	10.0.32.0/19	CIDR block for private subnet 2 located in Availability Zone 2.
Public Subnet 1 CIDR	10.0.128.0/20	CIDR block for the public (DMZ) subnet located in Availability Zone 1.
Public Subnet 2 CIDR	10.0.144.0/20	CIDR block for the public (DMZ) subnet located in Availability Zone 2.
Allowed Remote Desktop Gateway External Access CIDR	0.0.0.0/0	Allowed CIDR block for external access to the Remote Desktop Gateway instances.

Amazon EC2 Configuration:

Parameter	Default	Description
Key Pair Name	*Requires input*	Public/private key pair that allows you to connect securely to your instance after it launches.
NAT Instance Type	t2.small	EC2 instance type for the NAT instances. NAT instances are used only if the region doesn't support NAT gateways.
Remote Desktop Gateway 1 Instance Type	t2.large	EC2 instance type for the first Remote Desktop Gateway instance.
Remote Desktop Gateway 1 NetBIOS Name	RDGW1	NetBIOS name of the first Remote Desktop Gateway instance (up to 15 characters).
Remote Desktop Gateway 1 Private IP Address	10.0.128.11	Fixed private IP for the first Remote Desktop Gateway instance located in Availability Zone 1.
Remote Desktop Gateway 2 Instance Type	t2.large	EC2 instance type for the second Remote Desktop Gateway instance.
Remote Desktop Gateway 2 NetBIOS Name	RDGW2	NetBIOS name of the second Remote Desktop Gateway instance (up to 15 characters).
Remote Desktop Gatewau 2 Private IP Address	10.0.144.11	Fixed private IP for the second Remote Desktop Gateway instance located in Availability Zone 2.

Microsoft Remote Desktop Gateway Configuration:

Parameter	Default	Description
Admin User Name	StackAdmin	Administrator name for the new local administrator account.
Admin Password	*Requires input*	Password for the new administrator account. This must be a complex password that's at least 8 characters long.
Domain DNS Name	example.com	Fully qualified domain name (FQDN) of the forest root domain.

This architecture provides an empty application tier for instances in private subnets. If more tiers are required, you can create additional private subnets with unique CIDR ranges to accommodate other tiers.

To launch the AWS CloudFormation template into the US West (Oregon) region, launch the Quick Start.

Scenario 2: Deploy Standalone RD Gateway into an Existing VPC

The CloudFormation template for this scenario launches Windows Server 2012 instances into public subnets in two Availability Zones in an existing Amazon VPC. After launching this CloudFormation stack, you will have deployed the following infrastructure.

- Two Windows Server 2012 instances, each launched into independent public subnets
- Security group for the RD Gateway instances with an ingress rules permitting TCP port 3389 to the RD Gateway servers

This scenario deploys standalone RD Gateway instances; for domain-joined instances, see scenario 3.

Template Customization

The template for this scenario allows for rich customization of parameters at template launch. You can modify these parameters, change the default values, or, if you choose to edit the code of the template itself, create an entirely new set of parameters based on your specific deployment scenario. The template parameters include the following default values:

Parameter	Default	Description
AdminPassword	*Requires input*	Password for the new administrator account. This must be a complex password that's at least 8 characters long.
AdminUser	StackAdmin	User name for the new local administrator account.
DomainDNSName	example.com	Fully qualified domain name (FQDN).
KeyPairName	*Requires input*	Public/private key pair that allows you to connect securely to your instance after it launches.
PublicSubnet1ID	*Requires input*	ID of the subnet you want to provision the first Remote Desktop Gateway into (e.g., subnet-a0246dcd).
PublicSubnet2ID	*Requires input*	ID of the subnet you want to provision the second Remote Desktop Gateway into (e.g., subnet-a0246dcd).
RDGW1InstanceType	t2.large	EC2 instance type for the first Remote Desktop Gateway instance.
RDGW1NetBIOSName	RDGW1	NetBIOS name of the first Remote Desktop Gateway (up to 15 characters).
RDGW1PrivateIP	10.0.128.11	Fixed private IP for the first Remote Desktop Gateway.
RDGW2InstanceType	t2.large	EC2 instance type for the second Remote Desktop Gateway instance.
RDGW2NetBIOSName	RDGW2	NetBIOS name of the second Remote Desktop Gateway (up to 15 characters).
RDGW2PrivateIP	10.0.144.11	Fixed private IP for the second Remote Desktop Gateway.
RDGWCIDR	0.0.0.0/0	Allowed CIDR block for external access to the Remote Desktop gateway instances.
VPCID	*Requires input*	ID of the VPC (e.g., vpc-0343606e).

To launch the AWS CloudFormation template into the US West (Oregon) region, launch the Quick Start.

Scenario 3: Deploy Domain-Joined RD Gateway into an Existing VPC

This scenario is similar to scenario 2, except that it provides domain-joined RD Gateway instances in an existing VPC. It provides a few additional parameters for customizing this configuration.

Parameter	Default	Description
ADServer1PrivateIP	10.0.0.10	Fixed private IP for the first Active Directory server located in Availability Zone 1.
ADServer2PrivateIP	10.0.32.10	Fixed private IP for the second Active Directory server located in Availability Zone 2.

DomainAdminPassword	*Requires input*	Password for the domain administrator user. This must be a complex password that's at least 8 characters long.
DomainAdminUser	StackAdmin	User name for the domain administrator. This is separate from the default administrator account.
DomainDNSName	example.com	Fully qualified domain name (FQDN).
DomainMemberSGID	*Requires input*	ID of the Domain Member security group (e.g., sg-7f16e910).
DomainNetBIOSName	example	NetBIOS name of the domain (up to 15 characters) for users of earlier versions of Windows.
KeyPairName	*Requires input*	Public/private key pair that allows you to connect securely to your instance after it launches.
PublicSubnet1ID	*Requires input*	ID of the subnet you want to provision the first Remote Desktop Gateway into (e.g., subnet-a0246dcd).
PublicSubnet2ID	*Requires input*	ID of the subnet you want to provision the second Remote Desktop Gateway into (e.g., subnet-a0246dcd).
RDGW1InstanceType	t2.large	EC2 instance type for the first Remote Desktop Gateway instance.
RDGW1NetBIOSName	RDGW1	NetBIOS name of the first Remote Desktop Gateway (up to 15 characters).
RDGW1PrivateIP	10.0.128.11	Fixed private IP for the first Remote Desktop Gateway.
RDGW2InstanceType	t2.large	EC2 instance type for the second Remote Desktop Gateway instance.
RDGW2NetBIOSName	RDGW2	NetBIOS name of the second Remote Desktop Gateway (up to 15 characters).
RDGW2PrivateIP	10.0.144.11	Fixed private IP for the second Remote Desktop Gateway.
RDGWCIDR	0.0.0.0/0	Allowed CIDR block for external access to the Remote Desktop gateways.
VPCID	*Requires input*	ID of the VPC (e.g., vpc-0343606e).

To launch the AWS CloudFormation template into the US West (Oregon) region, launch the Quick Start.

RD Gateway Deployment Checklist

Now that we've covered the architecture considerations and deployment steps for running the RD Gateway role in the AWS cloud, we can use the following high level checklist to help ensure the deployment goes smoothly.

Note: If you've deployed your RD Gateways using the automated solution in this guide, steps 1-4 will already be complete.

1. Create an "RD Gateway" Security Group for Windows-based instances that will host the RD Gateway role. Create an ingress rule permitting TCP port 3389 from your administrator IP address.

2. Deploy Windows-based instances into public Amazon VPC subnets in each Availability Zone that will run the RD Gateway role. Associate these instances with the "RD Gateway" Security Group. Ensure that an EIP is assigned to each RD Gateway instance so that it is reachable directly from the Internet.

3. Connect to the RD Gateway instances via RDP and install the RD Gateway role.

4. On the RD Gateway, install an SSL certificate and configure RD CAP and RD RAP policies.

5. Create Security Groups for your Windows-based instances that will be located in private Amazon VPC subnets. Create an ingress rule permitting TCP port 3389 from the "RD Gateway" Security Group, CIDR range, or IP address. Associate these groups with instances as they are launched into the private subnets.

6. Ensure that the name for the RD Gateway endpoint (e.g. rdgw1.example.com) is resolvable by administrative clients.

7. Modify the "RD Gateway" Security Group. Remove the ingress rule permitting TCP port 3389. Create a new ingress rule permitting TCP port 443 from your administrators IP address.

8. Make sure that instances in private subnets are associated with a Security Group containing ingress rules permitting the RD Gateway server IP address to connect via TCP port 3389.

9. Configure administrative clients with the proper configuration settings. This includes installing the root certificate from each RD Gateway server on the client machines. When you use the CloudFormation templates, the default location for the root certificate will be *c:\ <servername>.cer* on each RD Gateway server.

Further Reading

- AWS Quick Starts
    - http://aws.amazon.com/quickstart/
- Microsoft on AWS
    - http://aws.amazon.com/microsoft/
- Secure Microsoft Applications on AWS
    - http://media.amazonwebservices.com/AWS_Microsoft_Platform_Security.pdf
- Building a Modular and Scalable Virtual Network Architecture with Amazon VPC: Quick Start Reference Deployment
    - http://docs.aws.amazon.com/quickstart/latest/vpc/
- Scenarios for Amazon VPC
    - http://docs.aws.amazon.com/AmazonVPC/latest/UserGuide/VPC_Scenarios.html
- Active Directory Domain Services on the AWS Cloud: Quick Start Reference Deployment
    - http://docs.aws.amazon.com/quickstart/latest/active-directory-ds/
- Amazon EC2 Windows Guide
    - http://docs.aws.amazon.com/AWSEC2/latest/WindowsGuide/

Appendix

Remote Desktop Connection Manager

Microsoft provides a free utility called Remote Desktop Connection Manager (RDCMan) that manages multiple remote desktop connections in a single user interface. This is a useful tool for managing your Amazon EC2 Windows fleet through an RD Gateway infrastructure running in the AWS cloud.

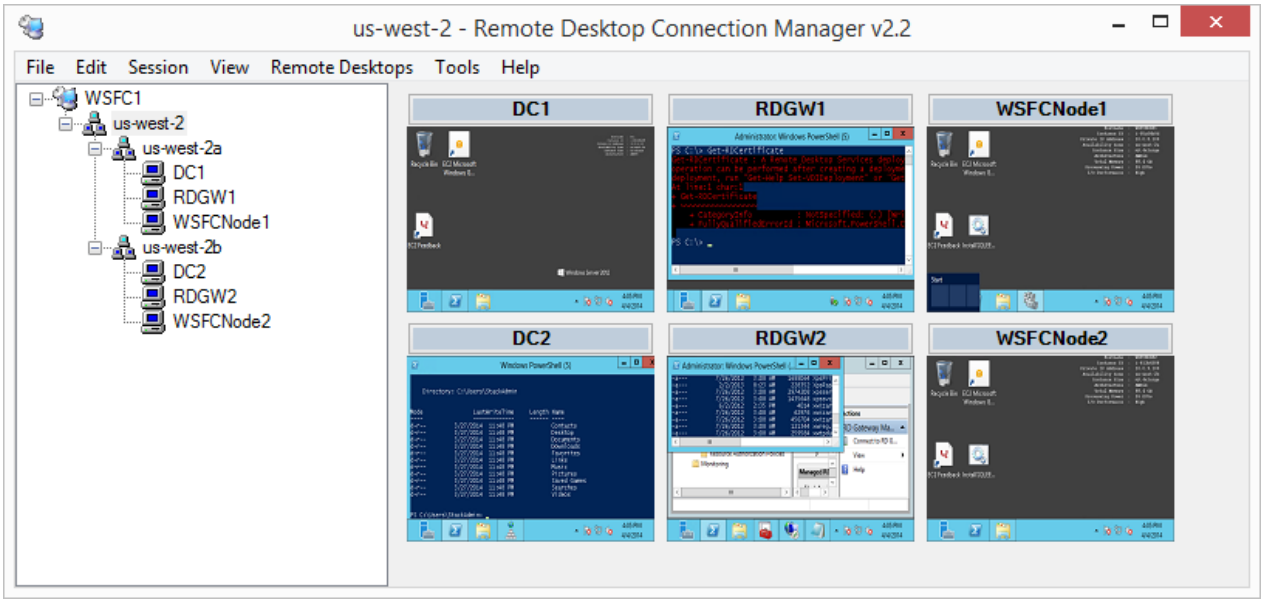

Figure 20: Managing Windows Instances with RDCMan

You can use this tool to manage multiple instances through one or more RD Gateways. You can define groups and server objects that correspond to the Availability Zones and Amazon EC2 instances running in the AWS cloud. You can download the sample connection file shown in Figure 20 from AWS and modify it to fit your environment.

Send Us Your Feedback

Please post your feedback or questions on the AWS Quick Start Discussion Forum.

You can visit our GitHub repository to download the templates and scripts for this Quick Start, and to share your customizations with others.

Document Revisions

Date	Change	In sections
July 2016	Updated the templates to use NAT gateways and an updated VPC configuration; added template for domain-joined RD Gateway instances.	Automated Deployment
September 2015	In the sample templates, changed the default type for RD Gateway instances from **m3.xlarge** to **m4.xlarge** for better performance and price.	Automated Deployment (template customization tables for sample deployment scenarios 1 and 2)
May 2015	Added Figure 19, which illustrates sample deployment scenario 1.	Scenario 1
March 2015	Optimized the underlying Amazon VPC design to support expansion and to reduce complexity.	Architecture diagram and template updates
November 2014	In the sample templates, changed the default type for **NATInstanceType** to **t2.small** to support the EU (Frankfurt) region.	Automated Deployment (template customization table for sample deployment scenario 1)