

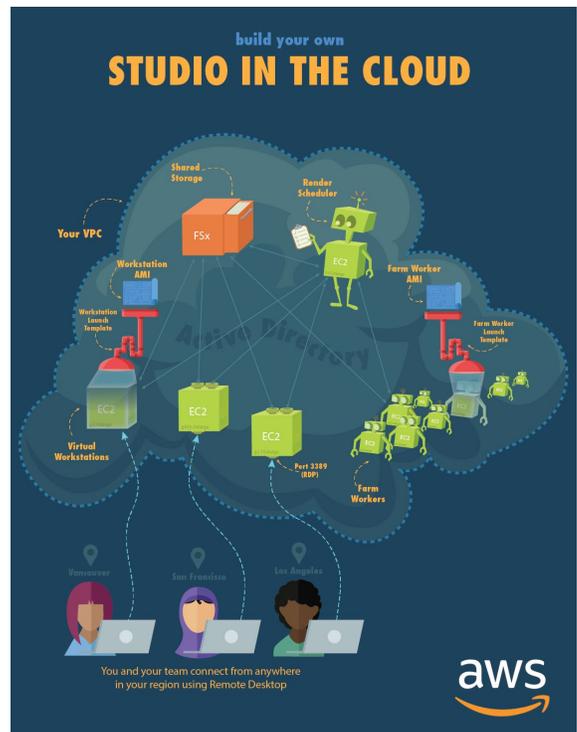
Tutorial 1: Getting Started with AWS Virtual Workstations

Estimated Time to Complete: 40 minutes

Overview

Welcome to tutorial 1 from our Studio in the Cloud series. This step by step series is written for creative studios looking to fully adopt the cloud for producing content. The series has been created by artists including animators and composers, so it's approach will resonate with an audience less familiar with cloud technology and using the AWS console. This tutorial series will walk you through setting up virtual workstations, cloud storage, and cloud rendering. By the end of the 7-part series, you'll have a fully cloud-based studio that leverages the scale, power, and convenience of AWS.

Each tutorial will start with an overview, like this one, as well as an estimate of the time required to complete. There will also be exercises throughout to confirm that you have correctly completed the previous steps. Supplemental information about the concepts covered in each tutorial and links to additional websites will appear at the bottom of the tutorial page.



Prerequisites

Before you begin, please check the following prerequisites to make sure that you'll be able to successfully complete the tutorials.

Region Support

Before you jump into creating your Studio in the Cloud, you should make sure that the geographic region you are in (or the one closest to you) supports all of the services and instance types required.

A Region represents a specific geographic area around the world where AWS maintains a network and servers. Each Region has one or more Availability Zone, which are isolated locations inside the Region. For the best performance (due to the current laws of physics), you want AWS services to be located as geographically close to you as possible. With over 20 Regions, there is a good chance that there is a Region located near you. Having multiple Availability Zones in a Region provides additional reliability.

At the time of publication (February 2020), the following AWS Regions fully support all of the services needed for Studio in the Cloud:

- N. Virginia (us-east-1)
- Ohio (us-east-2)
- N. California (us-west-1)
- Oregon (us-west-2)
- Singapore (ap-southeast-1)
- Sydney (ap-southeast-2)
- Tokyo (ap-northeast-1)
- Frankfurt (eu-central-1)
- Ireland (eu-west-1)
- London (eu-west-2)
- Stockholm (eu-north-1)

If you're not sure which Region is closest to you, please see the [Regions and Availability Zones](#) webpage. You may also want to use a website such as ping.psa.fun to measure the latency from your location to the different Regions. The Region with the lowest latency value is usually the one you want.

Please note: If your nearest Region does not have all of the required services, there may be another nearby Region that does. In addition, Regions are being updated all the time to add support for additional services and instance types. So in time, your closest Region may become supported. Please see the [Appendix](#) for information on the specific services that are required and links to webpages with the most up-to-date information on whether they are available in your nearest Region.

AWS Account

Now that you know whether your nearest AWS Region is supported, you should create an account. If you don't already have one: [Create an AWS account](#)

If you already have an AWS account, we have an extra question for you...is it your personal account or a linked account provided by your employer? If it's your personal account, then please be aware that there will be costs associated with using the AWS services required for these tutorials. We provide a rough estimate of those costs in the next section.

If your account is linked to your employer or another party, then you should check that they are prepared to accept responsibility for the costs involved. For more information on linked accounts and consolidated billing, please see this AWS documentation: [Consolidated Billing Process](#).

You will also need to verify that your account has sufficient permissions to access all of the AWS services required to complete the tutorials.

Account Permissions

If you are the owner of your account:

If you just created a new AWS account or are the owner of your account, you may be logging in as the root user of your account. How can you tell? If you're logging in using an email address, rather than a user name, you're using the root account.

 **Note:** Rather than continue to login as the root user, we strongly suggest that you create a new user and login with those credentials instead. Why? If your root user email and password are ever stolen, then the person who stole them will have full access to your account and there isn't anything you can do about it, unless you close the account. Creating a new user for yourself or others on your account allows you to have more control. You can specify exactly what permissions each user has and if that user information is ever lost, you can simply remove that user from your account.

If you're unfamiliar with how to create a new user and add permissions to it, we've provided instructions in the Appendix: [Creating an IAM User](#).

Even if you know how to create a user, please click the link above and follow the directions to add the specific permissions needed for this tutorial.

Accounts provided by a third party:

If your account is provided by your employer or a third party, you're likely already logging in as a user and not as the root. In this case, you do not need to create a new user, but you may need to request to have additional permission policies added in order to complete the tutorials. Please be aware that the owner of your account may have rules about which policies you can access.

Please click on the following link to a JSON file which lists of all the permissions needed for the tutorials: [Tutorials Permissions Policy](#). It's best to check with the administrator for your account and request that these permissions be added to your user.

On-Demand G Instance Quota

Many resources in AWS are subject to quotas, also called limits. Each resource has a quota that represents the maximum value of that resource you can use. These quotas are in place to ensure that AWS can provide highly available and reliable service to all our customers, but also to protect you against accidentally creating too many resources and incurring unexpected charges.

Our tutorials make use of GPU-enabled virtual workstations (G4 and G3 instances). If you are using a new AWS account or have never used GPU instances before, your quota for that instance type may be at the default value of zero. Before starting the tutorials, you should check the value of your G instance quota and request an increase, if necessary. Please see the [Appendix](#) for more instructions. Check out the [AWS Documentation](#) to learn more about service quotas.

Microsoft Remote Desktop

To keep things as simple as possible, we'll be using Remote Desktop to connect to your Windows instances. There are other options for connecting, such as [NICE DCV](#) from AWS or [Cloud Access Software](#) from Teradici. These solutions provide performance benefits over Remote Desktop, but

require some extra setup which is beyond the scope of these tutorials. For these reasons, we'll be sticking with Remote Desktop for now.

If your local computer is a Windows machine, Remote Desktop should already be installed. But if you're on macOS or Linux, you may need to install it. Here are links to download the latest version for [Windows](#), [macOS](#) and [Linux](#).

Estimated Costs

AWS services are charged using a pay-as-you-go model. That means that you only pay for the individual services you use, for the time that you use them. For example, virtual workstations are charged per hour for each hour that they are running (whether you are logged in or not), while storage is charged based on the type and amount of storage you use.

In addition, some services are free to use up to certain limits, while others may be free for a certain amount of time after you initially sign up for an account. After you have exhausted the Free Tier usage, you switch to paying for what you use. Please click on the following link for more information about the [AWS Free Tier](#).

The ultimate cost of your Studio in the Cloud will depend on how many hours it is running and how much of each service you consume. Items such as shared cloud storage and render hours can vary widely depending on your particular use case. In addition, costs for AWS services differ slightly from region to region. However, we understand that before even starting these tutorials you may want some idea of how much they will cost to complete. Assuming that you complete the tutorials in a time frame similar to what is estimated for each step, and that you require a minimum amount of storage and compute time, we estimate that the cost to complete the tutorials will not exceed \$150-200 USD.

For a more accurate estimate of your cost, we recommend using the [AWS Pricing Calculator](#), which factors in the pricing for your particular region. We have listed the AWS services used in the tutorials, along with approximate usage amounts, in the [Appendix](#).

Our estimate above also does not account for any existing credits that you may have in your account. In addition, if you decide to keep your Studio in the Cloud infrastructure running after completing the tutorials, you will continue to accrue costs based on how much of each AWS service you consume. Because the costs will vary depending on your individual needs, we do not attempt to estimate those costs here. The tutorials do not make use of any licensed software, so any licensing costs for additional software that you choose to use are also not covered.

Monitoring Your AWS Costs and Usage

As you work through the tutorials, you can keep an eye on your current costs using the [AWS Billing Console](#). In addition, AWS has many other tools for budgeting and tracking usage. Please see the following blog post for more information and helpful tips: [Keeping an Eye on Your AWS Costs and Usage](#).

Security

Throughout these tutorials we will make a point to call out security best practices, such as not logging in as the root user, which was mentioned above, as well as others, such as limiting access with security groups. However, the intention of the tutorials is introduce you to the key concepts and services of Studio in the Cloud, rather than dive deep on security.

 *Note:* While the security measures implemented here are sufficient for initial setup and testing of a cloud-based production pipeline, we recommend implementing extra security after completing the tutorials, according to your individual security compliance goals for production content. In the future we will publish an add-on tutorial focused specifically on how to add extra layers of security to your Studio in the Cloud.

Keeping Track of Important Information

Many times in these tutorials you will be asked to refer back to information about components you created in a previous step, including names, IDs, IP addresses and more. To make it easier to keep track of this information without having to constantly backtrack, we've created a cheat sheet for you to fill out as you go. Please follow this link to download it: [Important Information Cheat Sheet](#). It's a form that you can either fill out in a PDF editor or print out and fill out by hand.

Starter Exercise: Using the AWS Console

Ready to get started? We're going to begin by introducing you to the **AWS Console**, which is your main interface to all AWS services.

For this first exercise, we're going to perform one of the most basic AWS tasks: launching an EC2 instance. EC2 stands for **Elastic Compute Cloud**. It's the AWS service that you use to launch the virtual computers that you will use for your Studio in the Cloud.

An **instance** is a single virtual computer in EC2. Throughout these tutorials we'll be creating several instances for different purposes, including user management, render management, virtual workstations, and cloud farm workers.

For your first instance, we're going to keep things easy and use some of the default settings that already exist in your account.

Log in to Your Account

- If you haven't already, follow this link to [log in to your AWS account](#).
- As a reminder, make sure you log in as an IAM user and not as the root user for your account. If you need to create a new user, please see the directions in the Appendix: [Creating an IAM User](#).

Fill Out Your Cheat Sheet

- We'll put reminders in *italics* throughout these tutorials to prompt you to fill out the [Important Information Cheat Sheet](#) as you go.

- Now is a good time to fill out the cheat sheet with your AWS Account email address and account ID or alias, as well as your IAM user name.

⚠ Note: You'll also need to remember your password, but for security reasons we don't recommend writing that down on the cheat sheet.

- If you downloaded your credentials.csv file above, you should also note that file's location in the cheat sheet for later reference.

Set Your Region

Next, you should check that your Region is set correctly. After logging in, your current Region is listed in the top right corner of the AWS Console, to the left of the Support drop down menu.



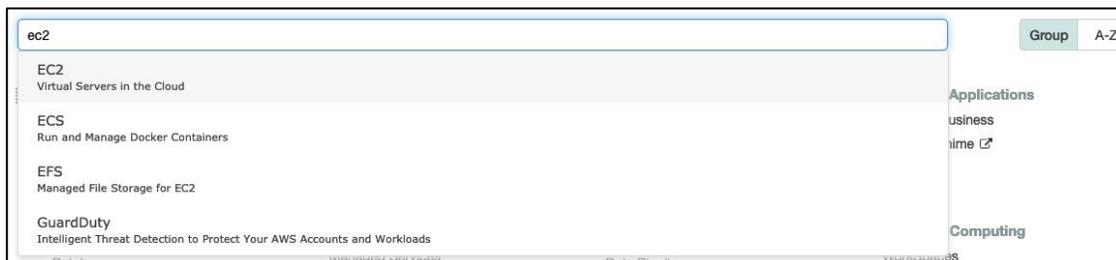
- To change your Region, click the down arrow next to the current Region and then select a different Region from the list.
- As a reminder, for these tutorials you must select a supported Region from the [list](#) above.
- Please take a moment to note the region you are using (e.g. N. Virginia (us-east-1)) in the [Important Information Cheat Sheet](#).

Launch an EC2 Instance

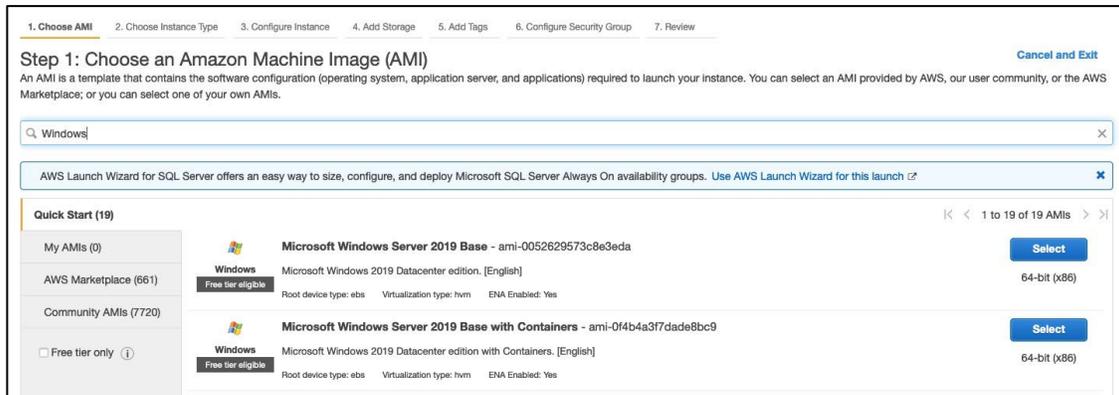
- At the top left of the webpage click the **Services** drop down menu, then under **Compute** select **EC2**.



- Note: Any time you are searching for a service, you can quickly **filter** the list of services by just typing the name of the service you want. Sometimes this is easier than searching for a particular category.



- On the **EC2 Dashboard** page that appears, click the **Instances** link in the left side panel.
- On the **Instances** page, click the **Launch Instance** button.
- For **Step 1: Choose an Amazon Machine Image (AMI)**, enter **Windows** in the search box and hit <enter>.
- You will see a long list of different Windows AMIs. Look for the one named exactly **Microsoft Windows Server 2019 Base**, it should be the first one in the list. Click the **Select** button next to it.



- For **Step 2: Choose an Instance Type**, leave it at the default and click the **Review and Launch** button.
- This will take you to **Step 7. Review Instance Launch**. A bunch of default settings have been filled in for you already. Feel free to look around at the details on this page. You may see a lot of unfamiliar terms, but don't worry, we'll go through everything you need to know later in this tutorial.

- Next click the **Launch** button at the bottom of the screen.

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

Step 7: Review Instance Launch

Please review your instance launch details. You can go back to edit changes for each section. Click **Launch** to assign a key pair to your instance and complete the launch process.

▼ AMI Details [Edit AMI](#)

 **Microsoft Windows Server 2019 Base - ami-0052629573c8e3eda**

 Microsoft Windows 2019 Datacenter edition. [English]
Root Device Type: ebs Virtualization type: hvm

▼ Instance Type [Edit instance type](#)

Instance Type	ECUs	vCPUs	Memory (GiB)	Instance Storage (GB)	EBS-Optimized Available	Network Performance
t2.micro	Variable	1	1	EBS only	-	Low to Moderate

▼ Security Groups [Edit security groups](#)

Security group name launch-wizard-2
Description launch-wizard-2 created 2019-12-09T17:49:38.172-08:00

Type i	Protocol i	Port Range i	Source i	Description i
<i>This security group has no rules</i>				

▶ Instance Details [Edit instance details](#)

▶ Storage [Edit storage](#)

▶ Tags [Edit tags](#)

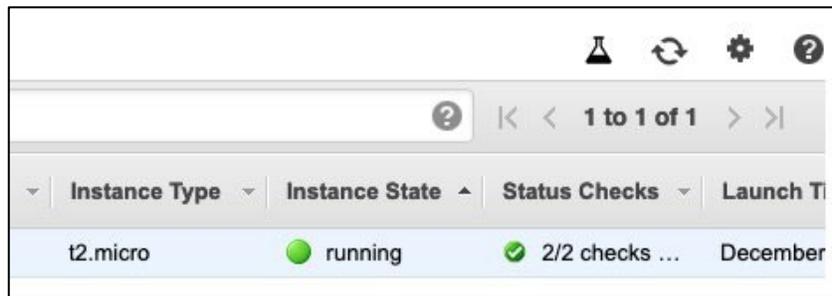
[Cancel](#) [Previous](#) [Launch](#)

- In the **Select an existing key pair or create a new key pair** popup window, select **Create a new key pair** from the first drop down menu and then enter a key pair name (e.g., `mystudio-keypair`)
 - A **key pair** consists of public and private key files that are used to encrypt data between two computers. AWS stores the public key file, but you need to store the private key file in order to connect to an AWS cloud workstation.
- After entering your key pair name, be sure to click the **Download Key Pair** button
 - *Important:* You must download the private key file when you create a new key pair and store it securely on your local computer. You will not have another chance to download the private key file, so save it in a safe place on your computer where you can find it again!

- You should also note the name of the key pair file in the [Important Information Cheat Sheet](#) for future reference.
- After downloading the private key file click the **Launch Instances** button.

View the Status of Your Instance

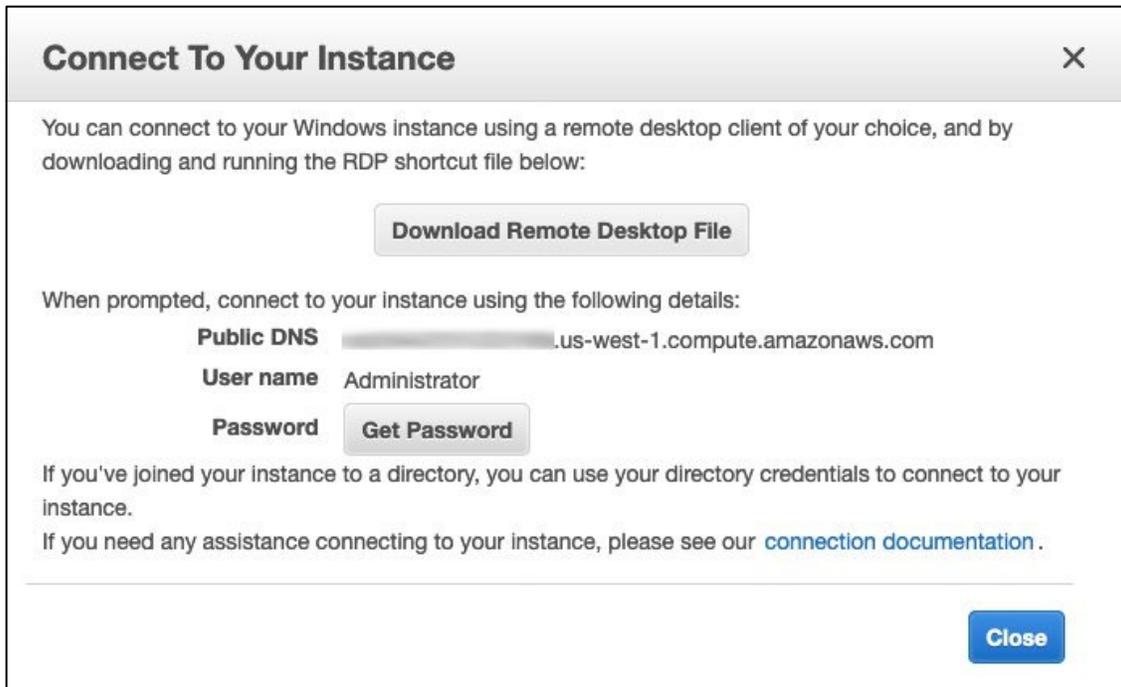
- On the **Launch Status** page click the **View Instances** button at the bottom right corner of the page
- This will take you back to the **Instances** page which will show a list of your running instances and their status. When you see the **Instance State** change from **pending** to **running** and the **Status Checks** change from **Initializing** to **2/2 checks passed**, your new instance is ready to be used. It may be between 5 to 10 minutes for the machine to be ready.



- Click on your instance to see more details in the panel at the bottom of the screen. You'll see many terms and acronyms like **VPC**, **Subnet** and **Security groups**, but don't worry, we'll go over all of those in our next tutorial.

Connect to Your Instance

- After your new instance is ready, select it from the list and then click the **Connect** button.
- In the **Connect To Your Instance** window that pops up, first click the **Get Password** button.



- In the **Connect To Your Instance > Get Password** window, click the **Choose File** (or **Browse**) button, then open the private key file that you created earlier.
- Next, click the **Decrypt Password** button.
- Hover your mouse over the password and then click the **copy** icon to the right of the password to copy it to your clipboard.



- Next, click the **Download Remote Desktop File** button.
- Your browser will download an .rdp file that the **Remote Desktop Connection** application will use to connect to your instance. Depending on your browser, you will either need to click on the downloaded .rdp file manually or click OK in another popup window when you're asked if you want to open it automatically.
- On Windows, click **Connect** in the popup window that appears. On a Mac, click **Continue** instead.
- Once you get a window asking for the Administrator password, paste the password that you copied into the password field and click **OK**.
- When a window pops up that says **The identity of the remote computer cannot be verified** click **Yes** (or **Continue** if on a Mac) to continue connecting.
- Finally, a new window will open that shows the desktop of your cloud workstation/instance.



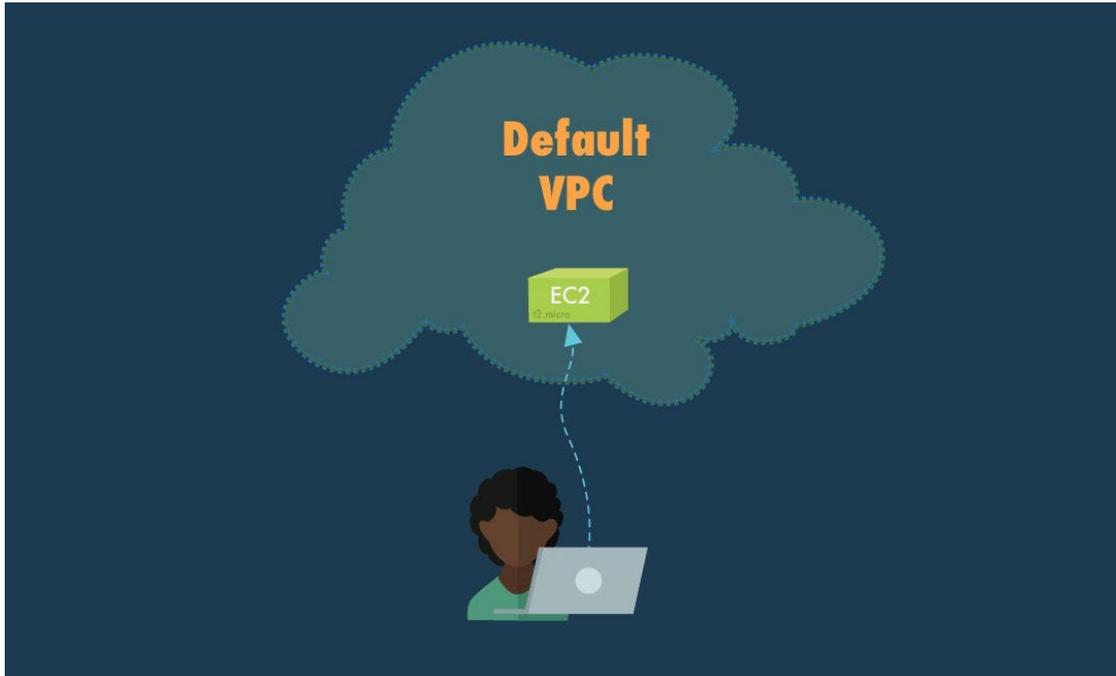
- You can interact with the Windows desktop of your instance and run applications just as if it were on your local computer. You may notice that there's not much installed on the instance, but don't worry about that, we'll show you how to install content creation apps like Blender in a later tutorial.
- Note: By default Remote Desktop tends to fill the entire screen that it runs on. If you have a single monitor, this can make it difficult to view both these tutorial pages and the Remote Desktop at the same time. You can click the "restore down" button and resize the Remote Desktop window as needed:



- However, you can also force Remote Desktop to start up at a resolution smaller than full size as well. We've provided Windows and MacOS instructions for that in the [Appendix](#).

Your First EC2 Instance

Congratulations! You have launched your first AWS EC2 instance!



Take a moment to marvel at what you have accomplished...but don't marvel too long. Once you've finished taking your instance for a test drive, you will want to disconnect and shut it down. Why is that? Move on to the next section to find out...

Cost Optimization

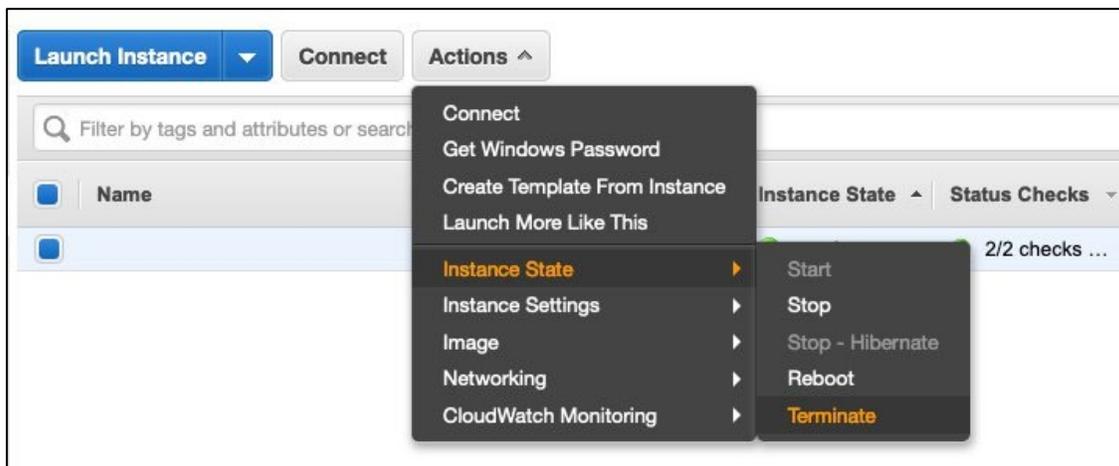
In an effort to keep your costs down as much as possible, at the end of each tutorial we'll provide instructions on which services you can shut down. This may be because they aren't needed anymore, like in the case of the instance you launched during the starter exercise above. Or it may be because it's going to be a day or two before you're able to move on to the next tutorial and you don't want to have running computers costing you money when you're not using them.

Below we will go over the procedure for stopping and/or terminating an instance. "Stopping" an instance is similar to shutting down your home or work computer. All running applications quit, but whatever data was stored on persistent storage (i.e., the instance's hard drive) is preserved. Once the instance is stopped, you are no longer charged any hourly fees for its use (although you are charged a small amount to store the data that was on the instance's hard drive). A stopped instance will still be listed in the EC2 instance list in the AWS Console and can be restarted at any time.

Terminating the instance is more final than stopping it. When an instance is terminated, it is stopped like above, but all the data that was stored in persistent storage is also deleted, so not only do you not incur any hourly charges, you also do not incur any charges for storage of data. You can think of it as first shutting down your computer, but instead of keeping it around for later, you send it to be recycled. Terminated instances will remain visible in the console for a little while, but unlike stopped instances, they cannot be restarted.

Terminate Your Instance

- From the **Start Menu** on your instance, click the **Power Icon** then select **Disconnect**. The Remote Desktop Connection app will close.
- Next, back in the **AWS Console**, select your instance.
- Then click the **Actions** button at the top of the list of instances and choose **Instance State**→**Terminate**. Then click the **Yes, Terminate** button.



The steps for stopping your instance are the same as above, except instead of Instance State→Terminate, you select Instance State→Stop. In this case, we don't need the start exercise instance anymore, so we can terminate it.

Throughout these tutorials, you'll be launching and connecting to many different instances, so you'll be starting and stopping instances on a regular basis. This first starter exercise is good practice for the steps yet to come.

In the next tutorial, we'll begin building the basic infrastructure for your Studio in the Cloud. Instead of using the default settings for your account, as we did here, we'll be adjusting the settings to optimize for a production workflow.

When you're ready, feel free to continue on to: [Tutorial 2: Creating a VPC and Active Directory for Your Studio](#).

Appendix

Tutorial Table of Contents

Tutorial 1: Getting Started with AWS Virtual Workstations - *this page*

[Tutorial 2: Creating a VPC and Active Directory for Your Studio](#)

[Tutorial 3: Setting Up an FSx File System and User Accounts](#)

[Tutorial 4: Building a Render Scheduler with AWS Thinkbox Deadline](#)

[Tutorial 5: Installing Applications and Creating a Workstation AMI](#)

[Tutorial 6: Setting Up a Linux Farm Worker and Spot Fleet Request](#)

[Tutorial 7: Onboarding New Artists and Sample Workflow](#)

Links to AWS Documentation

- [AWS Services By Region](#)
- [AWS Instance Types \(all Regions but Beijing & Ningxia\)](#)
- [AWS Instance Types \(Beijing & Ningxia only\)](#)
- [Consolidated Billing Process](#)
- [AWS Free Tier](#)
- [AWS Pricing Info](#)
- [AWS Pricing Calculator](#)
- [AWS Regions and Availability Zones](#)
- [What is Service Quotas?](#)
- [Launching an Instance Using the Launch Instance Wizard](#)
- [Connecting to Your Windows Instance](#)
- [Stop and Start Your Instance](#)
- [Terminate Your Instance](#)

Links to Other Resources

- ping.psa.fun - Useful website for estimating latency to different AWS Regions
- [Microsoft Remote Desktop for Windows](#)
- [Microsoft Remote Desktop for macOS](#)
- [rdesktop Remote Desktop Client for Linux/Unix](#)

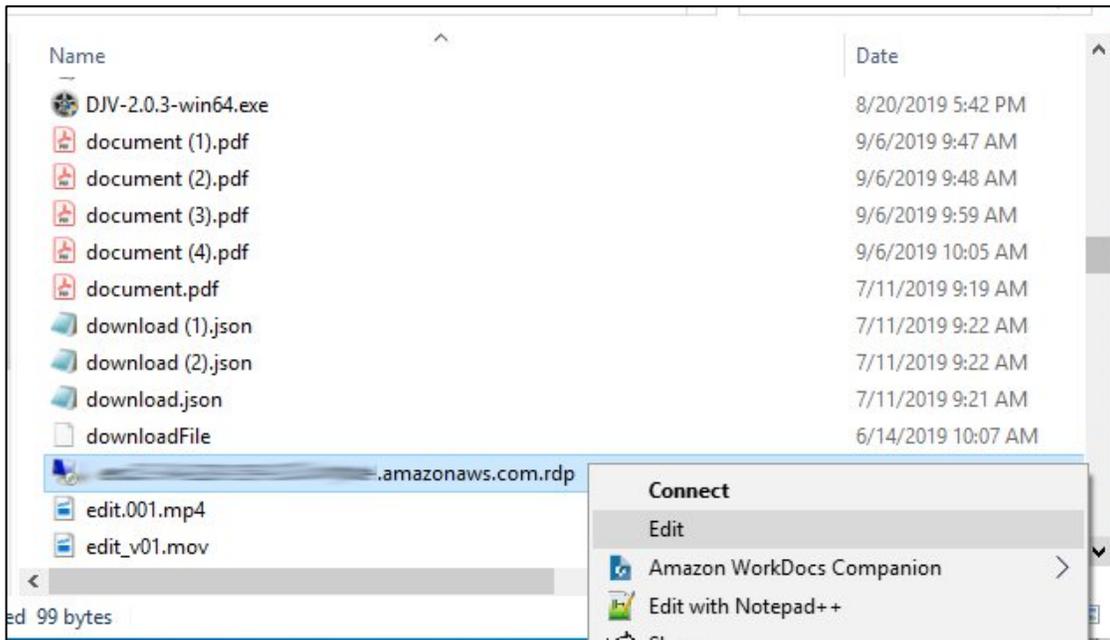
Downloads

- [Important Information Cheat Sheet](#)

Adjusting the Resolution of the Remote Desktop Window

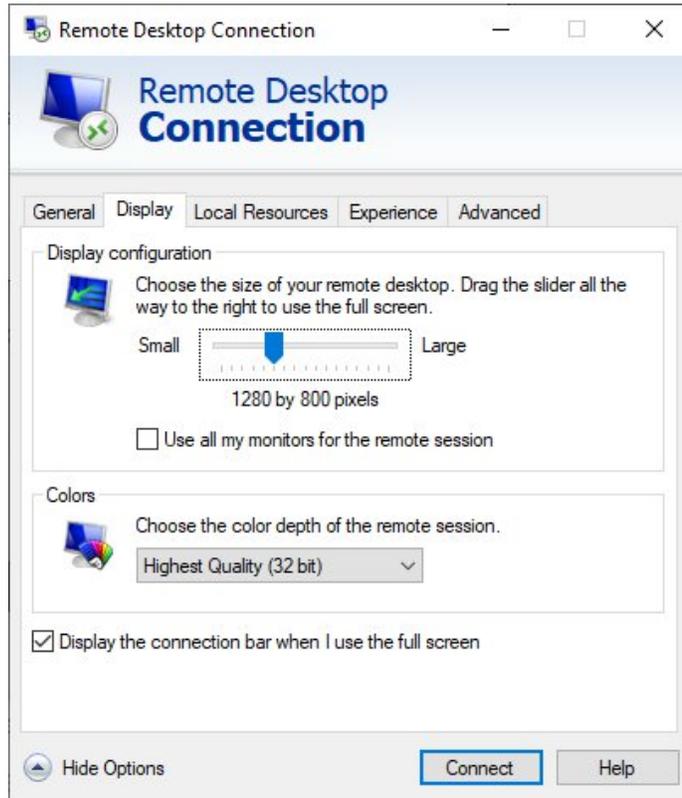
Windows:

- After clicking the **Download Remote Desktop File** to connect to your instance, if prompted by your browser, save the file.
- Find the .rdp file in your **Downloads** folder. It will be named something like ec2-35-39-399-854.compute-1.amazonaws.com.rdp



- Right-click on the .rdp file and choose **Edit**

- Click on the **Display** tab in the Remote Desktop Connection window



- Select a resolution smaller than you're monitor's resolution and then click **Connect**

MacOS:

- After clicking the **Download Remote Desktop File** to connect to your instance, if prompted by your browser, save the file.
- Find the .rdp file and open it with your favorite text editor
- Add the following two lines to the bottom of the file, where you set the width and height values to something smaller than your monitor's current resolution:
- Desktopwidth:i:1920
Desktopheight:i:1080
- Save the .rdp file and then launch Remote Desktop with that file

Required Services for Studio in the Cloud

AWS Service	Availability (as of 11/26/19)	Purpose	Usage for Tutorials
EC2 On Demand Instances	all regions	virtual computers in the cloud	<i>see instance types table below</i>
EC2 Spot Instances	all regions	discounted virtual computers ideal for render farm usage	<i>see instance types table below</i>
EC2 Auto Scaling	all regions except South America (Sao Paolo) and AWS GovCloud (US-East & US-West)	allows your render farm to automatically grow/shrink based on usage	<i>see instance types table below</i>
VPC	all regions	virtual network that contains all the parts of your studio	<i>no usage charges for tutorials</i>
IAM	all regions	grant/restrict access to AWS services	<i>no usage charges for tutorials</i>
FSx	limited regions	high speed managed storage	32 GB minimum
Directory Service	all regions except Asia Pacific (Osaka-Local)	manage user accounts and profiles for your studio	about 2-5 days
Secrets Manager	all regions except Asia Pacific (Osaka-Local) and China (Beijing & Ningxia)	store confidential information, like administrator passwords	1 secret

Required Instance Types for Studio in the Cloud

AWS Workstation Type	Availability (as of 11/26/19)	Purpose	Usage for Tutorials
m5.xlarge - Windows, 30 GB storage	all regions	studio creation and management	about 5-10 hours
m5.2xlarge - Linux, 300 GB storage	all regions	render farm workers	about 5-10 hours
m5.2xlarge - Windows, 100 GB storage	all regions	render scheduling	about 15-25 hours
g3.4xlarge or g4dn.4xlarge - Windows, 150 GB storage	limited regions	GPU-enabled Windows workstations	about 10-15 hours

Tracking Availability in Your Region

You can track the availability of the required services and instance types in your Region using the websites below:

[AWS Services By Region](#)

[AWS Instance Types By Region \(all regions but Beijing & Ningxia\)](#)

[AWS Instance Types By Region \(Beijing & Ningxia only\)](#)

Creating an IAM User

- Once logged in as the root user, go to the top left of the webpage click the **Services** drop down menu, then under **Security, Identity, & Compliance**, choose **IAM** or simply search for "IAM" in the search field



- In the navigation pane on the left, select **Users**
- Next click the **Add user** button
- Enter a **User name** for your new IAM user (e.g., bob)
- Check the check boxes next to both **Programmatic access** and **AWS Management Console access**
- Select **Custom password** and enter a password in the field
- Uncheck the check box next to **Require password reset**
- Click the **Next: Permissions** button
 - Click on **Attach existing policies directly**
 - Click the **Create policy** button. This will open a new browser tab
 - Next click on the **JSON** tab
 - <shift>+click on the image below to open a new browser tab with the text that needs to be entered into the JSON entry field:

A screenshot of a web browser window. The address bar shows the URL: https://studio-in-the-cloud-tutorial-1.s3-us-west-1.amazonaws.com/tutorial-1/sitc_iam_policy.json. The main content area displays the JSON content of the file, which is an IAM policy document. The JSON is formatted with indentation and line breaks. The policy allows actions on EC2 instances, Elastic Load Balancing, and various AWS Systems Manager (SSM) actions.

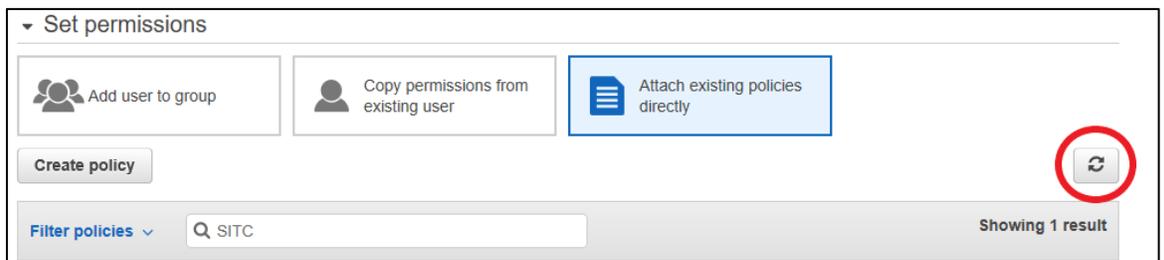
```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Action": "ec2:*",
      "Effect": "Allow",
      "Resource": "*"
    },
    {
      "Effect": "Allow",
      "Action": "elasticloadbalancing:*",
      "Resource": "*"
    },
    {
      "Effect": "Allow",
      "Action": [
        "ssm:CreateAssociationBatch",
        "ssm:CreateDocument",
        "ssm:DescribeAssociation",
        "ssm:GetDeployablePatchSnapshotForInstance",
        "ssm:GetDocument",
        "ssm:DescribeDocument",
        "ssm:GetManifest",
        "ssm:GetParameter",
        "ssm:GetParameters",
        "ssm:ListAssociations",
        "ssm:ListDocuments",
        "ssm:ListInstanceAssociations"
      ]
    }
  ]
}
```

sitc_iam_policy.json - <shift>+click on the image above to open the JSON file in a new tab

- Replace the existing text in the JSON entry field with the text from the file above. Note: The file is quite long, so please make sure you get everything.
- When you're all done, the JSON entry field should look like this:



- Click the **Review policy** button
- For **Name** enter **Studio-in-the-Cloud-IAM-Policy**
- For **Description** enter **Custom IAM policy with all permissions to complete the Studio in the Cloud tutorials**
- Click the **Create policy** button
- Return to the browser tab that you were using to create your user
- Refresh the policies list by clicking the refresh button on the right side of the page, opposite the "Create Policy" button

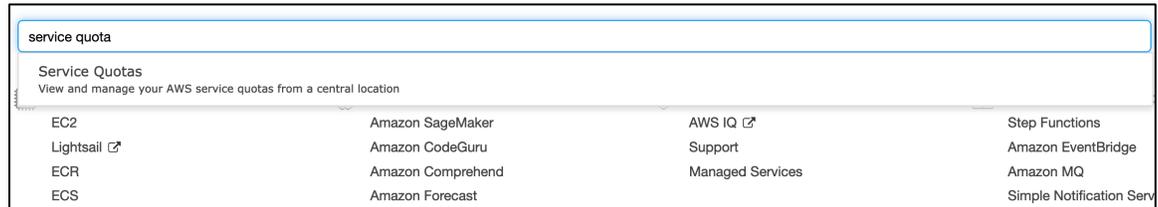


- In the policies search field, enter **Studio-in-the-Cloud-IAM-Policy** and then click the checkbox next to it to select it
- Click the **Next: Tags** button
- Click the **Next: Review** button
- Review the information for the user and then click the **Create User** button

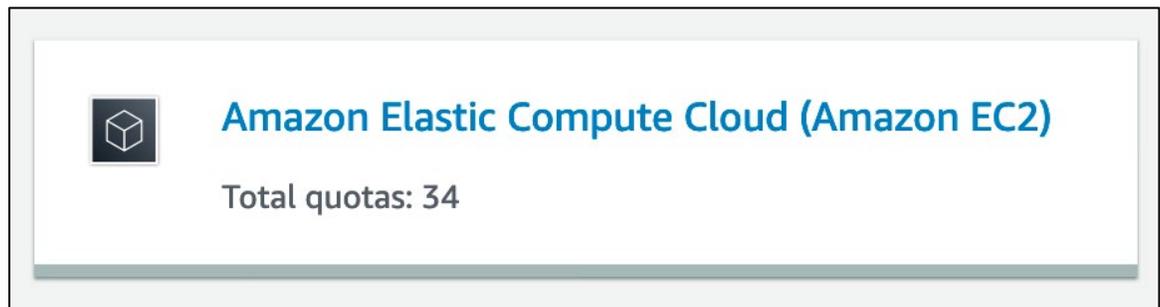
- On the next page, you'll be presented with the security credentials for your user. You'll need these credentials later in the tutorials, so it is important that you download them now by clicking the **Download .csv** button.
- Make note of the location of the .csv file so that you can find it later. *Also enter the location of the .csv file on the [Important Information Cheat Sheet](#) that we've provided.*
- When you're done, click the sign-in link at the bottom of the green box and login as your new user.

Checking Your G Instance Quota

- Login to your AWS account
- If you haven't already, check that your Region is set correctly. For instructions, please see the [Set Your Region](#) section above.
- Click the **Services** drop down menu next to the AWS logo at the top left corner of the AWS Console and type **service quota** in the search field.



- Select **Service Quotas** to go to the Service Quotas Dashboard.
- Click the **Amazon Elastic Cloud Compute (Amazon EC2)** card.



- In the "Filter by..." field, type in **G instance**
- Now you can review the quota you have available for your G instances. In order to complete these tutorials, you will need an **Applied quota value** of at least **16** because the **g4dn.4xlarge** instance we will be using in part of the tutorial uses **16 vCPUs**.

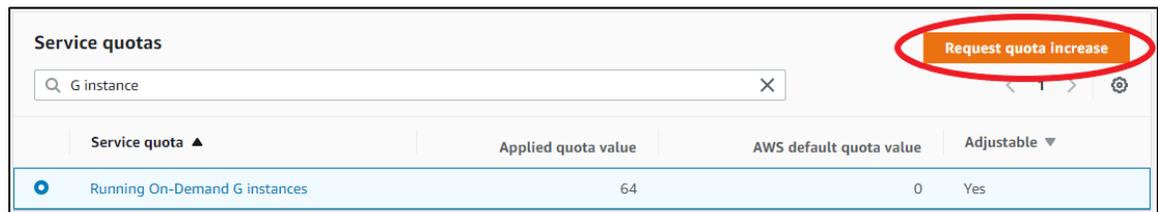
The screenshot shows the 'Service quotas' dashboard. A search filter 'G instance' is applied. The table below shows the quota details for 'Running On-Demand G instances'.

Service quota ▲	Applied quota value	AWS default quota value	Adjustable ▼
Running On-Demand G instances	64	0	Yes

- To see how many vCPUs are used by other G4 instance types, please see the [G4 Product Details](#) chart.
- **Note:** After completing the tutorials, you may need more than 16 vCPUs depending on how many additional GPU workstations you need to launch for your artists.
- To request an increase in your quota value, please follow the instructions below.

Requesting a Quota Increase

- Click the **circle** to the left of **Running On-Demand G instances**.
- Then click the **Request quota increase** button.



- Enter the new **Quota value** you would like to receive. **Note:** Make sure to enter a value of at least **16**.
- Click **Request**.
- Within a few minutes you will receive a new **Support Case** for your request. It can take anywhere from 12-48 hours to return a limit increase request.
- You can track the status of your request in the Service Quotas Dashboard, by clicking on **Quota request history** in the navigation panel on the left.
- You will see the status of your request in the list on the right.
- To see more details, click on the **status** value. (Immediately after submitting the request, it will say "Quota requested"). Then in the window that pops up, click on the **Support Center case number**
- If you have a contact at AWS who supports your account, you should also notify them that you have submitted a quota increase request.