

FTV501785

Flame on the Cloud: Remote Production Without Compromising the Quality

Jeffrey Ramirez

Autodesk

Learning Objectives

- Learn how to deploy a Single Autodesk Flame Family product on AWS
- Discover how Flame can collaborate with other Flame Family products (Flare, Flame Assist)
- Learn how to centralize the Project Data and increase your Productivity using Burn Nodes.
- Learn how to implement this into your VFX pipeline or workflow.

Description

With all the challenges we are facing during the pandemic, we've realized how constrained we can be when our access to our production site or office is limited. We learned to adapt with the new challenges and working remotely is one of them. Working remotely may have reduced the quality of our work due to lack of collaboration, technology and speed, to name a few challenges. We also learned that working remotely might be the way many will work going forward.

In this class, we'll discuss how we can help optimize your workflow and increase your collaboration remotely without compromising the quality of your work. Find out how Flame on the Cloud can help you achieve this and incorporate this workflow with your visual effects pipeline.

Speaker

[LinkedIn](#)

Jeffrey Ramirez is a Technical Support Specialist at Autodesk for 8 years, with a total of 17 years of experience in the Film and TV industry. Jeffrey's focus is to help customers resolve some challenges while using our Entertainment & Media solutions specializing with Creative Finishing Products or Visual Effects software like Flame, Flare, Flame Assist and Lustre.

Jeffrey started his career as a System Engineer in the Post-Production for TV series and reality shows. He also worked as a System Integrator for Post-Production and Broadcast department and a Post-Production Engineer with one of the top VFX (Visual Effects) company in South East Asia before joining Autodesk.

Aside from being a Technical Support Specialist in Autodesk, Jeffrey is also:

- a KCS (Knowledge Centered Support) Coach.
- a Geo-Escalation Lead for Creative Finishing Team.

Table of Contents

What is Flame?	3
AWS Instances Type and Configuration for Flame Family product	5
Deploying a Single Autodesk Flame Family product on AWS	6
Steps to deploy this configuration.	7
Collaboration between Flame Family products (Flare, Flame Assist)	8
Steps to deploy this configuration	9
Set Up an AWS Transit Gateway	9
Increase the Productivity using Burn Nodes and a Project Server	10
Project Server Configuration.....	11
Autodesk Burn Configuration	12
Considerations to implement Flame on the Cloud into your VFX pipeline or workflow	13
On premises components.....	14
Capacity and utilization	15
Logistics.....	15
Summary	16
Useful Links	18

What is Flame?

Flame® is a powerful 3D compositing, visual effects, and editorial finishing tool with an integrated environment that accelerates creative workflows.

- Build complex 3D scenes with an intuitive, node-based compositing system
- Deliver commercials, TV episodic, and films with AI-powered, pixel-perfect visual effects
- Perfect every shot with a robust, integrated color grading and finishing toolset

If you are a fan or has been amazed with TV commercials, TV series and Films high on visual effects (VFX), Flame is likely the tool behind them.

Watch the [Flame overview](#) and [Introduction to Flame on the Cloud](#) for more information.

Evolution

Flame was initially deployed on a high-end on-premise hardware, as technology evolves, Flame continues to adapt to take advantage of the newer hardware and software solutions. We have seen Flame being deployed on SGI, PC workstations (Dell, HP, IBM, Lenovo) and Mac.



With this type of on-premise hardware setup, it needs office space, hardware racks, network equipment, and other regular operational expenditures. We also need to be onsite to configure and maintain them.

The Good News

Flame Family Products are now a software only offering. It is not limited to a specific workstation as there are now several options and recommendations in the [Flame System Requirements](#) page for your flexibility including the “self-qualification”.

And what’s more, Flame now runs on the cloud, specifically on AWS (Amazon Web Services). Thanks to the effort of our engineering team that worked closely with [AWS](#) and system integrators to bring artists an alternative to working with Flame. This technology enables us to work leveraging the cloud without compromising the quality of our work.



AWS Instances Type and Configuration for Flame Family product

The table below are the list of the instance type that we will use for the Flame on AWS and the recommended configuration.

Instance name	Configuration	vCPU	GPU	Memory	Storage	Network performance
g4dn.8xlarge	Flame and Burn	32	NVIDIA T4 GPU with 16 GB VRAM	128 GiB	900 GB NVMe SSD	50 Gigabit
g5.8xlarge	Flame and Burn	32	NVIDIA A10G Tensor Cores GPU with 24 GB VRAM	128 GiB	900 GB NVMe SSD	25 Gigabit
r5.xlarge	Project Server	4	N.A.	32 GiB	EBS Only	Up to 10 Gigabit
c5n.9xlarge	NAS Controller	36	N.A.	96GiB	EBS Only	50 Gigabit

Storage	Configuration
AWS 4 x2 TB ST1 EBS	Direct attached
AWS ST1 EBS	NAS
AWS EBS GP3	Project Server

Note: AWS regularly updates their high-performance NVIDIA-based instance types. Consider the preceding as minimum requirements.

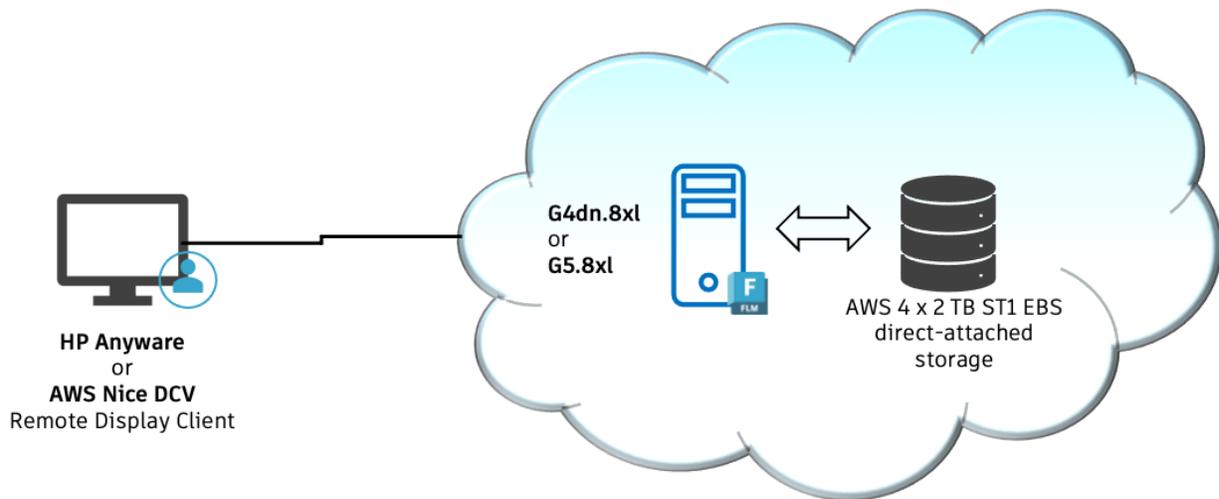
Deploying a Single Autodesk Flame Family product on AWS

This configuration is a great starting point to enable remote workflow leveraging the cloud technology. This is ideal for:

- Artist who mainly work alone and rarely collaborate with other artists.
- Freelancer.

Reference: [Autodesk Flame Family Cloud Reference Implementation Guide](#)

For the Single Flame Family product instance deployed on AWS, the media is stored on direct-attached storage local to each Flame instance. Project metadata is stored on the system disk of each Flame Family product instance.



For this configuration, you need:

- One Flame Family product instance, with its own media storage.
 - High performance GPU-based instances with NVIDIA GPU like **G4dn.8xl** or **G5.8xl**
- **Storage 500 GB** for the system disk.
- **Direct attached storage** using AWS 4 x 2 TB ST1 EBS volumes.
Note: In a standalone instance, or without a project server, the metadata can be stored on the system disk. Media is stored on an attached disk array or cloud NAS.
- Security Group.
- One remote display client for each of the Flame instances. [HP Anyware](#) or [AWS NICE DCV](#).

Steps to deploy this configuration.

1. Create the Amazon Machine Images (AMI) you need using [Creating Autodesk Flame Family AMIs](#)

Note: To simplify the deployment to the cloud, Autodesk provides a Rocky Linux 8.5 AMI. Available from the Flame Family 2023 [Linux \(ISO, AMI\) and DKU page](#), this AMI comes configured with the DKU 17.1, NVIDIA GRID drivers 460.106.00, and all the tools required to use Flame Family in the cloud.

2. Choose and deploy a storage solution. See [Media Storage for Autodesk Flame Family products Cloud Deployment](#).

To store media, you need fast storage capable of high throughput. This storage can be network or direct-attached. For standalone configuration, we will choose direct-attached storage.

- Direct attached storage using AWS 4 x 2 TB ST1 EBS volumes.

Note: When you create the instance, the option **Instance volume deletion on termination** controls what happens to the storage when you terminate the instance. If the volume is not deleted on instance termination, while you're no longer paying for the terminated instance, you're still paying for the storage.

To avoid paying for storage you no longer need, the instance volume can be deleted automatically when the instance is terminated by selecting “**Delete on Termination**”.

3. Create, configure and deploy your [Flame](#) (including Flame Assist, Flare, as required) instance. Flame software installation, media storage configuration, network requirements (such as the Machine ID, hostname, Framestore ID, etc). [See Autodesk Flame Family Product Configuration \(Standalone\)](#) for detailed steps.
4. Connect to your Flame using a remote display solution like [HP Anyware](#) or [AWS NICE DCV](#). Only these 2 solutions are tested so far by Autodesk.
 - **HP Anyware** is one of the remote display solutions tested by Autodesk to connect to Autodesk Flame Family products on the AWS cloud. HP Anyware clients are available for Windows, macOS, and Linux operating systems.
 - **AWS NICE DCV** is a remote display solution provided by AWS and is **free** to use on AWS instances. It is one of the solutions tested by Autodesk to connect remotely to Autodesk Flame on AWS. NICE DCV clients are available for Windows, macOS, and Linux.

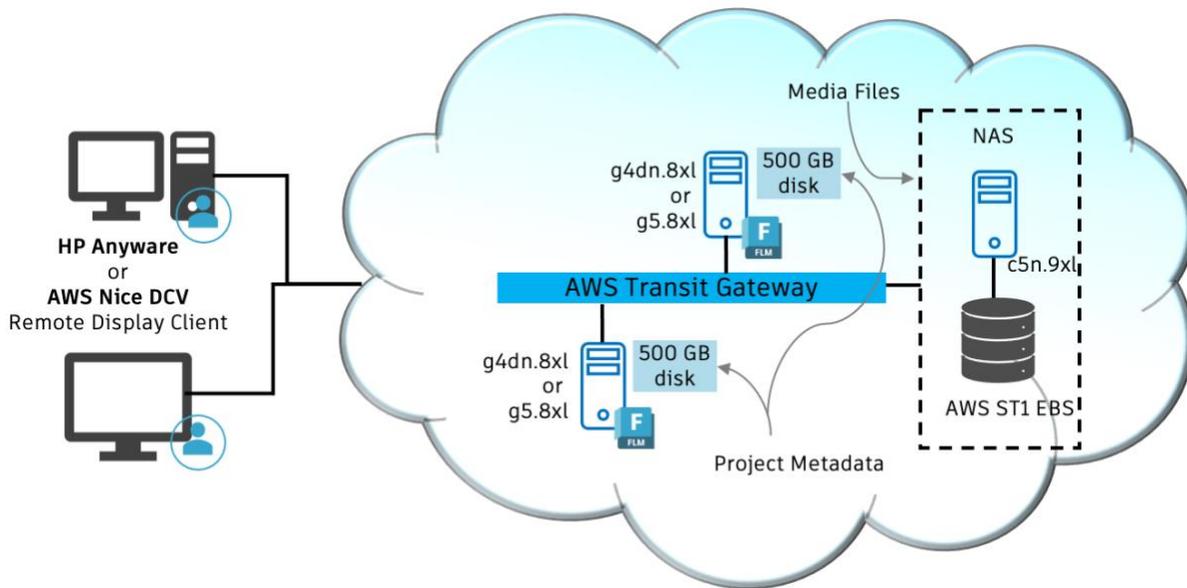
This setup is simpler as it does not require additional instances for NAS, Burn Nodes or Project Server. There will be no collaboration through the network but you can later scale it by setting up a AWS Virtual Private Cloud, which is what we will take a look next.

Collaboration between Flame Family products (Flare, Flame Assist)

This configuration adds NAS (Network-Attached Storage), AWS Virtual Private Cloud and AWS Transit Gateway to enable collaboration. This is ideal for:

- 2 or more artists working on the same project.
- Sharing project between Cloud instances and On-premise workstations.

In this scenario, multiple Autodesk Flame Family product instances are connected to a shared storage, enabling collaboration with local projects or Shared Libraries. Media is stored on a NAS shared with each Flame Family product instance. Each Flame Family product instance stores its project metadata on its system disk.



For this configuration, you need:

- At least 2 Flame Family product instances in the same Virtual Private Cloud.
 - Instance type: At least **G4dn.8xlarge** or **G5.8xlarge**.
- A NAS instance with at least **c5n.9xlarge** media storage.
- AWS Transit Gateway.
- One remote display client for each of the Flame instances. HP Anywhere or AWS NICE DCV.

Steps to deploy this configuration

The steps for deploying this configuration is almost the same with the previous scenario. This time, we need to configure the AWS Virtual Private Cloud Configurations to enable networking with other Flame instances and additional instance for NAS.

1. Create the Amazon Machine Images for every Flame instances. Refer to [Creating Autodesk Flame Family AMIs](#).
2. Configure your AWS cloud using [AWS Virtual Private Cloud Configurations](#). The Virtual Private Cloud (VPC) allows you to network Flame instances, a project server, and Burn nodes together in your cloud implementation.

Set Up an AWS Transit Gateway

To support the various networking capabilities of Flame Family products, you need to configure Transit Gateway service on your instances.

AWS Transit Gateway connects your Amazon Virtual Private Clouds (VPCs) and on-premises networks through a central hub. This simplifies your network and puts an end to complex peering relationships. It acts as a cloud router – each new connection is only made once.

3. Choose and deploy a storage solution. See the guidelines on [Media Storage for Autodesk Flame Family products Cloud Deployment](#).

To store media, you need fast storage capable of high throughput. For this scenario, we will choose NAS (Network-attached storage).

- NAS (Network attached storage) using AWS ST1 EBS volumes, as ST1 is optimized for sequential read-write workload. Refer to [Create a NAS Controller](#) for detailed steps.

Other solutions include:

1. [WekaIO](#)
2. [AWS FSx for OpenZFS](#)
3. [pixitmedia pixstor](#)
4. Create, configure and deploy your [Flame](#) (including Flame Assist, Flare, as required) instance. Flame software installation, media storage configuration, network requirements (such as the Machine ID, hostname, Framestore ID, etc).
5. Connect to your Flame using a remote display solution like [HP Anyware](#) or [AWS NICE DCV](#). Only these 2 solutions are tested so far by Autodesk.

With this configuration, the artists can easily collaborate by sharing the project and media through the network on the clouds and even on-premises workstation.

Increase the Productivity using Burn Nodes and a Project Server

This configuration adds **Burn** and a **Project server** to further enhance the collaboration and productivity. While a Project Server simplifies sharing projects for collaborative workflows, Burn helps by offloading render task from the Flame instance to enable artists to focus on their creative work. This is ideal for:

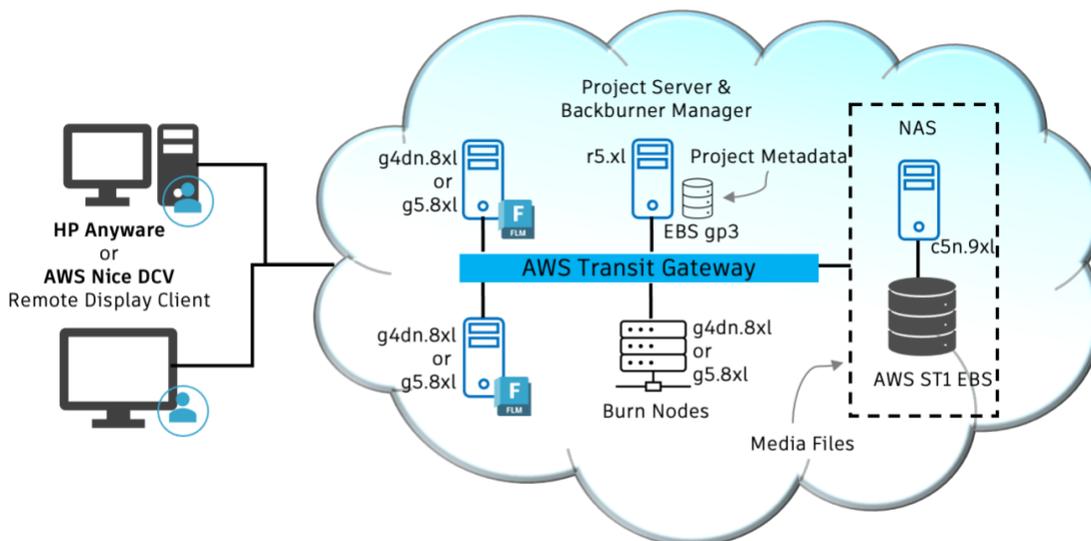
- 2 or more artists working on the same project.
- Sharing project between Cloud instances.
- Artist that need more time to focus with their creative task rather than waiting for the render task to finish on the Flame instance.

Burn allows you to render images in the background to free up Flame workstation for more creative tasks.

A **project server** eases collaboration and simplifies project management by eliminating the creation of project data on the Flame, Flare, or Flame Assist instances. The project data is stored on the centralized project server. There is no data stored on the Flame instances so it can be shutdown with no loss of data.

In this scenario, multiple instances are connected to a shared storage and all project data is created on the project server, enabling collaboration with Shared Libraries. The media is stored on a NAS storage shared with each Flame Family product instance. Project metadata is stored by the project server which is accessible by each Flame Family product instance.

We have to configure an AWS Transit Gateway to make collaboration possible between the Flame Family product instances, the project server, and the burn nodes. As you're using a project server, you must select the Project server when starting Flame Family product (Host Computer).



For this configuration, you need:

- A minimum of two Flame Family product instances.
 - Instance type: At least **G4dn.8xlarge** or **G5.8xlarge**.
- A NAS instance with media storage.
 - Media storage: **AWS ST1 EBS**
- Project Server and Backburner Manager.
 - Instance type: At least **r5.xlarge**.
 - Project storage: **EBS gp3**
- Burn Nodes as required.
- AWS Transit Gateway.
- One remote display client for each of the Flame instances.

Project Server Configuration

Note: Autodesk recommends a maximum of five clients per project server to minimize network congestion and issues with storage quality of service. This can be in any combination of Autodesk Flame Family products and Burn nodes, such as three Flames and two Burn nodes. The number of clients will depend on the Instance type.

As an example, the following minimal setup can expect to serve 3 Flame instances and 2 Burn node instances.

- Media storage: EBS media store
- Project storage: EBS gp3 project storage
- Project server: r5.xlarge for project server

On the other hand, the following, more expansive setup, could deliver up to 16 or 8 Flame instances and 8 Burn node instances:

- Media storage: Weka or FSx for OpenZFS
- Project storage: Striped EBS
- Project server: A more powerful instance with a 50Gbps network

1. Set up a project server instance on AWS. See [Project Server Configuration](#).
 - Instance type: At least **r5.xlarge**. The project server does not require a GPU since it's not used to decode media.

- Storage: You must attach two volumes to this instance: one for the OS and software of at least 20GB, and another for the project metadata storage.
 - The project metadata storage requirements are:
 1. **Volume Type:** gp3
 2. **Size (GiB):** 500
 3. **IOPS:** 3000
 4. **Throughput (MB/s):** set to default (125 MB/s)
 5. To prevent deletion of important project metadata, set the project volume to **not** delete on instance termination.
 - Security Groups:
 - Project Server
 - IGMP-multicast
2. Connect to the instance through a command line.
 3. Add some additional storage to store the projects metadata if required.
 4. Configure the instance as a project server.
 5. Configure the instance to use the networked storage.

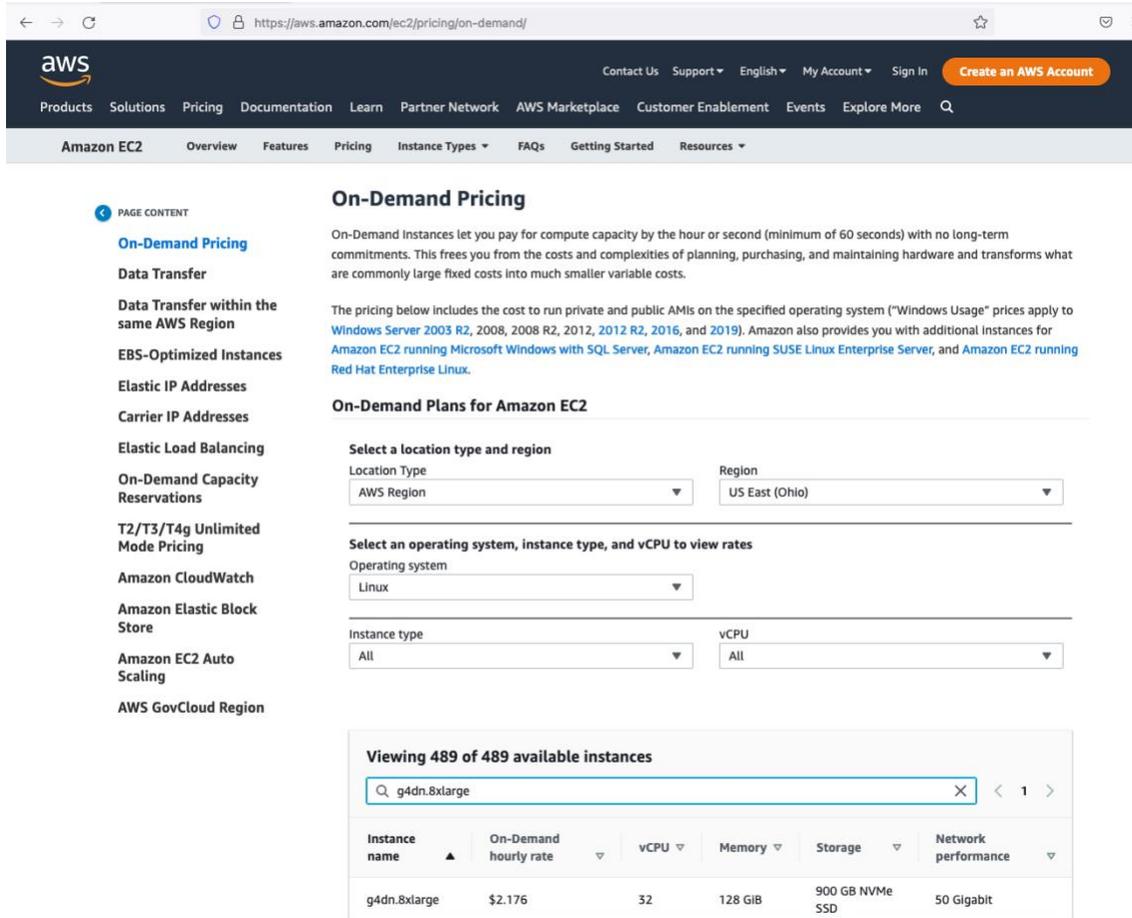
Autodesk Burn Configuration

1. Set up a Burn instance on AWS. See [Autodesk Burn Configuration](#).
 - Instance type: High performance GPU-based instances with NVIDIA GPU. At least **G4dn.8xlarge** or **G5.8xlarge**. *It must match your Flame instance type.*
 - Storage: 20 GB for the system disk is enough. Media is stored on an attached disk array or cloud NAS.
 - Security Group:
 - Burn
 - IGMP-multicast
2. Connect to the instance through a command line.
3. Configure the instance as a Burn node.
4. Configure the instance to use the networked storage.

With the use of Burn and Project Server, the work load of the Flame instances are reduced. If Burn is not required at some point, the instances can be stopped to reduce the cost.

Considerations to implement Flame on the Cloud into your VFX pipeline or workflow

To give you an idea about the AWS instances cost, here are the cost for some of the instances we used in the previous pages. This information are gathered from the [Amazon EC2 On-Demand Pricing](https://aws.amazon.com/ec2/pricing/on-demand/) website.



On-Demand Pricing

On-Demand Instances let you pay for compute capacity by the hour or second (minimum of 60 seconds) with no long-term commitments. This frees you from the costs and complexities of planning, purchasing, and maintaining hardware and transforms what are commonly large fixed costs into much smaller variable costs.

The pricing below includes the cost to run private and public AMIs on the specified operating system ("Windows Usage" prices apply to [Windows Server 2003 R2](#), 2008, 2008 R2, 2012, [2012 R2](#), [2016](#), and [2019](#)). Amazon also provides you with additional Instances for [Amazon EC2 running Microsoft Windows with SQL Server](#), [Amazon EC2 running SUSE Linux Enterprise Server](#), and [Amazon EC2 running Red Hat Enterprise Linux](#).

On-Demand Plans for Amazon EC2

Select a location type and region

Location Type: AWS Region | Region: US East (Ohio)

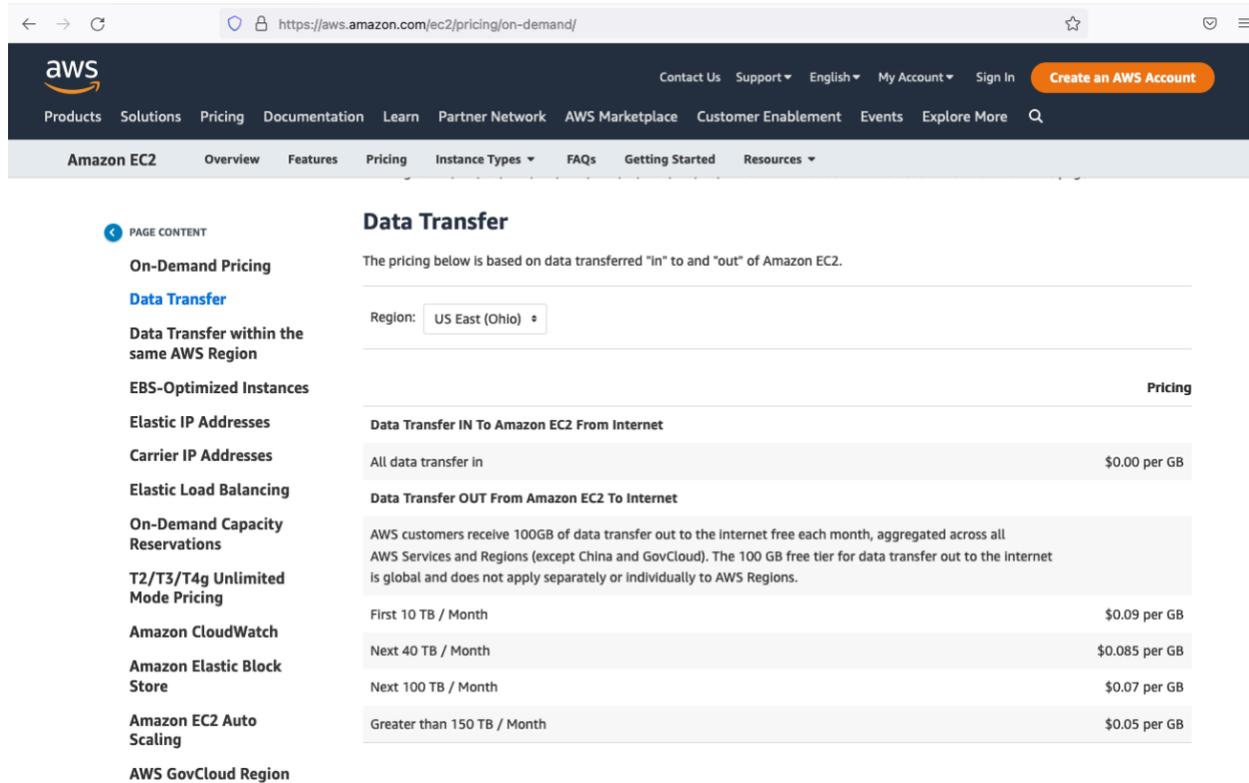
Select an operating system, instance type, and vCPU to view rates

Operating system: Linux | Instance type: All | vCPU: All

Viewing 489 of 489 available instances

Search: g4dn.8xlarge

Instance name	On-Demand hourly rate	vCPU	Memory	Storage	Network performance
g4dn.8xlarge	\$2.176	32	128 GiB	900 GB NVMe SSD	50 Gigabit



Data Transfer

The pricing below is based on data transferred "in" to and "out" of Amazon EC2.

Region:

	Pricing
Data Transfer IN To Amazon EC2 From Internet	
All data transfer in	\$0.00 per GB
Data Transfer OUT From Amazon EC2 To Internet	
AWS customers receive 100GB of data transfer out to the internet free each month, aggregated across all AWS Services and Regions (except China and GovCloud). The 100 GB free tier for data transfer out to the internet is global and does not apply separately or individually to AWS Regions.	
First 10 TB / Month	\$0.09 per GB
Next 40 TB / Month	\$0.085 per GB
Next 100 TB / Month	\$0.07 per GB
Greater than 150 TB / Month	\$0.05 per GB

Note: The instance availability varies depending on the region. The prices above are based on US West (Oregon) region). For more information on the pricing visit [Amazon EC2 On-Demand Pricing](#). For cost estimation, use [AWS Calculator](#).

Cloud computing is a significant shift from a traditional on-premises infrastructure. It is understandable that we weigh the benefits or the advantages and disadvantages before making our decision to adapt Flame on the Cloud in our workflow. Here are some key considerations to help us decide.

On premises components

1. Hardware costs
 - Server – including the workstation, rack, cables, spare parts, etc.
 - Storage – Disks, network cards, cables.
 - Network – Network switches, router, cable, ISP bandwidth costs.
 - 5-year upgrade cycle (usual refresh cycle).
2. Software costs
 - Operating System.
 - Licenses and subscriptions.
 - Management software.
 - Software upgrade.

3. Facilities costs
 - Server and workstation space.
 - Power and utilities.
 - Cooling / Air conditioning.
4. Manpower costs.
 - IT Technical support.
 - Facilities management.

Capacity and utilization

1. How many users are required. The cost reduces when the instance is idle or if not running. Some facilities invests with numbers of workstations and servers but there will be time that they are underutilize.
2. How long the workstation is needed. There are some project that will run within certain period only. For example, working on a movie project, short-film, advertisement, etc.
3. How many workstation are needed. Since AWS instance is on-demand, the quantity is scalable.
4. How much storage are needed. Since AWS instance is on-demand, the quantity is scalable.

Logistics

1. Travel cost for the user/artist. Does the user need to be on-site?
2. Shipment cost for the workstation. Some clients requires the user/artist and their Flame Workstation to be on-site.

Note: This list is the basic operating cost. This may vary in every facilities. For more information, see [AWS Cloud Financial Management](#).

These are just some of the considerations that we need to think about to help us decide whether on-premises are still viable for us, or if we can add **Flame on the Cloud** to our workflow, or maybe fully shift to the cloud.

Summary

Remote workflow will continue to be the way many will work going forward. Flame on the Cloud gives us the opportunity to leverage on the new technology that will help optimize your remote workflow. With its **speed, power, accessibility, scalability, security**, among many other benefits, you can now experience the full performance of Flame from about anywhere. Adopting this technology will benefit your organization with a broader business opportunities. With Flame deployed on the Cloud and essentially accessible from anywhere, an organization can have the ability to work from almost anywhere. It gives you the flexibility to recruit the finest talent around the world and collaborate with each other wherever they are. This is a testimony from one of early pioneer and adopter of Flame on AWS cloud, [Preymaker](#) founder Angus Kneale.

"Preymaker is all about having the finest talent using the best technology, and running Flame in AWS allows us to recruit and work with exceptional talent who live anywhere. Having Flame projects live in the cloud, with artists collaborating in multiple locations, we are able to create exceptional work for our clients. Our colorist in Los Angeles can start a project, with our Flame artist in London doing the conform, ready for our CGI team in New York to continue work. Ultimately, the cloud gives us the flexibility to execute highly complicated, demanding, and compute-intensive projects in a collaborative cloud-based workflow."

The Flame team has provided some pre-build components like AMI and the [Autodesk Flame Family Cloud Reference Implementation Guide](#) to help you get started. We also have resellers and AWS enabled System Integrators (SIs) that have successfully deployed Flame on the cloud and are equipped to help with your workflow, deployment, and configuration needs.

Equipped resellers:

- [CineSys.io](#)
- [Escape Technology](#)
- [Visual Graphics Inc.](#)

Equipped AWS enabled System Integrators (SIs):

- [Arch Platform Technologies](#)
- [Gunpowder](#)
- [Konsistent Consulting](#)
- [TrackIt](#)

Hopefully, this presentation helps you with your workflow. If you have questions regarding this class, please use the comment section in the AU page and I will try my best to answer as soon as I can. If you like this class, please help me to share to your peers and click the recommend icon. You may also visit the useful links below.

-  Add to playlist 
-  Share
-  Comment
-  Bookmark
-  Recommend
- [+ Follow](#)

Comments



Ask questions or comment here.



We have a be nice policy
Please be positive and constructive

POST

Useful Links

- [Autodesk Flame Family Cloud Reference Implementation Guide](#)
- [Running Flame Family on the Cloud FAQ](#)
- [Flame Family System Requirements page](#)
- [Introducing Autodesk Flame on AWS](#)
- [AWS Cloud Financial Management](#)
- [Edit in the Cloud on AWS](#)