

# File Gateway Service

UCSF IT

# UCSF IT File Gateway Service

- What is it?
- How much does it cost?
- How do I get it?
- How do I use it?
- How do I get support?

# What is the File Gateway Service?

- The file gateway service presents a Windows file share inside the UCSF Network. This file share can be made accessible to computers on the network via on-premises physical or WiFi networks, VPN or to EC2 instances in the AWS Secure Enterprise Cloud.
- Each file share is backed by AWS S3 object storage.
- All objects in S3 are encrypted.
- The S3 storage classes currently available are: S3 Standard, S3 Infrequent Access and S3 One Zone Infrequent Access.
- The main use cases: cold storage, large files, images, archive files, and migrating data to AWS for access by EC2 instances.
- File access control is governed by active directory.
- All user access is logged in AWS Cloudwatch logs in the event forensic discovery is needed.
- Direct AWS access is not permitted. Access to the files/objects must happen via the file gateway.
- In order to meet security compliance bucket versioning is on and must remain on. This means that if a file is updated, replaced or deleted a new version is created. Because of this small actively edited files are not advised as this will cause escalating costs.
- Deleting a file does not delete the data. The version of the file presented in the file share remains in the version history. This needs to be considered when estimating costs as the versions stored do incur costs at the same rate.
- The storage and access costs are recharged to the sponsor departments cost center.
- Costs are passed thru directly to the department and are set by AWS. <https://aws.amazon.com/s3/pricing/>
- The gateway is a virtual appliance hosted either on premises or in the cloud and each gateway can host ten independent shares each pointing to different buckets.

# How much will it cost?

- **Hosting costs:** The file gateway can host ten shares and each share incurs a cost equal to 1/10 of the hosting cost for the virtual appliance. For the base class gateway this is \$62 per month per share. If more performance is needed other options are available.
- **Storage costs:** the storage costs are passed directly from AWS at the rate they charge UCSF. This includes applicable and available discounts.
  - Standard: ~\$20.02 per TB per month
  - Infrequent Access: ~\$10.88 per TB per month
  - One zone Infrequent Access (reduced durability): ~\$8.70 per TB per month
- **Retrieval costs:** S3 Standard has no retrieval costs. IA and 1zIA cost ~\$8.70 per TB to retrieve.
- **Access request costs:**
  - IA and 1zIA \$0.01 per 1000 requests.
  - Standard \$0.005 per 1000 PUT \$0.0004 per 1000 GET

# Cost Example A

- Let's say you have 100 TB of images with an average image size of 5 GB for a total of about 20,000 images. These images need to be stored for 5 years and it is unlikely that they will be retrieved regularly. This set of images is a second copy and not the primary source of data. Estimated retrievals will be 1000 the first year. 500 the second year. 250 the third. 100 for the fourth and 50 on the fifth year. In this scenario we will use One Zone Infrequent access.

Cost Type	Per Month	Year 1	Year 2	Year 3	Year 4	Year 5	Total
Share Hosting	\$62	\$744	\$744	\$744	\$744	\$744	<b>\$3720</b>
Storage	\$850	\$10,200	\$10,200	\$10,200	\$10,200	\$10,200	<b>\$51,000</b>
Request	N/A	\$0.21	\$0.01	\$0.01	\$0.01	\$0.01	<b>\$0.05</b>
Retrieval	N/A	\$42.50	\$21.25	\$10.62	\$4.25	\$2.13	<b>\$80.75</b>
<b>Total</b>		<b>\$10,986.51</b>	<b>\$10,965.26</b>	<b>\$10,954.63</b>	<b>\$10,948.26</b>	<b>\$10,946.14</b>	<b>\$54,800.80</b>

# Cost Example B-1

- Let's say you have 10 TB of text files with an average size is 5 MB for a total of about 2,000,000 text files. These files need to be stored for 5 years and some of them will need to be retrieved regularly. This set of files is a primary copy. Estimated retrievals will be 500,000 files each year. In this scenario we will use S3 Standard storage class.

Cost Type	Per Month	Year 1	Year 2	Year 3	Year 4	Year 5	Total
Share Hosting	\$62	\$744	\$744	\$744	\$744	\$744	<b>\$3720</b>
Storage	\$190.92	\$2,291.04	\$2,291.04	\$2,291.04	\$2,291.04	\$2,291.04	<b>\$11,455.20</b>
Request	N/A	\$8.67	\$0.17	\$0.17	\$0.17	\$0.17	<b>\$9.35</b>
Retrieval	N/A	\$0	\$0	\$0	\$0	\$0	<b>\$0</b>
<b>Total</b>		<b>\$3,043.71</b>	<b>\$3,035.21</b>	<b>\$3,035.21</b>	<b>\$3,035.21</b>	<b>\$3,035.21</b>	<b>\$15,184.55</b>

# Cost Example B-2

- Let's say you have 10 TB of text files with an average size is 5 MB for a total of about 2,000,000 text files. These files need to be stored for 5 years and some of them will need to be retrieved regularly. This set of files is a primary copy. Estimated retrievals will be 500,000 files each year. In this scenario we will use Infrequent Access storage class.

Cost Type	Per Month	Year 1	Year 2	Year 3	Year 4	Year 5	Total
Share Hosting	\$62	\$744	\$744	\$744	\$744	\$744	<b>\$3720</b>
Storage	\$103.76	\$1,245.12	\$1,245.12	\$1,245.12	\$1,245.12	\$1,245.12	<b>\$6,225.60</b>
Request	N/A	\$20.50	\$0.50	\$0.50	\$0.50	\$0.50	<b>\$22.50</b>
Retrieval	N/A	\$24.41	\$24.41	\$24.41	\$24.41	\$24.41	<b>\$122.05</b>
<b>Total</b>		<b>\$2,034.03</b>	<b>\$2,014.03</b>	<b>\$2,014.03</b>	<b>\$2,014.03</b>	<b>\$2,014.03</b>	<b>\$10,090.15</b>

# Cost Example B-3 – inappropriate storage class

- Let's say you have 10 TB of text files with an average size is 5 MB for a total of about 2,000,000 text files. These files need to be stored for 5 years and ALL of them will need to be retrieved three times a month on average. This set of files is a primary copy. Estimated retrievals will be 72,000,000 files each year. In this scenario we will use Infrequent Access storage class.

Cost Type	Per Month	Year 1	Year 2	Year 3	Year 4	Year 5	Total
Share Hosting	\$62	\$744	\$744	\$744	\$744	\$744	<b>\$3720</b>
Storage	\$103.76	\$1,245.12	\$1,245.12	\$1,245.12	\$1,245.12	\$1,245.12	<b>\$6,225.60</b>
Request	N/A	\$92	\$72	\$72	\$72	\$72	<b>\$22.50</b>
Retrieval	N/A	\$2,988.28	\$2,988.28	\$2,988.28	\$2,988.28	\$2,988.28	<b>\$14,941.40</b>
<b>Total</b>		<b>\$5,069.4</b>	<b>\$5,049.40</b>	<b>\$5,049.40</b>	<b>\$5,049.40</b>	<b>\$5,049.40</b>	<b>\$24,909.5</b>

# How to obtain access to a hosted file share?

- The costs incurred by using the File Gateway service are recharged directly to the cost center of the department using the service. In order to properly conduct the recharge a CloudPay Account needs to be obtained from the Cloud Business Office.
  - [Open a CloudPay account here.](#)
  - (Please have a copy of your submission emailed to you for your records)
- After applying for CloudPay account a File Share request must be filled out for each share desired.
  - [For S3 Standard](#)
  - [Infrequent Access](#)
  - [One Zone Infrequent Access](#)
- Once approved the Technical Contact for the CloudPay account will be provided access to the share and will be responsible for managing access permissions for the data stored in the share.

# How to use the file share?

- The file share can be accessed with the standard method for the OS being used:
  - [Windows](#), [MacOS](#), [Linux](#)
- The Technical Contact will need to grant access to the individuals who need to use the share and will be responsible for submitting any needed firewall requests: <https://it.ucsf.edu/how-to/how-submit-firewallvpn-requests>
- Users and applications can then read and write to the share. For data no longer contained in the local cache there will be a delay for the file to open as it is pulled back across the network from S3. This pull will result in retrieval costs if IA or 1zIA storage classes are used.

# How to get support?

- For help in understanding the service and for help with the process please submit a consultation request with the [Customer Solutions team](#).
- Once the share is live support on the resource can be obtained submitting a [service ticket](#).
- Or email: [cloudsupport@ucsf.edu](mailto:cloudsupport@ucsf.edu)