

Amazon Elastic File System

Choosing Between the Different
Throughput & Performance Modes

July 2018

This paper has been archived

For the latest technical content, refer to the AWS
Whitepapers & Guides page:

<https://aws.amazon.com/whitepapers>



Notices

This document is provided for informational purposes only. It represents AWS's current product offerings and practices as of the date of issue of this document, which are subject to change without notice. Customers are responsible for making their own independent assessment of the information in this document and any use of AWS's products or services, each of which is provided "as is" without warranty of any kind, whether express or implied. This document does not create any warranties, representations, contractual commitments, conditions or assurances from AWS, its affiliates, suppliers or licensors. The responsibilities and liabilities of AWS to its customers are controlled by AWS agreements, and this document is not part of, nor does it modify, any agreement between AWS and its customers.

This paper has been archived

For the latest technical content, refer to the AWS
Whitepapers & Guides page:

<https://aws.amazon.com/whitepapers>

Contents

Introduction	1
Performance Modes	1
General Purpose	1
Max I/O	1
Selecting the right performance mode	2
Throughput Modes	3
Bursting Throughput	3
Provisioned Throughput	4
Selecting the right throughput mode	5
Conclusion	6
Contributors	6
Further Reading	7
Document Revisions	7

This paper has been archived

For the latest technical content, refer to the AWS
Whitepapers & Guides page:

<https://aws.amazon.com/whitepapers>

Abstract

Storage types can generally be divided into three different categories: block, file, and object. Each storage type has made its way into the enterprise and a large majority of data resides on file storage. Network shared file systems have become a critical storage platform for businesses of any size. These systems are accessed by a single client or multiple (tens, hundreds, or thousands) concurrently so they can access and use a common data set. Amazon Elastic File System (Amazon EFS) satisfies these demands and gives customers the flexibility to choose different performance and throughput modes that best suits their needs. This paper outlines the best practices for running network shared file systems on the AWS cloud platform and offers guidance to select the right Amazon EFS performance and throughput modes for your workload.

This paper has been archived

For the latest technical content, refer to the AWS
Whitepapers & Guides page:

<https://aws.amazon.com/whitepapers>

Introduction

[Amazon Elastic File System \(Amazon EFS\)](#)¹ provides simple, scalable, elastic file storage for use with AWS Cloud services and on-premises resources. The file systems you create using Amazon EFS are elastic, growing and shrinking automatically as you add and remove data. They can grow to petabytes in size, distributing data across an unconstrained number of storage servers in multiple Availability Zones. Amazon EFS supports Network File System version 4 (NFSv4.0 & 4.1), provides POSIX file system semantics, and guarantees open-after-close semantics.

Amazon EFS is a regional service built on a foundation of high availability and high durability and is designed to satisfy the performance and throughput demands of a wide spectrum of use cases and workloads, including web serving and content management, enterprise applications, media and entertainment processing workflows, home directories, database backups, developer tools, container storage, and big data analytics. EFS file systems provide customizable performance and throughput options so you can tune your file system to match the needs of your application.

Performance Modes

Amazon EFS offers two performance modes: **General Purpose** and **Max I/O**. You can select one when creating your file system. There is no price difference between the modes, so your file system is billed and metered the same. The performance mode can't be changed after the file system has been created.

For the latest technical content, refer to the AWS
General Purpose Whitepapers & Guides page:

<https://aws.amazon.com/whitepapers>
General Purpose is the default performance mode and is recommended for the majority of uses cases and workloads. It is the most commonly used performance mode, and is ideal for latency-sensitive applications, like web serving, content management systems, and general file serving. These file systems experience the lowest latency per file system operation and can achieve this for random or sequential IO patterns. There is a limit of 7,000 file system operation per second aggregated across all clients for General Purpose performance mode file systems.

Max I/O

File systems created in Max I/O performance mode can scale to higher levels of aggregate throughput and operations per second when compared to General Purpose file systems. These file systems are designed for highly parallelized applications like big data analytics, video transcoding and processing, and genomic analytics, which can scale out to tens, hundreds, or thousands of Amazon EC2 instances. Max I/O file systems do not have a 7,000 file system operation per second limit but latency per file system operation is slightly higher when compared to General Purpose performance mode file systems.

Selecting the right performance mode

We recommend creating the file system in the default General Purpose performance mode and testing your workload for a period of time to test its performance. We provide eight Amazon CloudWatch metrics per file system to help you understand how your workload is driving the file system. One of these metrics, *PercentIOLimit*, is specific to General Purpose performance mode file systems and indicates, as a percent, how close you are to the 7,000 file system operations per second limit.

If the *PercentIOLimit* value returned is at or near 100 percent for a significant amount of time during your test (see figure 1), we recommend you use a Max I/O performance mode file system. To move to a different performance mode, you migrate the data to a different file system that was created in the other performance mode. You can use Amazon EFS File Sync to migrate the data. For more information on Amazon EFS File Sync, please refer to the [Amazon EFS File Sync](#) section of the *Amazon EFS User Guide*.²




Figure 1

There are some workloads that need to scale out to the higher I/O levels provided by Max I/O performance mode but are also latency-sensitive and require the lower latency provided by General Purpose performance mode. In situations like this and if the workload and




Figure 2

applications support it, we recommend creating multiple General Purpose performance mode file systems and spread the application workload across all these file systems. This would allow you to create a logical file system and shard data across multiple EFS file systems. Each file system would be mounted as a subdirectory and the application can access these subdirectories in parallel (see figure 2). This allows latency-sensitive workloads to scale to higher levels of file system operations per second, aggregated across multiple file systems, and at the same time take advantage of the lower latencies offered by General Purpose performance mode file systems.

Throughput Modes

The throughput mode of the file system helps determine the overall throughput a file system is able to achieve. You can select the throughput mode at any time (subject to daily limits). Changing the throughput mode is a non-disruptive operation and can be run while clients continuously access the file system. You can choose between two throughput modes, **Bursting** or **Provisioned**. There are price and throughput level differences between the two modes so understanding each one, their differences, and when to select one throughput mode over the other is valuable.

Bursting Throughput

Bursting Throughput is the default mode and is recommended for a majority of uses cases and workloads. Throughput scales as your file system grows and you are billed only for the amount of data stored on the file system in GB-Month. Because file-based workloads are typically spiky – driving high levels of throughput in short periods of time, and low levels of throughput the rest of the time – file systems using Bursting Throughput mode allow for high throughput levels for a limited period of time. All Bursting Throughput mode file systems, regardless of size, can burst up to 100 MiB/s. Throughput also scales as the file system grows and will scale at the bursting throughput rate of 100 MiB/s per TiB of data stored, subject to regional default file system throughput limits. These bursting throughput numbers can be achieved when the file system has a positive burst credit balance. You can monitor and alert on your file system's burst credit balance using the *BurstCreditBalance* file system metric in Amazon CloudWatch. File systems earn burst credits at the baseline throughput rate of 50 MiB/s per TiB of data stored and can accumulate burst credits up to the maximum size of 2.1 TiB per TiB of data stored. This allows larger file systems to accumulate and store more burst credits which allows them to burst for longer periods of time. If the file system's burst credit balance is ever depleted, the permitted throughput becomes the baseline throughput. Permitted throughput is the maximum amount of throughput a file system is allowed and this value is available as an Amazon CloudWatch metric.

Provisioned Throughput

Provisioned Throughput is available for applications that require a higher throughput to storage ratio than those allowed by Bursting Throughput mode. In this mode you can provision the file system's throughput independent of the amount of data stored in the file system. This allows you to optimize your file system's throughput performance to match your application's needs, and your application can drive up to the provisioned throughput continuously.

This concept of provisioned performance is similar to features offered by other AWS services like provisioned IOPS for Amazon Elastic Block Store PIOPS (io1) volumes and provisioned throughput with read and write capacity units for Amazon DynamoDB. As with these services, you are billed separately for the performance or throughput you provision and the storage you use, e.g. two billing dimensions. When file systems are running in Provisioned Throughput mode, you are billed for the storage you use in GB-Month and for the throughput provisioned in MiB/s-Month. The storage charge, for both Bursting and Provisioned Throughput modes includes the baseline throughput of the file system in the price of storage. This means the price of storage includes 1 MiB/s of throughput per 20 GiB of data stored so you will be billed for the throughput you provision above this limit. For more information on pricing, see the [Amazon EFS pricing page](#).³

You can increase Provisioned Throughput as often as you need. You can decrease Provisioned Throughput or switch throughput modes as long as it's been more than 24 hours since the last decrease or throughput mode change.

File systems continuously earn burst credits up to the maximum burst credit balance allowed for the file system. The maximum burst credit balance is 2.1 TiB for file systems smaller than 1 TiB, or 2.1 TiB per TiB stored for file systems larger than 1 TiB. File systems running in Provisioned Throughput mode still earn burst credits. They earn at the higher of the two rates, either the Provisioned Throughput rate or the baseline Bursting Throughput rate of 50 MiB/s per TiB of storage.

You could find yourself in the situation where your file system is running in Provisioned Throughput mode and over time the size of it grows so that its provisioned throughput is less than the baseline throughput it is entitled to had the file system been in Bursting Throughput mode. In a case like this, you will be entitled to the higher throughput of the two modes, including the burst throughput of Bursting Throughput mode, and you will not be billed for throughput above the storage price. For example, you set the provisioned throughput of your 1 TiB file system to 200 MiB/s. Over time the file system grows to 5 TiB. A file system in Bursting Throughput mode would be entitled to a baseline throughput of 50 MiB/s per TiB of data stored and a burst throughput of 100 MiB/s per TiB of data stored. Though your file system is still running in Provisioned Throughput mode, its entitled to a baseline throughput of 250 MiB/s and a burst throughput of 500 MiB/s and will only incur a storage charge for a 5 TiB file system.

For information on maximum provisioned throughput limits, please refer to the [Amazon EFS Limits](#) section of the *Amazon EFS User Guide*.⁴

Selecting the right throughput mode

We recommend running file systems in Bursting Throughput mode because it offers a simple and scalable experience that provides the right ratio of throughput to storage capacity for most workloads. There are times when a file system needs a higher throughput to storage capacity ratio than what is offered by Bursting Throughput mode. Knowing the throughput demands of your application or monitoring key indicators are two important ways in determining when you'll need these higher levels of throughput.

We recommend using Amazon CloudWatch to monitor how your file system is performing. One of these metrics, *BurstCreditBalance*, is a key performance indicator that will help determine if your file system is better suited for Provisioned Throughput mode. If this value is zero or steadily decreasing over a period of normal operations (see *figure 3*), your file system is consuming more burst credits than it is earning. This means your workload requires a throughput to storage capacity ratio greater than what is allowed by Bursting Throughput mode. If this occurs, we recommend provisioning throughput for your file system. This can be done by modifying the file system to change the throughput mode using the AWS Management Console, AWS CLIs, AWS SDKs, or EFS APIs. When choosing to run in Provisioned Throughput mode, you must also indicate the amount of throughput you want to provision for your file system. To help determine how much throughput to provision, we recommend monitoring another key performance indicator available from Amazon CloudWatch, *TotalIOBytes*. This metric represents the sum of all numbers of bytes (data read, data write, and metadata) for each file system operation during a selected period. To calculate the average throughput in MiB/s for a period, convert the Sum statistic to MiB (Sum of *TotalIOBytes* x 1,048,576) and divide by the number of seconds in the period. Use Metric Math expressions in Amazon CloudWatch to make it even easier to see throughput in MiB/s. For more information on using Metric Math, see [Using Metric Math with Amazon EFS](#) in the *Amazon EFS User Guide*.⁵ Calculate this during the same period when your *BurstCreditBalance* metric was continuously decreasing. This will give you the average throughput you were




Figure 3

of *TotalIOBytes* x 1,048,576) and divide by the number of seconds in the period. Use Metric Math expressions in Amazon CloudWatch to make it even easier to see throughput in MiB/s. For more information on using Metric Math, see [Using Metric Math with Amazon EFS](#) in the *Amazon EFS User Guide*.⁵ Calculate this during the same period when your *BurstCreditBalance* metric was continuously decreasing. This will give you the average throughput you were

achieving during this period and is a good starting point when choosing the amount of throughput to provision.

If your file system is running in Provisioned Throughput mode and you experience no performance issues while your *BurstCreditBalance* continuously increases for long periods of normal operations, then consider decreasing the amount of provisioned throughput to reduce costs. To help determine how much throughput to provision, we also recommend monitoring the Amazon CloudWatch metric *TotalIOBytes*. Calculate this during the same period when your *BurstCreditBalance* metric was continuously increasing. This will give you the average throughput you were achieving during this period and is a good starting point when choosing the amount of throughput to provision. Remember, you can increase the amount of provisioned throughput as often as you need but you can only decrease the amount of provisioned throughput or switch throughput modes as long as it's been more than 24 hours since the last decrease or throughput mode change.

If you're planning on migrating large amounts of data into your file system, you may also want to consider switching to Provisioned Throughput mode and provision a higher throughput beyond your allotted burst capability to accelerate loading data. Following the migration, you may decide to lower the amount of provisioned throughput or switch to Bursting Throughput mode for normal operations. Monitor the average total throughput of the file system using the *TotalIOBytes* metric in Amazon CloudWatch. Use Metric Math expressions in Amazon CloudWatch to make it even easier to see throughput in MiB/s. Compare the average throughput you're driving the file system to the *PermittedThroughput* metric. If the calculated average throughput you're driving the file system is less than the permitted throughput, then consider making a throughput change to lower costs. If the calculated average throughput during normal operations is at or below the baseline throughput to storage capacity ratio of Bursting Throughput mode (10 MiB/s per TiB of data stored), then consider switching to Bursting Throughput mode. If the calculated average throughput during normal operations is above this ratio, then consider lowering the amount of provisioned throughput to some level in between your current provisioned throughput and the calculated average throughput during normal operations. Remember, you can switch throughput modes or decrease the amount of provisioned throughput as long as it's been more than 24 hours since the last decrease or throughput mode change.

Conclusion

Amazon EFS gives you the flexibility to choose different performance and throughput modes to customize your file system to meet the needs for a wide spectrum of workloads. Knowing the performance and throughput demands of your application and monitoring key performance indicators will help you select the right performance and throughput mode to satisfy your file system's needs.

Contributors

The following individuals and organizations contributed to this document:

- Darryl S. Osborne, solutions architect, Amazon File Services

Further Reading

For additional information, see the following:

- [Amazon EFS User Guide](#)⁶

Document Revisions

Date	Description
July 2018	First publication

¹ <https://aws.amazon.com/efs/>

² <https://docs.aws.amazon.com/efs/latest/ug/get-started-file-sync.html>

³ <https://aws.amazon.com/efs/pricing/>

⁴ <https://docs.aws.amazon.com/efs/latest/ug/limits.html>

⁵ <https://docs.aws.amazon.com/efs/latest/ug/monitoring-metric-math.html>

⁶ <https://docs.aws.amazon.com/efs/latest/ug/whatisefs.html>

This paper has been archived
For the latest technical content, refer to the AWS
Whitepapers & Guides page:

<https://aws.amazon.com/whitepapers>