



STORAGE WITHOUT BOUNDS



Compute,
Network
& Storage

Published Q2 2023

TABLE OF CONTENTS

Introduction	3
Understanding the Storage product line	4
Our purpose	10
The evolution of storage	12
Data growth	16
Storage types	17
The business of Storage	27
Our OKRs	28
Managing resources (Financials, Budget, and Headcount)	32
Securing our data	40
Principles of SRE	45
Understanding our users	46
The people of Storage	51
Meet the team	53
Appendix	76
Glossary	77
Storage hardware	82

INTRODUCTION

The Storage product line is responsible for all aspects of data storage at JPMorgan Chase, managing over 400 petabytes of on-premises and public cloud data. This in turn helps the firm serve tens of millions of clients across its lines of business.

The groundbreaking applications that our developers build use vast amounts of data, and that data must be housed in an easy-to-access, reliable, and secure environment.

Innovative solutions from Storage make this possible.

Continue reading to learn more.





Our mission: Deliver cutting-edge storage solutions that give our LOB partners access to the latest technological innovations, so that they can provide the best financial services in the industry.



UNDERSTANDING THE STORAGE PRODUCT LINE

OUR PURPOSE



Data is the engine of modern business. The world's most innovative organizations leverage data to separate themselves from the pack, producing and consuming huge amounts of information to lower costs, boost productivity, make hiring decisions, and improve employee and customer satisfaction. JPMorgan Chase is at the forefront of those companies pushing the envelope. By managing the firm's software and hardware-related storage infrastructure, the Storage product line plays a key role on this front, evolving our technology to handle these enormous data sets with a combination of on-premises, private cloud, and public cloud storage product offerings.

Key themes to continue delivering our purpose:




- Continue to increase velocity of public cloud storage features
- Maintain a stable and secure storage operational environment
- Make storage easy to use and feature-rich
- Ensure private cloud expenses are tightly controlled
- Make it easy for our customers to manage expenses of public cloud storage
- Train, certify, and expose our people to new learning opportunities giving them fulfilling careers within Storage and, with mobility, the wider GT

Top priorities against these themes:

- Reduce public cloud backlog to fewer than 40 items and complete important epics
- Retire spinning media (ScaleIO) to zero by December 2023 and continue to reduce Fibre Channel by 4% quarter on quarter
- Pen additional long-term, well-structured contracts with strategic vendors to contain expenses related to legacy datacenter exit dates
- Introduce new strategic block product to address supply chain challenges and enable future protocol and performance standards by June 30th
- Educate our application developer teams to be more cost efficient with a cost-avoidance goal of \$20M over the next year
- Continued growth of AWS trained workforce to 100% by year end (from 36% today)
- Be "always prepared" so that the product line achieves a satisfactory rating across all audits in 2023



CHOOSING THE RIGHT STORAGE



There are two types of data: structured and unstructured. Structured data is highly organized information, like names and addresses, that is stored in lists and is easy to search by humans or by program algorithms. Unstructured data is the opposite: information that isn't highly organized, like photos and videos, emails, and biometrics. The use of unstructured data has exploded in recent years, making the need for fast, efficient storage more important than ever.

When weighing different storage options, it is important to first determine if the data being stored is structured, unstructured, or semi-structured, as well as the individual needs of business applications. All products will not have the same data requirements. It's important to understand which storage options best suit various use cases. JPMC offers three storage types: object, block, and file.

Amount of storage

Organizations that produce large amounts of data will likely choose block or object storage as file storage becomes challenging when managing and retrieving a greater number of files.

Performance

What amount of latency can an application tolerate before performance suffers? All three storage options have low latency, but block offers the fastest performance.

Cost

Cost for block storage is often higher, which leads many businesses to opt for object storage as it is a less expensive alternative. Additionally, many cloud providers offer intelligent tiering with pay-as-you-go models that can scale with need.

Security

Stored data must have the most up-to-date security to protect against external threats.

