

Technical Insight Report

Mainframe Shifts in Cloud Storage Usage

By Randy Kerns

October 2022



Evaluator Group

Enabling you to make the best technology decisions

Overview

Public cloud storage usage for mainframe environments is expanding for a variety of reasons, changing the operational procedures for Information Technology in those enterprises. Evaluator Group has completed a research study to provide insights into the current state and expectations as use of cloud storage continues to evolve.

The use of cloud storage is dominated by storing backup and archive data from the mainframe. The growth in the amount of capacity from mainframe to cloud is greater than the organic growth of data according to information from the respondents in the research study.

The research project was limited to enterprises with mainframe environments and did not include service providers or consulting companies. Initially, a set of focused questions to characterize the use and motivation for cloud storage was completed by individuals with responsibilities that would clearly validate their understanding. This information gathering was followed with individual interviews to get additional background, explanations, and an outlook on strategy. The interviews helped in understanding the thinking behind decisions and lessons learned. The combination of data and qualitative commentary presents a current picture of cloud storage usage and expectations in the near term (generally within a three-year horizon).

Cloud Storage Usage

The results of the survey for cloud storage usage are informative, but the narrative through the follow-up interviews provides context to what the charts show. For each of the selected areas reported on below, explanations and more in-depth analysis help in understanding.

Understanding what was the motivating factor for use of cloud storage in mainframe environments gives some insights into the dynamics of IT operations. Some decision-makers in IT are very strategic in their thinking while others may be overburdened to the point of reacting to directions from both internal and external sources (external meaning seeing what others were doing).

IT Considerations

Mainframe environments continue to grow in the amount of information. Decisions regarding storing of information consider the type of data, access requirements, maintaining current security level, and costs associated.

Addressing the storage needs with cloud storage is most likely for:

- Backup data used for disaster recovery
- Archive of data that is unlikely to be accessed frequently

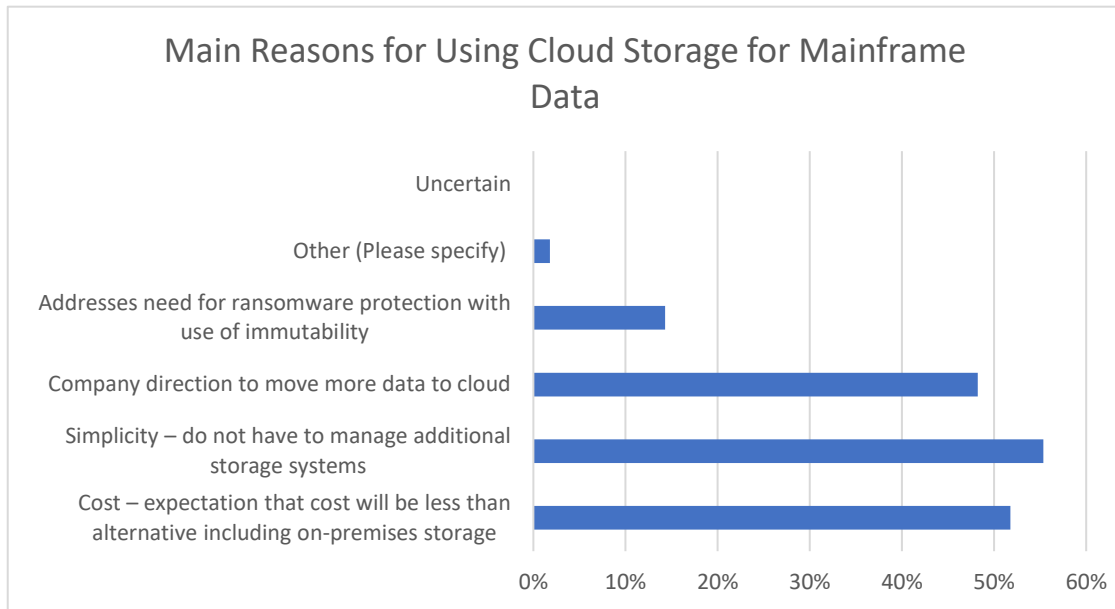


Figure 1: Reasons for use of Cloud Storage

Simplicity as the main reason for use of cloud storage was echoed through the interviews with the consensus that adding more storage systems introduced additional work, training, and effort for deployment required by the staff. For all interviewed, there was difficulty in addressing staffing requirements. Noted were budget and headcount restrictions and the lack of people with readily applicable skills to be productive immediately.

Many of the interviewees also said there was an expectation that overall costs would be less than on-premises storage but a general conclusion was that this turned out to not be the case unless detailed data management (and cost accounting) were done to ensure that data placement met cost expectations.

Company directives to utilize public cloud storage was another selection by a substantial percentage of respondents. Interviews revealed this was more of a general directive about use of the cloud as a company strategy with the expectation there would be cost savings and an increase in agility, specifically the ability to react more quickly to changes. There were several interesting observations here:

- Mainframe applications rushed to be run on the public cloud (either ported or an equivalent application used) suffered from the ability to guarantee performance and, despite expectations to the contrary, storage costs were prohibitive because of the class (expensive performance level) of storage required and the time-based charges for the amount that was resident. This led to added effort to move application execution back to the mainframe environment for more predictable operation.
- New or replaced applications were candidates to be deployed on public clouds.
- The security team took a very conservative view and were roadblocks in use of public clouds.

- Transitioning away from the mainframe was a major investment but a few were well down that path, reporting that applications were being rewritten (not ported) with new container-based architectures that they believed were more flexible and easier to maintain. The fact that mainframes support containers with Linux environments was not mentioned.




Figure 2: Cloud Storage for Backup or Archive Data

Backup and archive usage for mainframe data is well understood and a common factor among most respondents and interviewees. The delineation between what is backup and what is archive is not as clear since, in general, much of the data that is archived is done with backup software and the backup copy is retained. When interviewed, most respondents viewed their backup data on cloud storage as primarily for disaster recovery or recovery from a specific point while the most recent backup is still on-premises. Backups were based on using software to make a copy of volume data that would go to on-premises storage such as an IBM TS7700 and then potentially to cloud storage or directly to cloud using an S3 connector. Another case for backup was a snapshot on an IBM DS8900 that was tiered to cloud based on criteria with Transparent Cloud Tiering. Archive data was most commonly data that had a very low probability of subsequent access. The important consideration was that the same current operational practices for protecting data were used, even if the target was cloud storage.

Evaluator Group – Expectations of Change

The conversation and survey focus were on use of cloud storage in mainframe environments and the responses were clear about use for backup as primarily a disaster recovery (DR) copy. There will be value in the use of granular restores such as at the dataset level rather than the complete recovery expected in

a disaster situation. The conversations and questions did not delve into this usage but this will be of value, especially in selective restores for recovery from cyber-attacks.




Figure 3: Amount of Archive and Backup Data in the Cloud

The amount of data that mainframe operational environments stored on clouds gives a perspective on significance. The data shows that more than half of respondents reported storing more than one half of a petabyte of data in clouds. That is a large amount of data for mainframe environments and the expense of storage and management, either on-premises or in a public cloud, would be a noticeable amount in an overall budget.

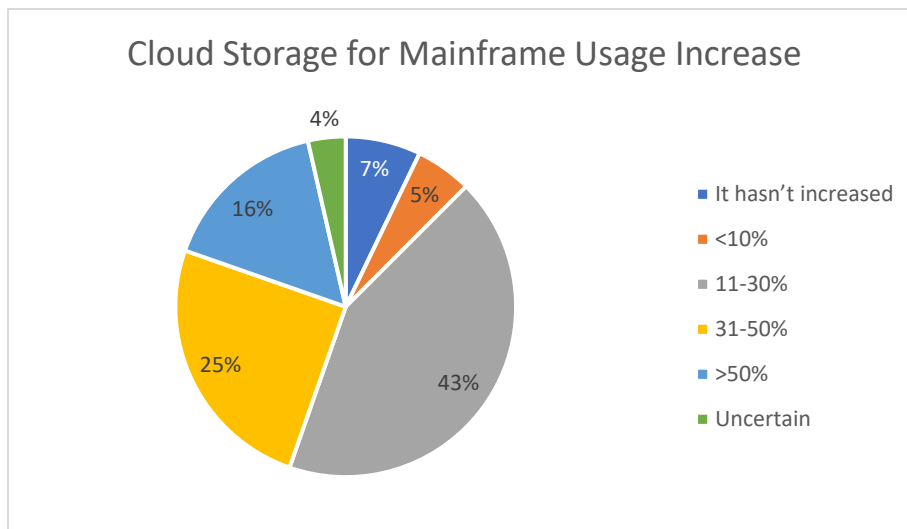


Figure 4: Growth of Cloud Storage in Mainframe Environment

A question that is invariably asked is “How much is the use of cloud storage increasing?” The answer is relatively clear from the survey data; in more than two thirds of the mainframe environments, it is more than 11%. The numbers tell a story of increasing amounts of mainframe data being stored in clouds. The narrative around this increase is aligned with the earlier motivation discussion: it is easier to acquire and utilize cloud storage than on-premises systems. Backup and archive were the primary types of data cited by every interviewee and the growth was primarily due to expanded usage.

Mainframe Storage in the Cloud and Effect on Tape Storage

The first guess by many is that mainframe data stored in a public cloud, specifically backup and archive data, was a trade-off from storing the data on tape that is on-premises. Understanding this dynamic is interesting and, from interviews, a very strategic decision.




Figure 5: Tape Usage

A large percentage of the respondents are using tape. With mainframe operations having evolved from reliance on tape, this was expected. Interviews indicated that even though some are not using tape, they are still using tape images stored on disk-based systems such as an IBM TS7700 system. The commentary about moving away from operational practices around tape or tape images was that it would be very difficult to change.

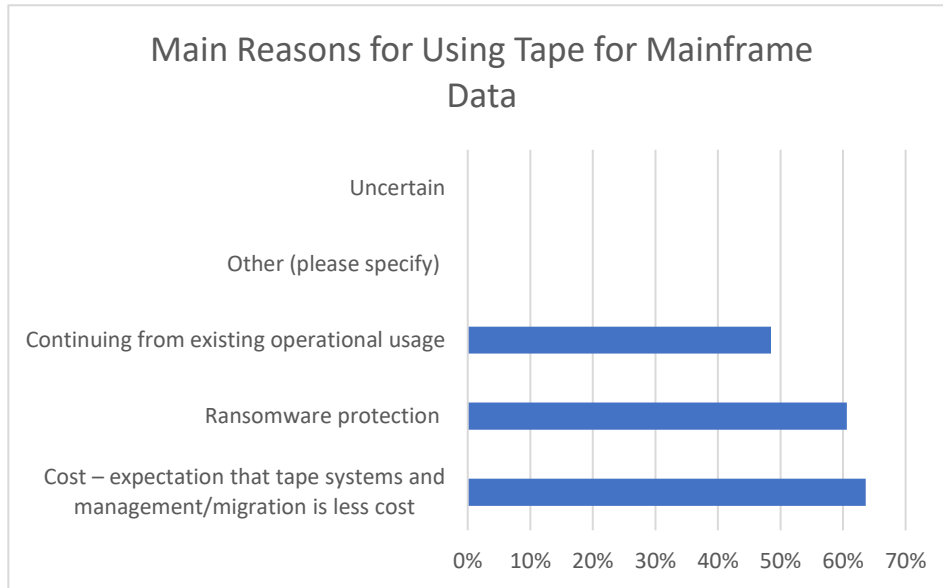


Figure 6: Reasons for Use of Tape

There are interesting contradictions in the reasons for using tape reported in the survey and the more expository answers from the interviews. The first was continuing current operations being brought up as a primary reason during interviews was not the most popular reason selected in the survey. This may have been borne out of frustration surfaced during conversations that operational procedures could not be changed without great difficulty. In the survey, the cost of storing data on tape media was also a commonly selected reason for using tape. When asked during the interviews about forward migration to new technology tape when earlier generations are no longer in new manufacture, the comments were that migration was something done on a regular basis but the cost of migration was not typically included with tape storage costs as reflected in the survey response.

The ransomware question in the survey and the follow-on interviews indicated that the ability to have a physical air gap, while not specifically planned for that purpose, was a defense given when challenged about ransomware recovery. A lengthier discussion from the interviews about security and ransomware is in a later section of this report.




Figure 7: Tape Usage for Backup and Archive

Comparing mainframe usage for backup and archive to tape usage in the survey does show a difference. When asked in the interviews, the response was that operational practices were to store data on tape when the processing was complete with the expectation it would not be needed again unless there were some unforeseen circumstances. The financial institution interviewees said that in general data was retained for seven years and then discarded (tapes destroyed).




Figure 8: Data Stored on Tape

Another comparative was how much data was stored on tape. It should be noted again that the tape usage respondents here all were also storing data on cloud storage. The results were similar with more than half storing greater than one half of a petabyte of data.

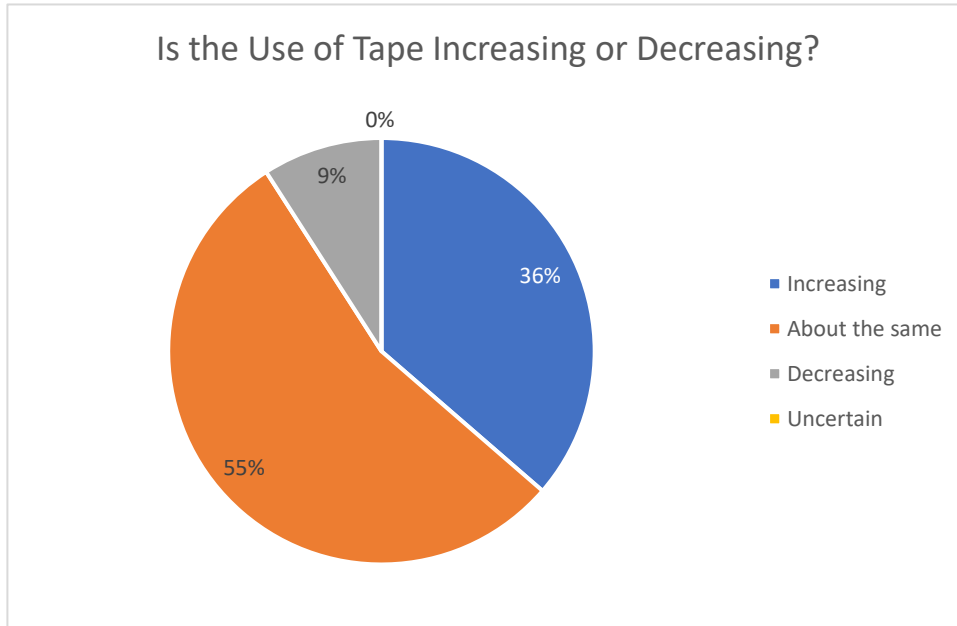


Figure 9: Tape Usage Increasing or Decreasing

For those using tape, about a third indicated the use was increasing while the majority said it was staying the same. The discussion in interviews that followed indicated the increase was not due to new applications or usages, just continued organic growth. The interviews indicated that flat or declining usage was due to expiring data and not storing data generated by new applications data on tape, instead moving it to cloud storage.

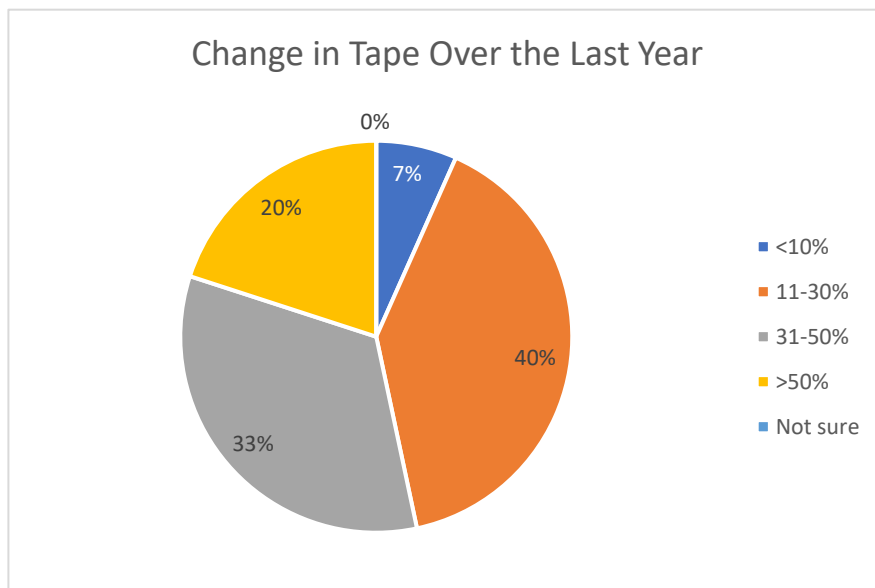


Figure 10: Tape Usage Increase

For those using tape and where there was an increase, the amount of increase was more than expected based on the interview responses. Based on the comments, the expected increase would have been less than 30%, focusing on organic growth.

These numbers should make it clear that in a mainframe environment the use of tape is still significant and growing, at least in the amount of data stored.




Figure 11: Next Generation Tape

Because of the potential cost of making a transition to a new generation of tape technology, understanding whether a transition was planned and, if so, when it would happen was an interesting topic. Discussing the ‘no transition plans’ in interviews, the assumption that no transition plans would equate to moving away from the use of tape was correct. For the rest of the current tape users, the commentary indicated that most had a regular cadence for making the transition, even if the total cost was not fully understood or incorporated into the cost of storage.

The use of cloud storage for mainframe data was seen as a more modern approach for backup and archive, much simpler to acquire than virtual tape or physical tape with required technology transitions, and would help in executive directives about use of public cloud. However, use of existing tape was so ingrained in the operations that making a change was a major effort unlikely to have budget assigned.

Operational Considerations

Discussions regarding use of cloud storage and tape storage included the topic of operational changes required. Changes were extra effort and added potential for errors. Those enterprises employing cloud storage for mainframe data that had made changes reported that the simplest means was to use

software that masked differences in virtual tape environments. The flexibility of being able to use either target was mentioned as highly valuable. In general, backup was the primary use case in the discussion.

Other Considerations

A concern mentioned by almost all the interviews was regarding security for mainframe data moved to cloud storage. A small number who did not use cloud storage, reported that their security teams would not allow usage.

Ransomware continues to be a major topic in the industry but most of those interviewed felt that the mainframe had greater overall security due to it being primarily a closed environment with continuing attention to security overall. The plan for recovering data from an attack was an interesting discussion. In general, the enterprises that used tape or virtual tape set into WORM mode (where recent copies were retained on-premises) believed they had covered the need for an operational air gap. Conversely, only a few of those interviewed were using immutable settings for cloud storage. This seemed to be a contradiction in the concern for security and the ransomware recovery strategy. Probing further, the belief was ransomware recovery would be done from on-premises copies. Also mentioned was a concern for the performance or time required to recover if data had to come from cloud storage. This was an important topic for most because recovery from the cloud using connections where they had nominal influence or control left them deciding to rely on recovering from on-premises copies.

Evaluator Group – Expectations of Change

From the information in the study, for most respondents the effective use of immutability is focused on the on-premises storage – setting WORM mode for physical and virtual tape. While that will not change, the use of immutability in cloud storage will become more common. The ability to control immutability such as S3 Object Lock through tools such as the Broadcom S3 connector make it such that it does not add to the complexity and operational workload.

Another element in security strategy was in regards to encryption of data. Universally, even if applications encrypted data, data at rest encryption in storage systems and cloud storage was used as well as data in flight encryption. There was no variation in the responses.

Evaluator Group – Expectations of Change

With the introduction of quantum-safe encryption, IT operations will be heavily influenced to utilize solutions that employ new technology. While not called out during the research, Evaluator Group believes this will be an important advancement in security that will eventually become a common practice.

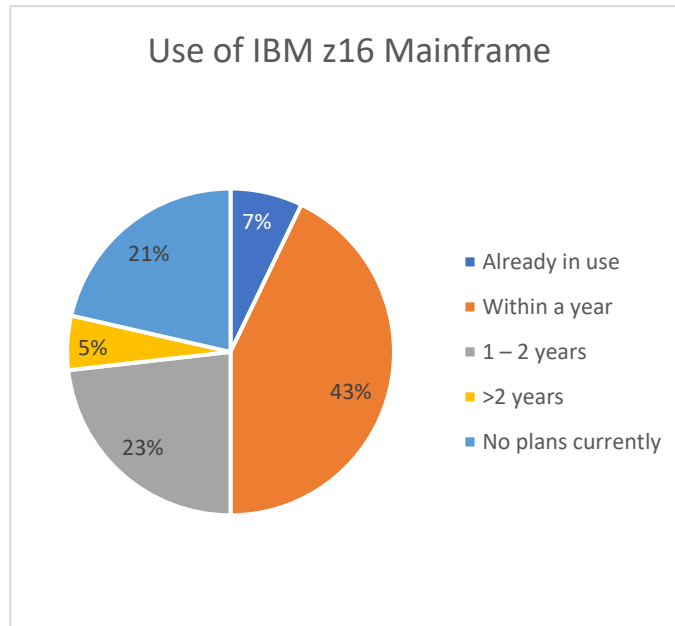


Figure 12: Use of IBM z16 Mainframe

A final area queried was in regards to transitioning to the new IBM z16 mainframes. Most enterprises had a plan for making the step to the newest model. When those with no plans were asked why, the response was that there was an executive direction to move away from the mainframe, mostly in executing on private clouds and in a few cases on public clouds. Because of that executive directive, planning for the next generation mainframe could not be done. It should be noted that those interviewed who had not already proceeded along the path to move away from mainframe usage, believed it would take a significant amount of time and required a budget that had not been fully understood.

Respondent Positions

For any survey and follow-on interviews, it is important to understand who was responding. The most useful information for the details presented would come from practitioners, those who work with the systems and software on a regular basis. Best described as 'This is their job', these people would have the knowledge to give accurate answers. Indeed, a few times during the interviews, if the person did not know, someone else was queried to give more precise answers.

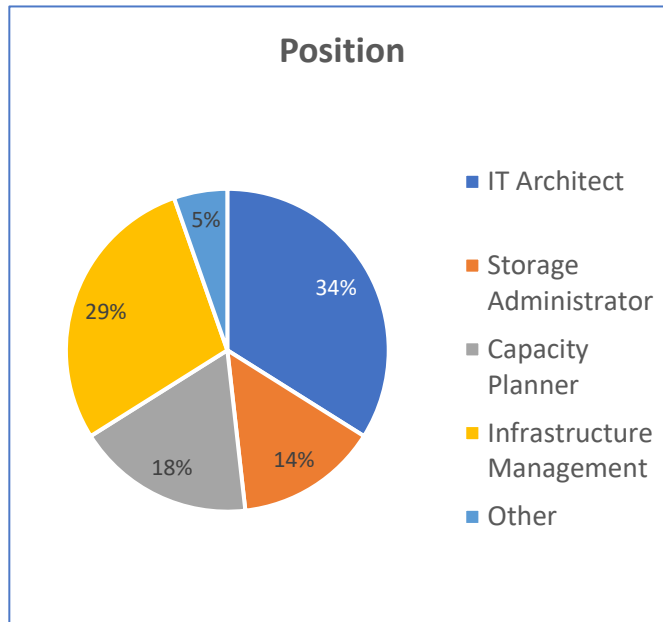


Figure 13: Respondent Positions in Mainframe Environment

The information in the charts and narrative was about companies that were using cloud storage for mainframe data. Those that did not were not included in the statistics, focusing the questions on more details around that usage and gathering information on how the use of physical tape was changing. The percentages using cloud storage for mainframe data vs. not is shown in the following chart.

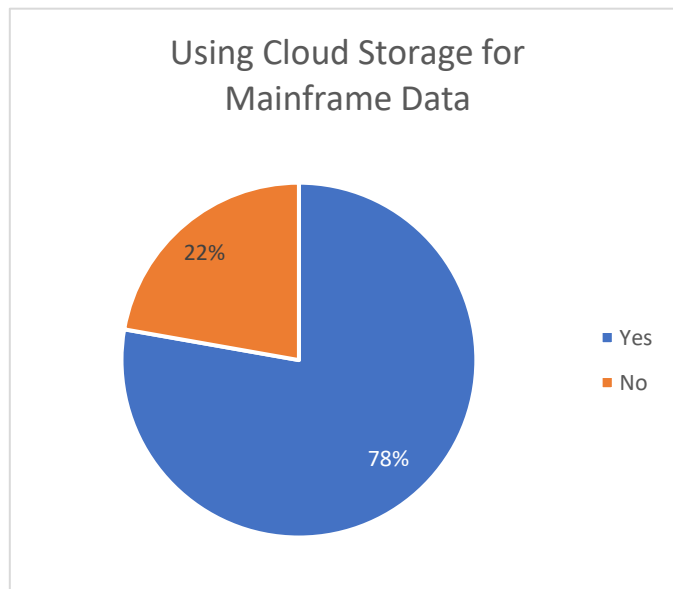


Figure 14: Using Cloud Storage for Mainframe Data

Summary

This Evaluator Group research project has increased the understanding of cloud storage usage for mainframe data and has yielded several insights. The insights, coupled with the interview discussions, lead us to making recommendations for IT in modernizing their operations for mainframe storage.

Insights from Research

- **Simplicity** – Primarily used because of simplicity in providing storage resources for backup and archive of relatively inactive data, the use of cloud storage is increasing, faster than the organic growth rate of mainframe data.
- **Cloud Storage vs. Tape Usage** – In some environments, cloud storage usage was at the expense of physical tape. But the current operational practices where tape is used will continue and, in fact, increase the amount of data on tape. While it might be assumed that the physical removability of tape may be an answer for air gap requirements for ransomware recovery, it was not a reason for any change in plans. It was convenient, however, to cite as one of the security attributes for IT organizations.
- **Operations** – The flexibility for IT operations to use either cloud storage or tape (physical or virtual) without operational changes was an important consideration. Significant change in operations were experienced by organizations that used cloud storage from applications running on zLinux.
- **Future use of Cloud Storage** – The research and interviews show that use of cloud storage will continue to increase, and that tape usage will continue in most cases. Evaluator Group expects that more applications will directly store data on public clouds or access data from public clouds. Using cloud storage as a means to interchange information will be more important for enterprises where public cloud usage is allowed.

Call to Action for IT

IT operations can benefit from the insights gained with initiatives in optimizing operations and addressing enterprise goals for modernization.

- **Economics** – The ability to expand storage capacity using cloud options as needed can reduce storage cost and the time required to deliver necessary resources compared to adding on-premises storage systems. Effectively delivering storage-as-a-service shifts the requirement for capital expense to operational expense.
- **Security** – Concerns regarding security of data continue to increase and storing of data whether it is on premises or in a public cloud will be subject to additional security measures. Encryption of data has been a requirement, but greater concerns are leading to the need for quantum-safe encryption algorithms to be used. File-level and whole-

volume encryption protect from exfiltration and theft of valuable information, common in ransomware attacks.

- **Operational Simplicity** – A software solution that enables use of public clouds for mainframe storage is simpler to introduce than a hardware-based solution. Software that can direct data access from current targets such as virtual tape to public cloud without operational changes eases the burden on IT operations and enables the economics and security requirements. Supporting both volume and file-level storage and access expands the value and adds to the operational simplicity.

A prime example of a solution to provide these high-value capabilities (economic, security, operational simplicity) is the Broadcom CA 1™ Flexible Storage™ solution. Many mainframe IT operations already have access to CA 1 Flexible Storage and can use it to modernize operations. Benefitting from the value provided should be a near-term initiative.

It is clear from the research that IT operations are challenged when changes are required because of staffing limitations. The calls to action noted are initiatives that will yield high value without encountering a high bar of implementation complexity and cost.

About Evaluator Group

Evaluator Group Inc., an Information management and data storage analyst firm, has been covering systems for over 20 years. Executives and IT Managers rely upon us to help make informed decisions to architect and purchase systems supporting their data management objectives. We surpass the current technology landscape by defining requirements and providing an in-depth knowledge of the products as well as the intricacies that dictate long-term successful strategies.

Copyright 2022 Evaluator Group, Inc. All rights reserved.

No part of this publication may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying and recording, or stored in a database or retrieval system for any purpose without the express written consent of Evaluator Group, Inc. The information contained in this document is subject to change without notice. Evaluator Group assumes no responsibility for errors or omissions. Evaluator Group makes no expressed or implied warranties in this document relating to the use or operation of the products described herein. In no event shall Evaluator Group be liable for any indirect, special, consequential, or incidental damages arising out of or associated with any aspect of this publication, even if advised of the possibility of such damages. The Evaluator Series is a trademark of Evaluator Group, Inc. All other trademarks are the property of their respective companies.