Cubbit

Storage buying guide for IT managers using Veeam

Stefano Baldi
VP R&D at Cubbit

Blaine Harnack
Head of Technical Services & User Data Management at Cubbit

Storage buying guide for IT managers using Veeam

Introduction

While talking with the IT managers of our client companies, two key points always emerge: the widespread use of Veeam as a backup solution and the emerging need to define the fundamental criteria when purchasing a storage solution.

Simultaneously, we are witnessing an exponential growth in unstructured data. In today's world, data is the new oil, with production showing no signs of slowing down. Unstructured data is expected to triple by 2026.

However, such rapid growth comes with risks, as evidenced by considerable incidents in recent years: the fire at OVHCloud that took millions of services offline (Reuters), a glitch in Amazon's cloud that disrupted access to hundreds of sites, including McDonald's and Delta Airlines (Washington Posts), and the global ransomware attack in early 2023 that made thousands of web pages unavailable (TechCrunch).

Unfortunately, security and disaster recovery are not the only issues to consider. The exponential growth of data poses compliance problems, risking to compromise the digital sovereignty of companies: non-compliance with GDPR can result in fines of up to 10 million euros and the direct closure of the company.

It is crucial, therefore, to identify the purpose of the storage solution being considered in the face of such risks. If local backup requires constant uptime and efficient restoration, a disaster recovery solution comes into play when business continuity is compromised. Like insurance, it must provide assistance when every other system fails. It must therefore not only comply with regulations, but also be hyper-resilient to any type of disaster, be it ransomware or a natural catastrophe.

This brings us to the challenge of choosing the most appropriate solution for specific needs. This guide will outline all aspects to consider, providing a comprehensive overview to facilitate an informed and efficient decision.

1. How to understand how much storage space is needed

Accurately determining the cloud storage space required involves carefully considering specific requirements as well as the nature and volume of data to be managed. It is not merely a quantitative matter but also includes objective variables, such as the importance of mission-critical data for the company's operational continuity, and subjective factors, like confidence in the reliability of internal business processes.

In light of these elements, it is advisable to proceed with an analysis of the redundancy systems to be implemented, the timeframe to be covered by the backup (weeks? months?), and the level of granularity to be achieved (should you aim for weekly or daily coverage?).

Let's consider an example where we need to keep a six-month history of a virtual machine. Note: The following numbers refer to backup sizes, considering compression and deduplication, which are generally much smaller than the ones required by the original VM.

In this case, for the first two weeks, we will perform:

1. A full backup (FB) every Monday morning (100 GB per week).
2. Incremental backups (IB) each evening (10 GB per week).

For the remaining two weeks, we will perform a full backup every Monday morning (100 GB per week).

For the first month, we will have:

100 GB (FB - 1st week) +
100 GB (FB - 2nd week) +
100 GB (FB - 3rd week) +
100 GB (FB - 4th week) +
10 GB (IB - 1st week) +
10 GB (IB - 2nd week) =

for a total of 420 GB

Therefore, our backup strategy for virtual machines (VMs) includes the following steps:

1. Weekly backups: Each week, we generate one Full Backup (FB) on Monday and five Incremental Backups (IBs) throughout the week, one per day. Each day, these backups are then copied to the cloud.
2. Two most recent weeks (granularity is daily): We store the complete set of backups (2 FB + 10 IBs) of the two most recent weeks. After this period, we delete the ten Incremental Backups but retain the two Full Backups. This way, for the two most recent weeks, we will always be able to restore the VM's situation on any given day and thus retain all relevant full and incremental backups.
3. Two-month retention (granularity is weekly): After the most recent two weeks, we keep all eight Full Backups, created each Monday, for up to two months. This allows us to restore the VM to any week within the two-month period using the Full Backups.
4. Two to six months (granularity is monthly): Beyond two months and up to six months, we only keep four specific Full Backups — those created on the last Monday of each month. This allows us to restore the VM to the state it was in at the beginning of each month within this timeframe.
5. Older than six months: We delete any backups that are older than six months, ensuring we don't use more space than necessary.

We will then need storage with a total space of:

420 GB (1st month) +
400 GB (2nd month) +
400 GB (1 full backup per month for the remaining 4 months) =

for a total of 1.2 TB

On the other hand, if we are particularly confident in our business processes — for example, if we are sure to quickly detect a data loss issue — or if we have non-critical data for business continuity, such as extremely detailed logs with summarized versions stored, we can consider reducing the size of the history.

We can also consider an extreme case, where we limit ourselves to maintaining two weeks of history for our virtual machine compared to the six months in the example above.

In this case, we would need a space equivalent to:

100 GB (FB - 1st week) +
100 GB (FB - 2nd week) +
10 GB (IB - 1st week) +
10 GB (IB - 2nd week) =

for a total of 220 GB

Lastly, it is essential to consider the cost. The price of object storage solutions designed to archive large amounts of data in a scalable manner is becoming increasingly affordable. Unlike local storage, whose price tends to increase over time.

If you need further information on the topic covered in this chapter or have doubts or questions regarding the example above, you can contact our team of experts to receive a custom free quote for the storage space you need.

2. Disaster recovery vs business continuity

When it comes to backups, performance is usually the first element considered. However, there are many factors to observe, especially concerning the storage system's role within the company's IT ecosystem. For example, speed undoubtedly represents a critical element if the focus is on a simple local backup. On the other hand, resilience and reliability carry more weight than performance for a second or third line of backup.

It is important to emphasize the need for cloud storage to offer a high degree of protection against threats such as localized disasters and ransomware while ensuring high reliability. Second or third-line backups function as an actual life insurance policy, a security measure that an increasing number of IT managers rely on.

In a catastrophic scenario that jeopardizes operational continuity, such as a fire in a data centre or the failure of multiple local backups, resilient storage capable of surviving regional-scale disasters can be the only adequate safeguard. The fact that a complete restoration from this type of storage can take several days is secondary. What is crucial is that the latest archive containing the data for business recovery remains intact, available, and protected.

Therefore, despite performance seeming like the most important criterion, data security is undoubtedly the top priority.

3. Upload performance

While, as we have seen, download restoration times are of secondary importance in the context of disaster recovery, upload performance is crucial. To keep a backup consistently updated, it is essential to be able to upload critical data within a reasonable timeframe.

This is why, for most companies, bandwidth is the primary point of concern. If the bandwidth is too low, the company will never be able to secure all its data, remaining constantly exposed to ransomware attacks and disasters.

For example, let’s assume an available bandwidth of 100 MB — a common scenario: many companies have headquarters in an area not covered by ultra-wideband. During a regular workday, the company will already saturate it by about 50%. Even if it chooses to perform nightly backups to utilize 100% of its bandwidth outside of working hours, it will still reduce the available bandwidth for other systems, such as security alarms, video surveillance, remote support, etc., limiting their functionality.

These issues can be circumvented with advanced backup strategies implemented by many vendors, including Veeam. Although still underutilized, these strategies are highly efficient, such as the Veeam Forever Forward Incremental Backup.

This strategy has three advantages:

1. Reduces file sizes through deduplication.
2. Updates only those fragments of the full backup that change over time.
3. Optimizes restoration times by downloading a ‘synthetic’ full backup (i.e., composed of original and updated fragments) that is then stored in the cloud as a single file.

In this way, when it is necessary to recover data in the face of a ransomware attack, restoring the synthetic backup is sufficient, eliminating the need to restore hundreds of incremental backups. This procedure can be completed within a minimal timeframe.

It is advisable to opt for storage that ensures high reliability and then, based on specific requirements, consider the opportunity to improve performance using tools provided by Veeam. A backup is only as reliable as the storage that houses it. The tools provided by Veeam make it possible to optimize times and objectives based on the specific network configuration being utilized.

If you need further clarification on the topic, you can schedule a consultation with our team of experts.

4. S3 compatibility

Talking with our customers, the importance of the S3 standard is evident. We are in the era of hybrid and multi-cloud: 64% of companies use more than one cloud provider. To ensure easy and efficient integration across different platforms, it is essential that the storage solution we choose adopts the S3 standard, thus avoiding manual operations that require a significant investment of time and attention and are prone to high error rate.

Thanks to S3 compatibility, developers and IT managers can distribute their storage space among multiple cloud services such as Amazon, Google, and Cubbit quickly and without having to rewrite the source code. Another benefit of S3 compatibility is the ability to integrate S3-compatible object storage with backup solutions popular in IT departments, such as Veeam, Synology, Commvault, and MSP360, with just a few clicks.

For all these reasons, S3 compatibility is an essential feature of any storage solution.

5. Storage cost

With the introduction of Veeam v12, it is now possible to perform backups directly to S3-compatible cloud capacity tiers without going through a local performance tier such as NAS or hard disks. This innovation, coupled with local storage having an increasingly higher cost in the market, allows for a reduction in CapEx and frees up local space by moving older backups directly to the cloud. This provides solidity and great flexibility in defining the most suitable storage strategy, whether hybrid, cloud-first, or anything else.

It is also essential to consider the lifecycle of the hardware used, whether you prefer a local or cloud solution. In this perspective, geo-distributed cloud object storage (see the demo with Veeam V12) ensures significant savings by recycling the existing hardware infrastructure, thereby extending its lifespan and simultaneously reducing CapEx and CO2 emissions — in addition to the already amortized cost of hardware that would otherwise be decommissioned.

6. Customer support and team expertise

The purchase of a storage solution should not be limited to evaluating the product alone. From various reports by IT professionals, it is increasingly clear how crucial it is to have a team with expertise in the field, capable of dedicating time to listening and responding quickly and in the customer's language.

It is increasingly common to find that large companies provide a standardized service and adopt an impersonal and distant approach towards their customers. The lack of human interaction is particularly felt during crucial moments — for example, during the solution installation or when emergency recovery needs to be activated.

It is no coincidence that the Managed Service Provider (MSP) approach is becoming more popular than the Value-Added Reseller approach. This is indeed a relationship of continuity between the provider and the customer, extending beyond a single transaction to include ongoing consultation, both technically and in advising on new products.

For this reason, it is crucial to evaluate the availability offered by the support team, the language of communication, and the knowledge and expertise that this team has developed and can provide to its customers.

7. Hyper-resiliency and security

We have mentioned the need for highly resilient and secure systems several times, but what factors should be considered?

A fundamental element is certainly resilience. In this context, the metrics to take into account include:

- Durability: a value expressing the probability that the data will not be lost. It must be at least 99.999999999% (the standard ‘11 9s’).
- Availability: a value expressing the probability that the data is immediately accessible when attempting to retrieve it. Availability must be at least 99.95%.

In addition to this first statistical element, there are some features that are important to consider.

For protection against ransomware:

- S3 Object Locking: a feature that allows you to make data immutable by setting an expiration date. Within this period, your data can’t be encrypted, altered, or deleted by anyone or anything, be it a ransomware attack or human error.
- S3 Versioning: a feature that allows you to maintain multiple versions of the same object over time. Each version is marked with a unique identifier, enabling easy recovery and restoration of specific versions of the stored data.

For protection against localized disasters, the most common solution is data redundancy. For on-premises storage services, this process must be done manually and involves additional costs, such as purchasing new hardware and possibly new physical space to accommodate it.

On the other hand, when it comes to cloud storage, there is a growing trend towards automatic data redundancy. Some services, having to cover redundancy costs between their physical data centres, offer it as a premium option, while others incorporate it by design into their service.

Unlike traditional cloud storage, Cubbit's technology, for example, does not rely on centralized data centres. Instead, files are encrypted, fragmented, and replicated across multiple geographic locations — safe from ransomware and localized disasters. Thanks to the built-in geo-redundancy in the service, maximum durability and resilience are guaranteed at a fraction of the cost of alternatives on the market.

If you want to learn more, start a free trial of Cubbit's geo-distributed cloud storage now.

8. Digital sovereignty and compliance

While high security and resilience are invaluable in the storage domain, digital sovereignty and compliance become essential in sectors such as Public Administration and Healthcare, where increasingly restrictive regulations prevent IT managers from implementing reliable disaster recovery and business continuity plans.

An example of this dynamic was provided by the IT manager of a genetics institute, one of our clients. They explained how challenging it is to find an off-site backup service for data protection from ransomware and disasters that is also compliant with existing regulations regarding DNA — which is why, before discovering Cubbit, the company had limited itself to on-premises storage solutions.

It is also essential to consider that, following a cyber attack, any company, from ICT to manufacturing, is subject to legal scrutiny. In such cases, if anything is yet to be done in terms of compliance and resilience, hefty fines are imposed, often leading to the company's closure.

For this reason, it is increasingly necessary to adopt solutions that comply with all data management regulations and standards, such as GDPR, ISO 27001, and ACN (formerly AgID). In short, it's important to adopt a solution that allows for data localization and geo-fencing within the boundaries required by the nation's law while ensuring total control of the data and minimized egress fees — so that data can be retrieved without issues.

Cubbit enables 200+ European companies to add a secure, GDPR, and ISO 27001 compliant cloud backup to their S3-compatible client, such as Nakivo, Veeam, or Synology.

Giuseppe Giovinazzi, IT manager of Trasporti Pubblici Parma, explained the benefits of adopting Cubbit in a case study.

Conclusion

Purchasing a storage solution for IT managers using Veeam is a crucial decision that requires careful consideration of various factors. In this guide, we have examined the key elements to evaluate during the purchasing process.

First and foremost, it is essential to understand your storage needs, considering the amount of data to be stored and the period you want to retain it. In this context, it is important to consider the granularity of backups and the system redundancy needed to ensure business continuity.

Furthermore, when purchasing a storage solution for backup, it is fundamental to realize that you are investing in a disaster recovery solution rather than a business continuity solution. System resilience and recovery capability are crucial to ensure data protection in case of disasters or ransomware attacks.

Upload performance is another critical aspect. The availability of adequate bandwidth and technologies like Veeam's Forever Forward Incremental Backup can help ensure reasonable data upload times.

Customer support and the expertise of the team are other valuable elements. Ongoing consultation and effective communication with the support team can make a difference in ensuring an optimal storage solution and an overall positive experience.

Finally, hyper-resilience and security are key requirements for a reliable storage solution. Storage durability and availability must be taken into account along with features like S3 Object Locking and S3 Versioning to protect data from threats such as ransomware.

In conclusion, choosing a highly resilient solution supported by an expert team and compliance with security requirements ensures critical data protection and operational continuity.

As a European company offering geo-distributed cloud object storage, we position ourselves as your partner in this delicate decision. We are ready to listen and delve into your use case: book a 15-minute call with our experts now.

About Cubbit

Cubbit is Europe’s first geo-distributed cloud storage: hyper-resilient, sovereign, 100% S3 compatible — with no extra costs for redundancy.

Unlike traditional cloud services, where data is stored in a few centralized data centres, data on Cubbit is encrypted, fragmented, and replicated across multiple locations within a single country or geographic perimeter — safe from ransomware and localized disasters.

The service is compliant with GDPR and AgID, and certified with ISO 9001 and ISO 27001. It supports object lock, versioning and IAM policy.

Thanks to S3 compatibility, you can integrate Cubbit in just a few clicks with any S3-compatible stack (e.g., Veeam, Nakivo, Synology,) enabling **immutable backup** lines with maximum security.

Today, Cubbit technology is trusted by 5,000+ customers, partners, and companies in 70+ countries, including globally renowned partners, institutions, and enterprises like Exclusive Networks (a $4.9B global distributor in cybersecurity, listed at Paris Stock Exchange) and Leonardo (a $14B world leader in defence and cybersecurity).

If you want to learn more and discover how Cubbit can help protect your company’s data and backups, you can activate a cardless free trial. And if you need custom support or consultation on storage and cybersecurity issues, you can book a free 15-minute consultation.

About the authors of this guide

Stefano Baldi is Vice President of R&D at Cubbit. An engineer by training, passionate about software and AI, in the second half of the 2000s he started working on Big Data, and high-reliability storage systems for major Italian telcos. Ten years later he founded and led as CTO two different startups, producing patented hardware and software solutions for cloud object storage, both on cloud and on-premise — solutions which were later adopted in several industries, including space, defence, and healthcare. Today, Stefano is responsible for innovation of products, methods, workflows and know-how in cloud storage.

Blaine Harnack is Head of Technical Services & User Data Management at Cubbit. A Canadian native, Blaine has more than 10 years of experience in IT, communication, and research. Since 2019, he has been responsible for Cubbit's technical services team and software quality assurance processes — a team that works closely with customers and users to ensure the best experience when using geo-distributed cloud storage.

Cubbit

Via della Zecca 1 - 40121, Bologna (BO), Italy

business@cubbit.io
www.cubbit.io