



Vawlt “Supercloud Storage” Report

A Broadband-Testing Report

First published October 2024 (V1.0)

Published by Broadband-Testing

E-mail : info@broadband-testing.co.uk

Internet: [HTTP://www.broadband-testing.co.uk](http://www.broadband-testing.co.uk)

@2024 Broadband-Testing

All rights reserved. No part of this publication may be reproduced, photocopied, stored on a retrieval system, or transmitted without the express written consent of the authors.

Please note that access to or use of this Report is conditioned on the following:

1. The information in this Report is subject to change by Broadband-Testing without notice.
2. The information in this Report, at publication date, is believed by Broadband-Testing to be accurate and reliable, but is not guaranteed. All use of and reliance on this Report are at your sole risk. Broadband-Testing is not liable or responsible for any damages, losses or expenses arising from any error or omission in this Report.
3. *NO WARRANTIES, EXPRESS OR IMPLIED ARE GIVEN BY Broadband-Testing. ALL IMPLIED WARRANTIES, INCLUDING IMPLIED WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NON-INFRINGEMENT ARE DISCLAIMED AND EXCLUDED BY Broadband-Testing. IN NO EVENT SHALL Broadband-Testing BE LIABLE FOR ANY CONSEQUENTIAL, INCIDENTAL OR INDIRECT DAMAGES, OR FOR ANY LOSS OF PROFIT, REVENUE, DATA, COMPUTER PROGRAMS, OR OTHER ASSETS, EVEN IF ADVISED OF THE POSSIBILITY THEREOF.*
4. This Report does not constitute an endorsement, recommendation or guarantee of any of the products (hardware or software) tested or the hardware and software used in testing the products. The testing does not guarantee that there are no errors or defects in the products, or that the products will meet your expectations, requirements, needs or specifications, or that they will operate without interruption.
5. This Report does not imply any endorsement, sponsorship, affiliation or verification by or with any companies mentioned in this report.
6. All trademarks, service marks, and trade names used in this Report are the trademarks, service marks, and trade names of their respective owners, and no endorsement of, sponsorship of, affiliation with, or involvement in, any of the testing, this Report or Broadband-Testing is implied, nor should it be inferred.

TABLE OF CONTENTS

.....	i
TABLE OF CONTENTS	1
BROADBAND-TESTING.....	2
EXECUTIVE SUMMARY	3
WHY DOES THE WORLD NEED VAWLT?.....	4
VAWLT OVERVIEW.....	5
VAWLT: HANDS ON	7
Installation, Deployment and Management.....	7
Volume Creation	10
Vawlt Storage Agent.....	11
Trash Control And Snapshot Rollback	12
Immutability.....	14
Use Cases	16
Integration Use Case.....	17
IN CONCLUSION	18
Figure 1 – The Vawlt Architecture	6
Figure 2 – The Vawlt Management Interface: Partner vs. Customer	7
Figure 3 – Creating a Template – CSPs.....	8
Figure 4 – Configuring Trash Policy And Immutability	9
Figure 5 – Creating Volumes.....	10
Figure 6 – Order Summary	10
Figure 7 – Storage Agent Overview Screen	11
Figure 8 – File Transfer Throughput	11
Figure 9 – File Restore Example	13
Figure 10 – Snapshot Rollback Example.....	14
Figure 11 – Immutability Examples	15
Figure 12 – Immutability No File Deletion Example.....	16
Figure 13 – Veeam Integration.....	17

BROADBAND-TESTING

Broadband-Testing is an independent testing operation, based in Europe. Broadband-Testing interacts directly with the vendor, media, analyst, consultancy and investment communities equally and is therefore in a unique space within IT.

Testing covers all aspects of a product/service from business rationalisation in the first instance to every element – from speed of deployment and ease of use/management, through to performance and accuracy.

Testing itself takes many forms, from providing due diligence for potential investors through to public domain test reports.

Broadband-Testing is completely vendor neutral and independent. If a product does what it says on the tin, then we say so. If it doesn't, we don't tell the world it does what it cannot do... The testing is wholly complementary to analyst-related reports; think of it as analysts getting their hands dirty by actually testing what's on the latest hype curve to make sure it delivers on its claims and potential.

Broadband-Testing operates an **Approvals** scheme which prioritises products to be short-listed for purchase by end-users, based on their successful approval, and thereby short-cutting the evaluation process.

Output from the testing, including detailed research reports, articles and white papers on the latest IT technologies, are made available free of charge on our web site at [HTTP://www.broadband-testing.co.uk](http://www.broadband-testing.co.uk)



EXECUTIVE SUMMARY

- The cloud era is well and truly established and here to stay, but there are significant issues with a cloud-based strategy – or even a hybrid approach – that need resolving.
- Unanticipated high (and often unmanageable) costs of a single or multi-cloud strategy are now well documented, but there are several further, and equally important issues: security, data governance, performance, ease of management, data repatriation... It's a long list of negatives.
- Given the amount of data even a small to medium company is generating and having to manage nowadays, it's not as simple as "taking everything back On Premise" (OnPrem). Gone are the days when even a mid-sized company had its own "machine room" onsite, with mainframe and network servers and storage providing the entire data repository for the company. Yes, OnPrem data storage very much has a role to play, but it is not a solution in its own right.
- Taking a DIY approach to creating and managing a hybrid/multi-cloud environment is an enormous challenge; too many different options to manage, too many variables – and how do you secure that environment?
- With its "supercloud" or "cloud of clouds" approach to the issue, Vawlt is offering a "one-stop shop" approach to creating that ideal hybrid/multi-cloud data environment.
- Moreover, it is looking to do that in an incredibly secure fashion, with the added benefits of reducing costs, easing the management overhead (enormously) and improving performance. In other words, more for less – can it be possible?
- We carried out an extensive, multi-site test and found that all the Vawlt claims were completely substantiated. We also liked the way the Vawlt platform is designed to accommodate both the service providers and their customers, or indeed an enterprise customer running its own internal data management. In other words, it is a true single solution, regardless of whether you are selling or receiving those cloud services. One obvious, additional benefit of this approach, is that any updates are only impacting upon a single platform, not multiple vendor products and services.
- From a pure security perspective, the Vawlt platform goes a long way towards making your multi-cloud/hybrid infrastructure effectively ransomware immune.
- Overall, while there is clearly an extensive roadmap in place for further developing the Vawlt platform, it appears to us that the company is well ahead of the competition right now.

WHY DOES THE WORLD NEED VAWLT?

Whatever our thoughts on cloud computing – and let's face it, it's just outsourcing data centre capacity, so nothing new – it is here to stay; at least for the foreseeable future as well as that can be depicted in the IT world.

There are two ways of substantiating this claim. The first is to look at what the analysts are saying. Gartner is forecasting worldwide public cloud spending in excess of \$675 billion this year and that's just the end user companies. It anticipates future annual growth around the 22% mark. Last year, IDC forecast that the public cloud market would reach \$1.35 trillion by 2027, with an average annual growth rate of 27%. You can choose to believe what you want, but clearly companies are throwing money at the cloud.

The second way of supporting the continuing growth in cloud theory is to use simple common sense. Latest indicators show that data is still growing exponentially, with over 400 million terabytes of data being created daily. This, in turn, equates to around 147 zettabytes per year. When I started in IT, I'm pretty sure the word "zettabytes" didn't exist! The point is, that's a lot of data to manage locally, so I don't think cloud computing is disappearing anytime soon.

However... those very same analysts citing endless market growth are equally observing a kickback against public cloud and the number one objection is cost. Several reports have forecast a figure in excess of 80% of companies citing high – or even uncontrollable - cost as their primary cloud challenge and that optimising cloud spend can become a full-time job in its own right. But it's not just costs that are causing problems for the ITOps guys. Many of the promises of a better world cited by the major public cloud providers initially, are proving to be inaccurate. These include:

- Cost-effective
- Limitless and instant scalability
- Easy data management
- Totally secure and resilient environment

The reality, however, has been somewhat different. Costs are not only exceeding expectations, but proving impossible to predict and manage. If you look at the menu of options and related costs for a service such as AWS' S3, for example, you will see why. All the "simplicity" aspects of cloud computing – given the huge range of options, the management complexity of supporting a multi-cloud strategy and the impact this has on "easy scalability" – are proving to be a logistical nightmare.

Security at the cloud itself may be at "Fort Knox" levels, but that data has to be accessed remotely and therein lies the opportunity for a cyber-attack. Going for a single CSP/single ISP connectivity strategy immediately creates single point of failure potential. Moving to a multi-cloud strategy opens up more potential security issues.

There's also the compliance aspect that no CSP talked about initially. Data governance and data provenance have become major potential obstacles: how can you be sure your data is compliant, when you don't even know where it actually is residing at any given point?

But managing a hybrid environment, combining cloud with OnPrem, in order to potentially safeguard data that might be at risk from compliance standards by keeping it in-house, creates an even greater security – and data management – challenge. And, here's the final "elephant in the room": data repatriation. Who needs ransomware when you look at the typical costs of repatriating your data from a CSP? And then there's the complexity...

If you think about the original mainframe concept, and why it was so successful, the primary ingredient was "simplicity". One data storage silo to manage and secure, one management interface and the ability to control and secure all user access to that data – and applications – at all times, from that same interface. So, imagine a scenario where that concept is applied to the world of a multi-cloud/hybrid data strategy. That is what Vawlt is looking to achieve with its "supercloud" or "cloud of clouds" platform. And that's my observation, not a line from Vawlt's marketing data sheet.

Let us then look at what the Vawlt solution actually consists of, followed by a hands-on session where we put the product to the test.

VAWLT OVERVIEW

Vawlt is essentially a platform that forms the central hub for creating and managing data Volumes, users, access control, and deployment.

Every data Volume can be backed by a set of different cloud storage providers, chosen by the user, which effectively creates a cloud of clouds, or supercloud environment. This applies equally to private storage environments, regardless of storage type or vendor, with the same feature set and service guarantees. The platform automates all interaction, removing the need to administer multiple interfaces and systems. Instead, you simply interact with one – Vawlt.

Moreover, the platform has been designed – with control plane and data plane – to support both service providers/partners and their customers/enterprise in that single platform and interface. The Vawlt Platform is where all the configuration is carried out: for example, volume creation, access control, etc), while the data plane is the basis of the Software Agent. Each party therefore has access to what is relevant to them, simplifying using Vawlt even further.

A fundamental part of the Vawlt architecture is the software agent, that creates the entry points for the data-plane. This manages the interaction between Volume configurations and the storage interfaces themselves – such as an S3-Compatible API, or a File System. Rather than directly uploading any written data to the Volume, the Vawlt agent stores it in a local cache on the disk. At this point, before being asynchronously uploaded to the cloud providers, the agent performs a series of security operations on that data.

- It is encrypted using Vawlt's proprietary "zero-knowledge end-to-end encryption protocol". This is designed so only the owner of that data has access to it.
- The data is split into redundant segments using erasure-code techniques to maximise storage efficiency.
- Each segment is then uploaded to different cloud providers using Byzantine-quorum systems, so as to maximise fault tolerance.

In order to reconstruct the original data, only a subset of the segment is sufficient, meaning that data is accessible, even if some of the CSPs are offline and unavailable, have suffered data loss or corruption, or have even become the victim of a cyberattack themselves. Vawlt doesn't just apply Byzantine-quorum systems for storing data across multiple CSPs, it also uses them for storing metadata and managing concurrent access to the same Vawlt Volume. In other words, Vawlt provides a globally distributed namespace, fully supported by Byzantine quorum systems, with no reliance on a single entity and no single point of failure or weakness, therefore.

The architecture is designed so that data is never uploaded through Vawlt servers – indeed the data never touches the Vawlt servers, as we'll see later – removing yet another possible point of failure or delay. In practise, it means that scalability, performance, availability and security are all optimised, resolving many of the issues associated with using cloud-based storage.

Given the enormity of the ransomware issue right now, the Vawlt agent extends the security to provide specific ransomware mitigation features, such as Immutability and Versioning, safeguarding data integrity and providing protection against unauthorised modifications. Flexibility is provided in the form of configurable trash policies, so deleted data can be retained for a specified period before permanent deletion, providing an added additional protection against accidental or malicious data loss.

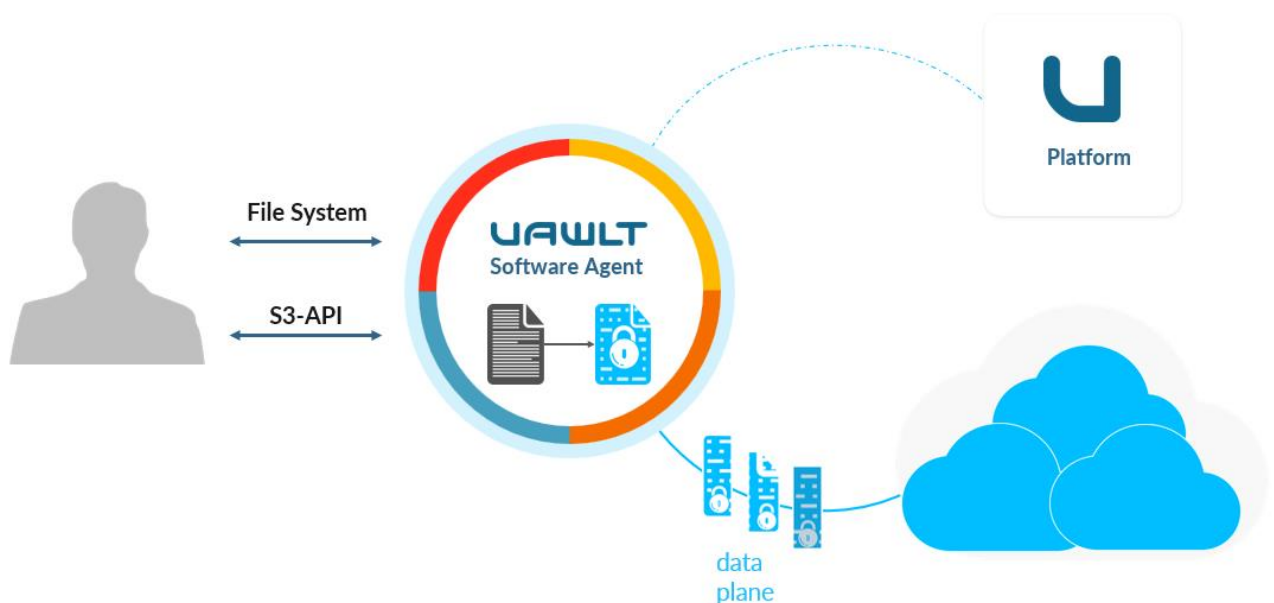


Figure 1 – The Vawlt Architecture

As would be expected, in addition to its multi-cloud, distributed data and flexible backup capabilities, Vawlt supports all the classic data management features, such as data synchronisation, snapshot and Rollback, and ensures data sovereignty for compliance regulations requirements.

Vawlt Versus A DIY Approach

What Vawlt is offering is realistically only otherwise achievable by taking a Do-It-Yourself approach to creating a multi-cloud/hybrid storage strategy. But taking a DIY approach is an enormous challenge; you have to research pricing and management options for each CSP, then manage each on in an ongoing basis, including data transfer between those different CSPs, ensure pricing stays in line with expectations - good luck with that one! And then you still have to secure that environment – good luck with that one too!

With its “supercloud” approach to the problem, Vawlt is offering a “one-stop shop” approach to creating that ideal hybrid/multi-cloud data environment, including letting you know exactly how much each set of options you select, in terms of mix of CSPs, geo-locations, data quotient and every other relevant variable, will cost. There are no hidden costs whatsoever. And just one management interface. Think about it...

VAWLT: HANDS ON

Installation, Deployment and Management

As we noted earlier, one of the principal benefits of the Vawlt platform, is that has been designed for both the partner/service provider and their customers to use, sharing the same management interface but with access to different options. For example, the sidebar on the UI provides more – and different – options for the partner (left-hand graphic) to those of the customer (right-hand graphic).

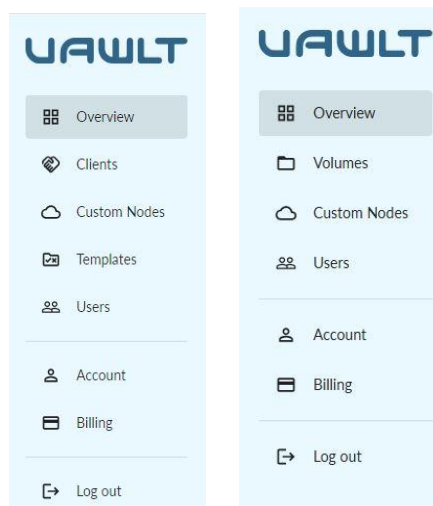


Figure 2 – The Vawlt Management Interface: Partner vs. Customer

From an installation perspective, it really is as simple as logging into an account set up by Vawlt, with the only download requirement being the Software Agent and any required documentation – both available from the UI.

Looking at the Partners perspective, deployment is also very straightforward, even if their customer is a large enterprise, for example. One reason for this is that Vawlt provides a template option for constructing a series of “ready to go” templates that can be designed exclusively for a one-off customer, or so that they can readily be reused – with minimal editing – again and again for a variety of different customers (following the 90% of options works for 90% of customers rule). Basically, as a partner, you first create a client profile – basic administrative details, including billing information such as currency/exchange rate, percentage mark-up (not for customer eyes!) and frequency of billing. There is also the option to provide a free, limited time trial.

A custom node option allows the creation of your own storage locations, in addition to those CSPs supported by Vawlt by default. For partners - for example MSPs - that already have a private cloud solution in place, and consequently are effectively competing with Public Clouds, the custom node option provides a powerful – and “politically correct” - extension of the Vawlt solution.

It allows them to add their own cloud as a private cloud and create a template combining them with public clouds in a way that looks more like they are collaborating than competing; keeping all partners happy! Here you will also create the set of users, company accounts and general billing/payment/invoice information. Creating templates is where you combine the key features of Vawlt: the base template – Hot, Warm, Custom or Archival - minimum storage duration (in days) and Time To First Byte (TTFB), the geographical dispersion of the data, then which CSPs/locations are available/going to be selected for primary through to standby (four in total).

The screenshot shows the 'Creating a Template' interface with the following sections:

- GENERAL INFORMATION**
 - Template Name: GDPR European Compliant Premium
 - Base Template: Custom
 - Template Color: #21BFFC
- PROVIDERS CONFIGURATION**
 - Minimum Storage Duration: 30 days
 - Time To First Byte (TTFB): Up to Hours
 - Geographical Dispersion: Worldwide
 - Providers:
 - AWS S3 (Any)
 - Azure (Any)
 - Google (Any)
 - IBM IBM (Any)
- INTERFACES**
 - File System:
 - S3:
 - Unlocked:

The 4th provider will only be used if one of the three main providers is faulty at a given time (e.g. downtime, data corruption).

Figure 3 – Creating a Template – CSPs

- Any
- AWS S3**
- Google
- Azure
- IBM
- OVH
- Oracle
- Backblaze

So, for example, in a worldwide deployment, you might choose AWS S3 in the USA, Google in Germany, Azure in Asia and Backblaze in the Netherlands. With a European geo-dispersion, for example, your location options will then be restricted to Europe – totally logical. It is worth reiterating that templates define the rules that later can be optionally applied on the volumes. As noted above, a partner can add their own private cloud, after which they can create a template which obligates its use, after which they can freely choose the remaining three clouds, thereby enabling them to maximise their business opportunities, while still being very flexible. With the data repositories chosen, you then configure some of the critical – and key – options of Vawlt; Trash Policy, Versioning (in AWS S3 it is a means of keeping multiple variants of an object in the same bucket. You can use the S3 Versioning feature to preserve, retrieve, and restore every version of every object stored in your buckets), and Immutability controls.

While the option is available to disable the Immutability function, realistically – as a primary feature of Vawlt – this is going to be enabled and the choices thereafter depend entirely on what level of data protection/control you wish to exert. Many of these features have a direct impact on other Vawlt features; for example, the Trash Policy defines Volume Rollback capabilities, based on the number of days trash is maintained for, before permanent deletion. Another important partner feature of Vawlt is the capacity for them to define their own alerts so they can be instantly notified of any important aspect; for example, when a volume is almost full and they need to urgently contact the customer to make an upsell and create more storage. Most important of all are the Immutability options for the objects inside the Volumes. When enabled, you can create default settings in the template, but these can also be edited – per Volume – on the Vawlt Software Agent. Immutability options include – for file and folder - object deletion, rename permissions, modify permissions, including append-only and overwrite-enabled options, and how those Immutability rules should be applied. The 'by time' policy means that the properties above will be applied after the time configured passes after the last modification on the object (you can specify a number of days or immediately after the object is created on the Volume) while by manual action means that you need to activate the Immutability configuration manually in the Software Agent for the objects you want to make Immutable.




Figure 4 – Configuring Trash Policy And Immutability

Volume Creation

The next critical factor is creating customer Volumes.

The screenshot shows a configuration form for creating a volume. It is divided into three main sections: IDENTIFY VOLUME, CAPACITY, and NODES CONFIGURATION. In the IDENTIFY VOLUME section, the Name is 'Report Test' and the Type is 'GDPR European Compliant Low Cost'. The CAPACITY section shows a Storage capacity of '1 TB' and a Download Quota of '130 GB FREE'. The NODES CONFIGURATION section lists four providers: Oracle Frankfurt, Germany; OVH Gravelines, France; Backblaze Amsterdam, Netherlands; and Wasabi London, United Kingdom. A note at the bottom states: 'The 4th provider will only be used if one of the three main providers is faulty at a given time (e.g. downtime, data corruption)'.

Figure 5 – Creating Volumes

Here is where your starting point can be one of the previously created templates, which you can now edit, as well as adding information such as the required storage capability (for that Volume) and the download quota. As you make your selections, a window on the right of the screen shows you the breakdown and absolute costs of that chosen Volume configuration.

The screenshot shows an 'Order Summary' window. It lists the following details for the volume 'Report Test': Storage capacity: 1 TB; Interfaces: File System; Geographical Dispersion: Oracle Frankfurt, Germany; OVH Gravelines, France; Backblaze Amsterdam, Netherlands; Wasabi London, United Kingdom; Grant Access: tedderdoguk@gmail.com; Trash Policies: 10; No. of file versions: 2; Metadata: none; Total to pay: €18.40 (+ VAT).

Figure 6 – Order Summary

Thereafter you can edit the Trash Policy, Versioning and Immutability for that specific Volume, specify which users can access that Volume, and add custom metadata tags, should you wish to do so. Once defined, the Volume will be created as a background task (typically took only a few seconds during testing).

Note: We also put to the test creating a volume combining the public cloud with an OnPrem node; in this case a private OnPrem S3 node using MinIO (an open-source, high-performance object storage server compatible with S3 and designed for large-scale workloads such as AI/ML, data lakes, and databases). We then simulated the accidental deletion (data corruption would have been a similar use case) of data by an end user on that private node and were able to prove that the data was still completely intact and healthy. Human error is still the major contributor to data loss, so being able to protect those users from errors is a huge benefit.

Vault Storage Agent

The final piece of the Vawlt interface jigsaw is the Storage Agent, which can be downloaded from the main UI.

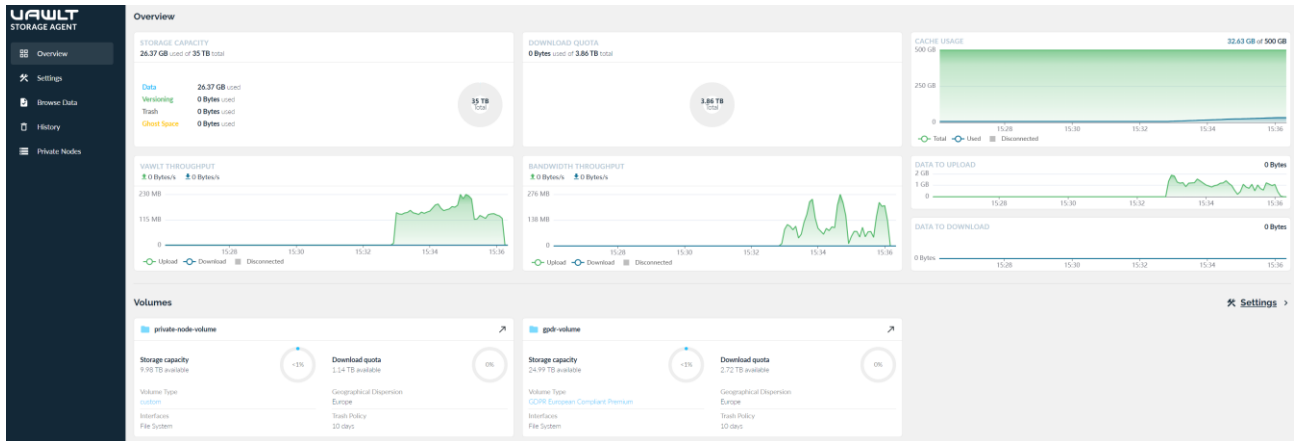


Figure 7 – Storage Agent Overview Screen

The main dashboard provides real-time information such as storage capacity usage, download quota usage, cache usage, Vawlt throughput and bandwidth throughput (the former shows the actual performance/data pattern generated by the Vawlt Agent and is the most relevant), data to upload/download activity and Volumes information. In the Settings section on the left, users can adjust configurations related to data volume accessibility, storage interfaces, and cache management. Regarding the storage interfaces settings, the user can, for example, turn on/off the virtual file system and the S3 endpoint, as well as edit settings as the Vawlt mount-point location for filesystem and port and SSL for S3. This provides a high degree of user customisation.

New cache locations can be created and managed (cleared) and cache size set. The default is to use all available space but you can use a minimum size of 1GB. Again, it is important to note that the cache size neither influences the amount of storage you can store in a volume, nor the maximum file size. As an example, you can store a 10TB file using a cache of 1GB. As with any data transfer, caching plays an important part in optimising performance on the Vawlt platform. Throughout the testing we were impressed by the performance capabilities of Vawlt.




Figure 8 – File Transfer Throughput

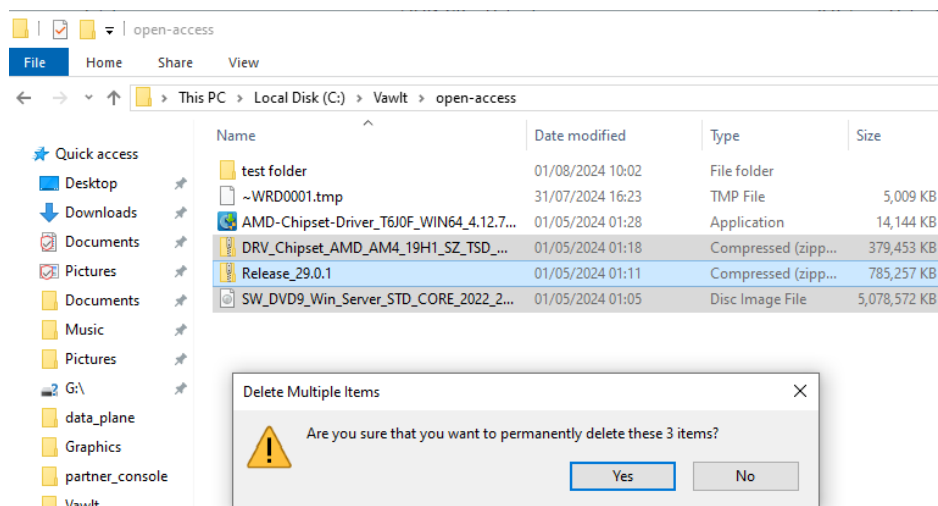
This was regardless of the actual data activity in question – and what the actual source and destinations were, both literally and geographically. Caching is put to great use in the Vawlt platform. For example – and this is using basic Internet provider capabilities – when transferring around 5GB of data between folders on a Volume, we saw throughput figures of around 3Gbps = impressive, even by caching standards! We were equally impressed with the throughput Vawlt can achieve when uploading or downloading data across the various CSPs that make up the Supercloud. With sufficient computational resources, Vawlt can easily reach multi-gigabit speeds for large-scale data transfers, as proven in our tests.

In the rest of this hands-on testing report, we will describe two of the main features of Vawlt: **Trash Control and Snapshot Rollback**, and **Immutability**. Additionally, Vawlt comes equipped with a large set of features that makes data storage and management much more efficient such as bandwidth throttling, volume automatic air-gapping, access permissions per prefix, cross-platform support, file versioning, etc.

Note: A data synchronization tool is also available that can be used to copy data from a source location to a volume. This tool can be configured for a single bulk copy or periodic jobs that synchronize changes performed in the source location.

Trash Control And Snapshot Rollback

In addition, to providing real-time activity information and configuration editing, here is where you access Volume specific functions, such as trash control and Rollback capabilities via the History tab on the sidebar menu. Throughout the testing we created Volumes with different trash policies and put the trash control (restore) and snapshot Rollback capabilities to the test, with no failures. As a simple example, here we created a Volume “open-access”, added files, deleted three, then restored one back. The actual requirement, from a user standpoint, could not be simpler.



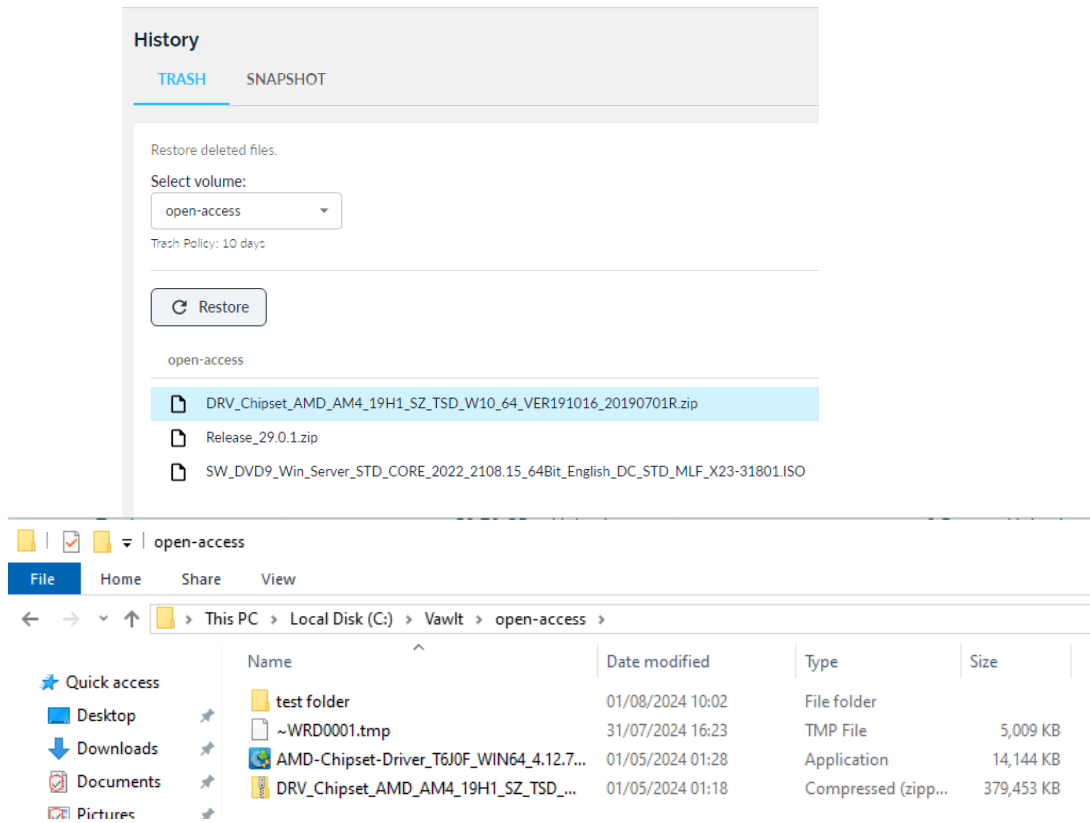



Figure 9 – File Restore Example

Having the ability - for each Volume - to define a custom retention policy for your files, meaning you can define a period when your files are still available for recovery, even if they're deleted, makes for very flexible and secure data management, even protecting the users from their own errors!

Directly tied in with this data history control is the snapshot Rollback capability. One more, over the test period we performed many different Rollbacks on different data Volumes, all without any issues. Again, from a user perspective, it is very straightforward. Under the History/Snapshot tag, you simply select which data Volume you want to perform a Rollback procedure on, then zoom in the specific time period, and each Volume data change will be highlighted with a blue marker.



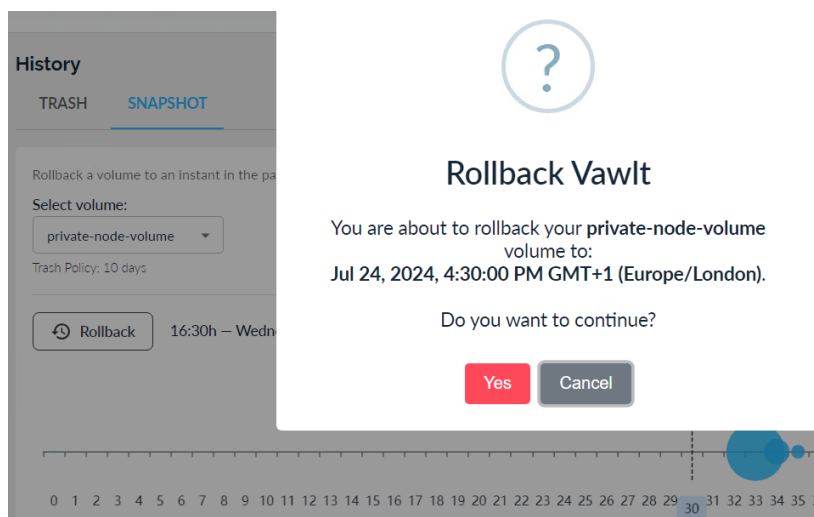


Figure 10 – Snapshot Rollback Example

You then simply choose a point in time to enable the Rollback function and click on the Rollback button. Where this is especially important is in the case of – for example – a ransomware attack; in the event of a large-scale attack, this could involve literally millions of files being encrypted and renamed, yet this potentially business-ending scenario can be resolved by one simple – and quick – rollback procedure.

The Vawlt solution then runs the Rollback as a background function. As with the other processes we tested, these were completed very quickly, though this will obviously depend on the size of the Rollback data. It’s important to note that all procedures and actions – for example, creating a volume and providing access rights - can be confirmed by email updates to the user, as defined during the configuration, allowing for hands-free management from anywhere, but meaning the user is kept up to date at all times about the system status – and their own.

Immutability

We’ve already defined the Immutability capabilities of Vawlt. During the testing, we created every variation on file and folder Immutability we could think of, with different data Volumes and put it to the test. Every data operation – whether we were deliberately blocking an action, or enabling free data manipulation, worked perfectly. As just one example, we created a Volume “budget-safe-files” with an immutability rule defined by kicking in at a future date. We tested the rule by deleting and renaming files and folders before the policy was enforced by the system without problems. With the rule enforced we were still able to rename folders, but every other action was blocked.

We then created a reverse Immutability rule with a Volume called “home-office”, where we allowed folder modifications but not renaming. We were able to add files to the folders, but were blocked from renaming them:

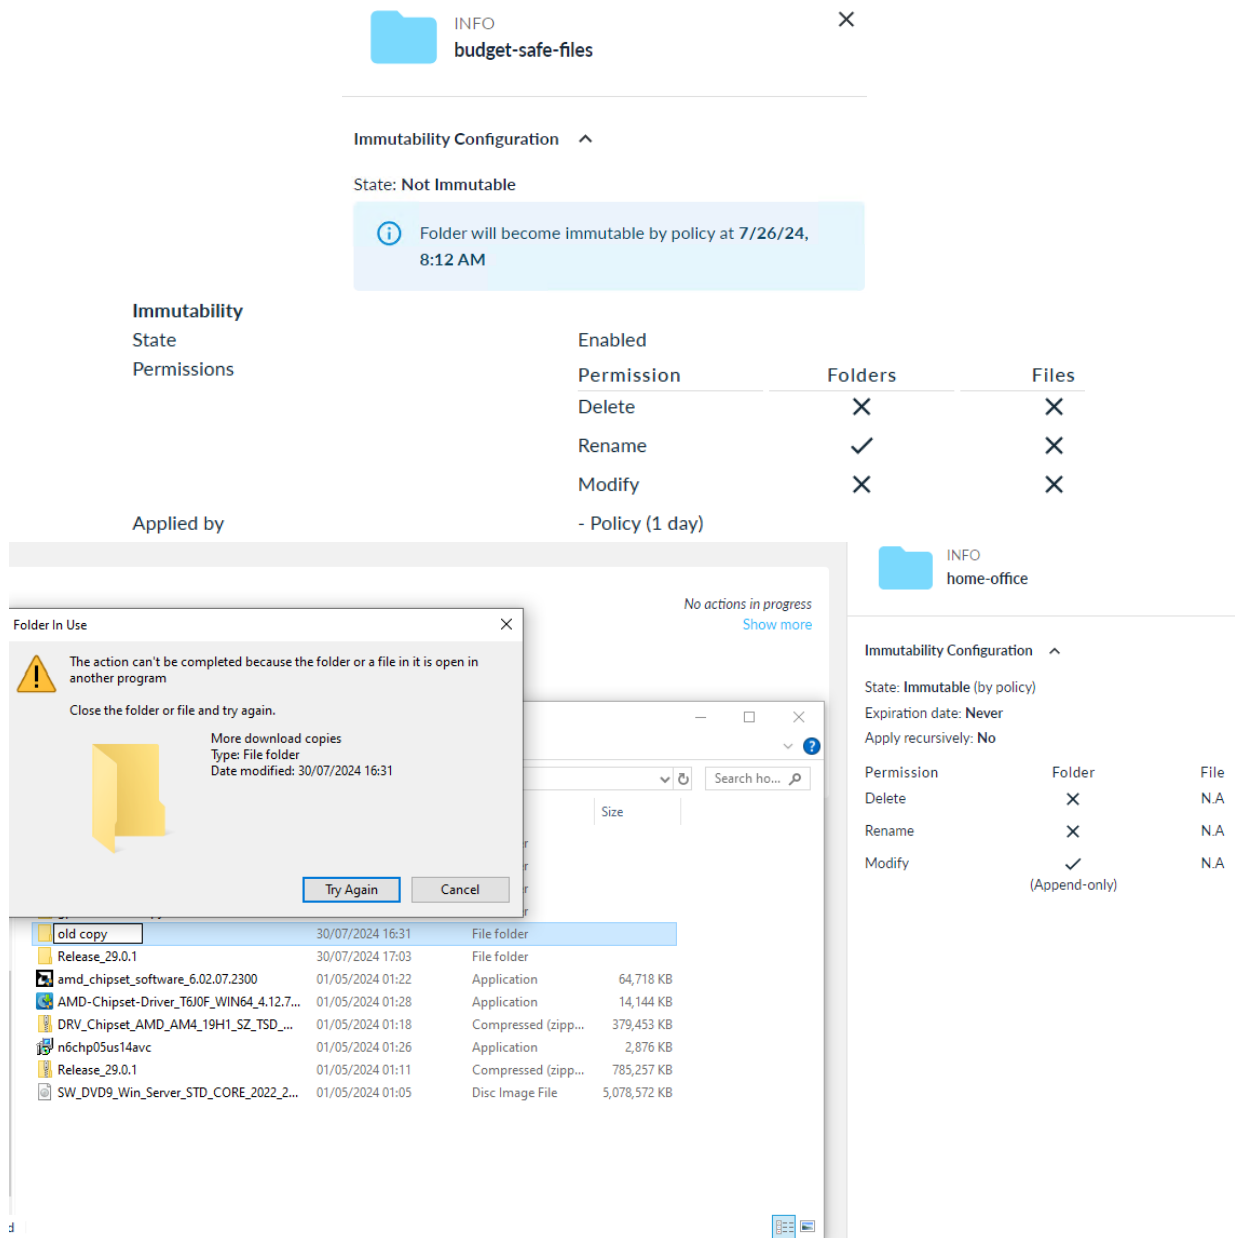


Figure 11 – Immutability Examples

A final example showing file deletion Immutability is where we created a Volume called “short-term-store” where we also allowed file renaming and appending of data, but not modification of existing data.

The point is, the Vault platform allows you to be very flexible in terms of the precise combination of actions allowed – and not allowed – on any file or folder, depending on the user.

Volumes / short-term-store

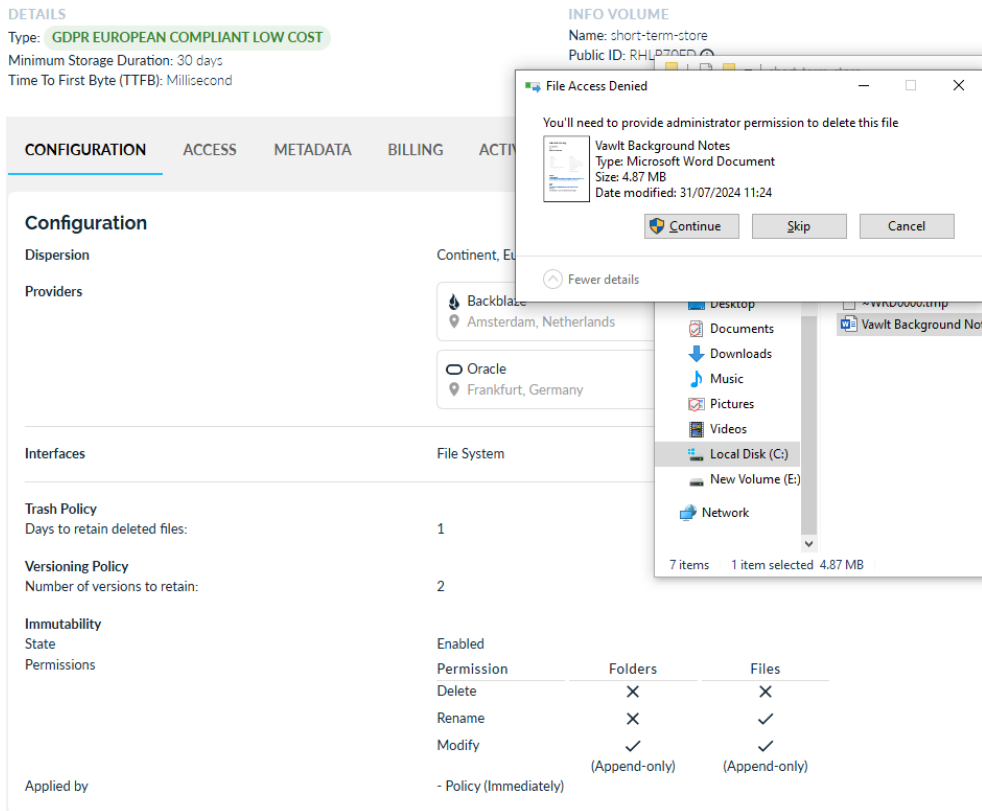


Figure 12 – Immutability No File Deletion Example

While data Immutability is powerful in its own right – both for protecting the data and protecting users from their own errors, that latter issue can easily result in an unintended security vulnerability, maybe even leading to a Ransomware attack. This is not scaremongering – it really does happen, as the many and well-documented examples that have been making news across Europe this summer and autumn prove.

In creating an Immutable file architecture, Vawlt doesn't simply provide a company with improved defences against ransomware attacks, but it effectively comes free of charge, as well as being fully integrated. No cost – no disruption. Factor in the extensive data restore and snapshot rollback capabilities, as well as the distributed data encryption architecture on which Vawlt is based and the product becomes almost as much a security platform as a data storage and management tool.

Use Cases

As we have noted, the use cases for Vawlt's platform are many and varied.

Primary uses cases can be identified as:

- Active and long-term archiving
- Ransomware protection (as noted above)

- Freeing up data centre resource
- Virtual NAS deployments
- Being a cloud-enabler for legacy applications
- Creating a secure cloud gateway
- Offsite backup and disaster recovery (including via integration with 3rd party products)

Integration Use Case

Since Vawlt is running an ongoing integration program; we tested this with VEEAM’s data software, creating a synchronised, scheduled job to perform backup. In conjunction with Veeam we were able to create a Vawlt S3 backup and schedule it as a recurring backup job.

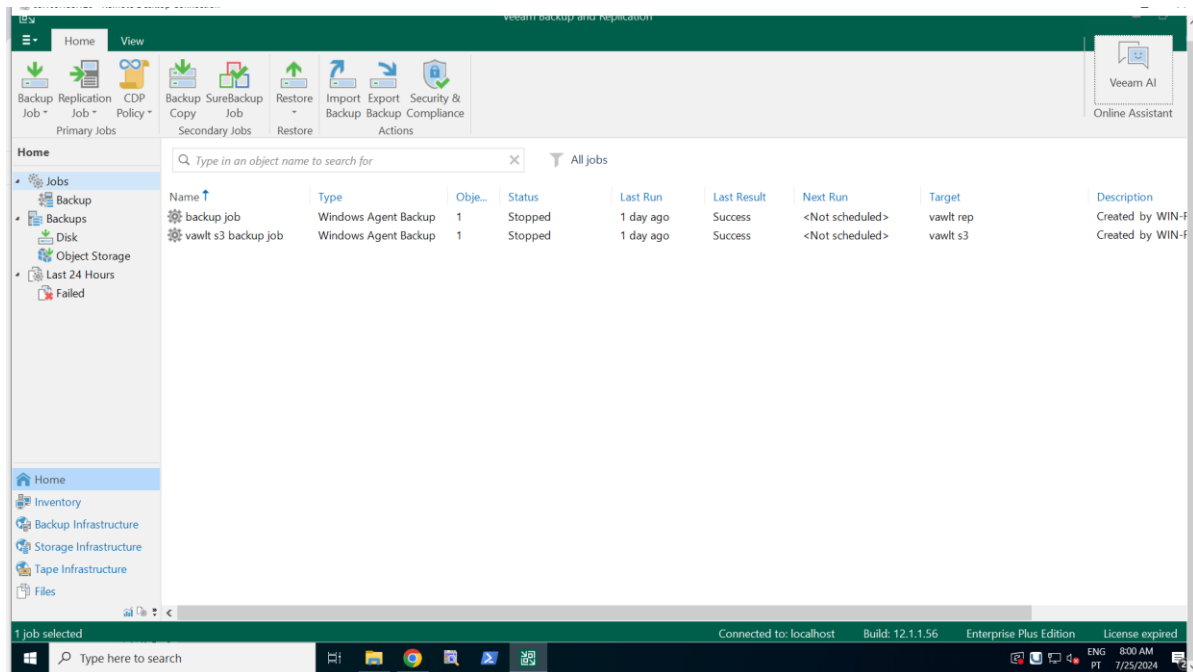


Figure 13 – Veeam Integration

Note, this integration program already extends to other vendors such as Commvault, Veritas NetBackup and Storware and will be expanded over time. Due to Vawlt providing both S3-compatible and file system interfaces, integrating it with *any* backup solution is very straightforward. Given that most backup solutions are capable of storing data using at least one of these interfaces, it makes Vawlt a highly compatible and flexible option for ongoing, seamless integration.

IN CONCLUSION

While cloud computing is undeniably here to stay for the foreseeable future, it is equally undeniably flawed in many ways, not least in terms of costs and manageability.

What Vawlt is attempting – and it is very ambitious – is to create a kind of “supercloud” that makes a multi-cloud/hybrid data storage approach not only far easier to manage, but also far more secure and significantly less expensive. That is surely the definition of “ambitious”? However, the company has achieved what it set out to do. Yes, the product will progress and there is a very logical roadmap being followed, but even right now it is definitely out on its own as far as we can see.

Its “supercloud” approach really is offering a “one-stop shop” solution in enabling the creation of that ideal hybrid/multi-cloud data environment most businesses across the globe are trying – and often failing – to implement. Moreover, it is looking to do that in an incredibly secure fashion, with the added benefits of reducing costs, easing the management overhead (enormously) and improving performance – all qualities we questioned at the beginning of this project but which have all been proved conclusively during the testing.

And, from a security perspective, it is worth reiterating – in a year that has seen relentless, successful ransomware attacks in both the private and public sectors – that Vawlt also brings a very high degree of what is effectively ransomware immunity to the table.

We also liked the way the Vawlt platform is designed to accommodate both the service providers and their customers, or indeed an enterprise customer running its own internal data management. It is a true single solution, regardless of whether you are in the business of selling or acquiring those cloud services.

Overall, then, we are extremely impressed by both the braveness of Vawlt and what it has achieved. If you are a service provider looking for a multi-cloud platform to generate (lots of) business from, or are a company looking for the multi-cloud/hybrid Holy Grail, then it would be unthinkable to not consider the Vawlt platform.

