



Areas IT professionals need to consider when building a private storage cloud

White Paper

wanstor

Contents

- + **Introduction**
How do you start building your own private storage cloud?
- + **Creating elasticity**
The first area a business should evaluate in terms of their private storage cloud
- + **Choosing a platform**
The key to efficient hardware deployment
- + **Required APIs**
A minor management evolution in the storage appliance world
- + **Workflow matters**
The final piece in the puzzle when delivering a private cloud solution
- + **Stack Deployment**
Building storage automation into frameworks for effortless provisioning
- + **Public Cloud Integration**
Private clouds and Public clouds as a hybrid
- + **Final thoughts**

Introduction

Storage as a service has advanced considerably over the past 3 years.

So how do you start building your own private storage cloud?

To start with, businesses should first of all take a step back and review what cloud computing really means to them.

The standard definition of cloud includes elasticity to grow and shrink resources as and when they are needed, services delivered on demand defined in abstract terms rather than on physical hardware, multi tenancy to support multiple clients, access for requesting resources with little or no manual intervention, and detailed reporting and billing for services consumed over a period of time.

All of these points are must-haves for public cloud, and at Wanstor we believe private storage clouds should reflect these same features. For example, business users should be able to request storage capacity without having to worry about how it is delivered.

This means service catalogues should be refreshed to reflect this change in user behavior and be more focused on service metrics. Using terms like I/O density, latency, throughput, data availability and resiliency.

Multi tenancy in the public cloud is fairly obvious for what it means, in the private cloud it should cover security and performance isolation. The right security makes sure data isn't visible between private storage cloud users, while performance features like quality of service make certain each user receives a consistent service level, irrespective of system load.

On-demand access guarantees customer requests for resources with minimal intervention from IT. Reporting has to cope with a more granular measurement of storage utilisation, including the capability to report on individual teams or business areas.

Creating elasticity

The first area businesses should look at in terms of their private storage cloud is *elasticity*.

Elasticity comes in two scenarios: firstly from the customer's ability to expand and contract usage on demand and, secondly, for system administrators to be able to deploy more infrastructure as demand warrants.

While some system administrators worry deployed hardware won't be used if end users can easily return storage, in reality this rarely happens as data is growing exponentially in many businesses.

The challenge here is to continue meeting demand by deploying new hardware into the data centre, whilst managing the technology refresh cycle without impacting application availability.

Achieving the *'just-in-time'* deployment of new hardware is both an art and a science for most IT departments, most of which have limited amounts of cash and staff at their disposal.

They must compromise in their inability to provide infinite resources by predicting when demand justifies new hardware acquisition.

This is where subject knowledge comes into play for IT professionals. Demand forecasting requires engagement with lines of business to plan potential future projects and their storage needs.

The science comes into collecting information on storage growth. Many IT environments use thin provisioning, which means physical storage capacity increases over time as data is written to allocated space.

Because reservations for planned consumption of storage are rarely fully utilized immediately (e.g. a 1 TB request may only use 100 GB on its first day and be sized for 2 year growth), file systems and object stores will naturally increase in utilisation as applications write more data.

This makes it essential to have accurate and detailed tools to measure storage consumption over time, while using that data to generate meaningful growth projections.

Additionally, determining when to deploy new hardware requires understanding and managing vendor lead, hardware deployment and configuration times.

Insight into potential future storage resource requirements means demand can be more easily anticipated, especially in non-core products like object or high performance storage

In most businesses, IT still owns these issues, which are, of course, under the cloud service provider when purchasing public cloud storage.



Choosing a platform

Having the right storage platform in place is key to deploying new hardware with efficiency.

Having the right storage platform in place is key to deploying new hardware with efficiency. Scaling out choice of services over one particular storage technology can make new deployments relatively simple, as all an IT Manager has to do is add hardware to the existing configuration to increase capacity.

Most modern object and block scale-out products perform some level of rebalancing, redistributing data to make use of new capacity and gain the most performance from the hardware.

Monolithic scale-up architectures can be harder to manage because of scalability limits, while older legacy storage systems may not naturally load balance to make use of new physical capacity.

This means legacy architectures have to be more carefully designed to load balance the distribution of logical resources over physical hardware. Many of these platforms come with tools to move large data units around within the storage platform, mitigating some of the balancing problems.

Multi-tenancy and quality of service become key features to consider when choosing a storage platform for a private storage cloud.

Looking at service metrics offered by cloud service providers, we can see that performance is rated on IOPS and throughput, with some mention of I/O latency. These service levels are deemed available to users whether the cloud service provider is running fully loaded or not.

Quality of service therefore, becomes really important, either as a tool to make sure end users get the performance they require or to limit performance to the level they have paid for with new systems.

APIs required

In recent years, the storage appliance world has seen a minor management evolution.

Traditionally businesses managed storage manually through GUIs and some command line interface interaction, with *commit* phases to make changes. Command line interfaces enabled the storage administrator to script the provisioning and decommissioning process, allowing a degree of automation.

However, creating the scripts was a time-consuming process. Over the years, vendors moved to implement APIs that make storage programmable and to set configurations via authorised API calls. Configuration data can also be easily extracted now, with some storage platforms producing very detailed metrics.

APIs also enable automation, taking human processes out of provisioning storage. Now, storage can be mapped to hosts through one or two API calls. Some platforms implement APIs natively, others have built API wrappers around existing API tools.

The crucial requirement here is to ensure APIs, Command Line Interfaces and GUIs operate in sync together rather than as silo'd pieces which step over one another.

```
{
  "login": "",
  "id": 1,
  "avatar_url": "https://com/images/error/",
  "gravatar_id": "",
  "url": "https://api.com/users/",
  "html_url": "https://com/",
  "followers_url": "https://api.com/followers",
  "following_url": "https://api.com/{/other_user}",
  "gists_url": "https://api.com/gists{/gist_id}",
  "starred_url": "https://api.com/starred{/owner}",
  "subscriptions_url": "https://subscriptions",
  "organizations_url": "https://api.com/users/orgs",
  "repos_url": "https://api.com/users/repos",
  "events_url": "https://api.com/events{/privacy}",
  "received_events_url": "https://received_events",
  "type": "User",
  "site_admin": false,
  "name": "",
  "company": "Hub",
  "blog": "https://hub.com/",
  "location": "",
  "email": "@hub.com",
  "hireable": false,
  "bio": "...",
  "public_repos": 2,
  "public_gists": 1,
}
}
```

Workflow matters

The final piece of the puzzle to deliver private cloud storage is executing *workflow processes*.

User requests have to be serviced in a way that allows a request to be validated and then implemented. Public clouds achieve this validation process through users providing a credit card or other payment method for billing. After that, services are configured through a web portal or API.

In a business, the traditional process for requesting storage has been to have internal processes that manage requests manually, provisioning storage to hosts based on details in service tickets.

The owner of the ticket takes responsibility to make sure the line of business is allowed to purchase the storage being ordered and then takes care of fulfillment.

In a private cloud environment, the IT departments aim should be to make this process as automated as possible. Tools which allow you to build workflow processes around storage automation and other resource provisioning are crucial.

There are many vendor tools out there which will address these areas, but at Wanstor we have found it is best to select a few vendors, test the tools to see if they can deliver for your business and then think about deployment and in life management.

Following this approach eliminates mistakes at an earlier stage, so businesses benefit from the right process reporting tools sooner.

Many businesses will need to consider the changes in billing that a private storage cloud can introduce. If billing and chargeback are not implemented, then there is no issue because the IT department will continue to take the hit on the cost of delivering the service, and charge per project.

However, if new resources have to be paid for, then some changes to financial practices (which may include IT paying directly for hardware) will need to be introduced to allow service-based billing of business units to cover purchasing and administration costs of a private cloud.



Stack Deployment

Taking a wider view of storage requirements for private cloud, IT teams can build storage automation into frameworks such as OpenStack to take the effort out of provisioning.

Initial OpenStack deployments had no persistent storage capability, so a number of projects took place to manage the integration of external storage arrays.

Many hardware and software companies already support all of the OpenStack storage APIs. Additionally, storage vendors can write plug-ins that enable the OpenStack framework to provision and map storage LUNs on demand.

Many hardware and software companies already support all of the OpenStack storage APIs. A Cinder support matrix lists supported vendor features in each release of the OpenStack platform.

figure 1 : Vendors that currently support OpenStack

Public cloud integration

Moving forward, businesses need to recognise that the cloud world will not be exclusively public or private, but a hybrid of the two.

As a result, there will be requirements to move data and applications between public and private cloud infrastructure, with the latter offering additional data protection and increased availability.

Products for moving applications and data between private and public cloud locations are starting to come to market. Object storage vendors are starting to provide the ability to archive on-premises data into the cloud, while crucially retaining the ability to search across all content as if it were in a single view.

For data protection, Druva and Zerto both offer products that allow you to back up and restore on-premises virtual machines (VMs) in the public cloud

The conversion of the VM image, and injection of additional drivers, is handled in software as part of the backup and migration process.

Velostrata goes one step further by allowing the booting of VMs in the public cloud for cloud bursting.

This could be used to run an application on a VM with greater resource capacity than is available on site, or move workloads to the public cloud to cope with increased demand. Once the surge in demand has passed, the VMs can be returned back on site.

In the meantime, virtualization vendors are starting to partner with cloud vendors to facilitate the migration of applications into the public cloud.

For example, VMware now has partnerships with AWS and IBM. It has also introduced Cross-Cloud Architecture as a way to manage multiple cloud deployments.

Microsoft Azure is more of a one-stop shop offering the same functionality of Azure in the cloud to be run in a private data centre with a link to the public Azure cloud.

Final thoughts

We believe that a successful cloud storage programme begins with behaviours of the IT team and IT partners who supply your business. Private cloud providers need to have a positive *can do* attitude and be focused on putting in place service delivery models which empower IT teams and users in a business.

Only by truly understanding the business objectives and how private cloud storage can help the business achieve these objectives can a private cloud storage programme be successful.

At Wanstor, we provide a range of private cloud services to help businesses at every stage of their journey. We can help you with planning, deployment and in life management of private clouds, giving you access to best practices we have deployed with other customers.

By working in partnership with Wanstor, businesses can adopt private clouds that bring tangible value to the organisation.

For more information about Wanstor's cloud services please contact us on **0333 123 0360** or email us at info@wanstor.com and one of our cloud computing experts will give you a call back.

+ + +

