Storage in the Cloud

2e2 in conjunction with HP

It would seem fair to say that Cloud, combined with the concepts of shared infrastructure, is not about to go away. Cloud means many things, but most importantly it means flexibility, agility and reduced cost. Shared infrastructure provides significant cost benefits and efficiencies from pooled resources, allowing organisations to buy less, utilise less and devote less time on management resource. This article examines the implications for storage systems in a Cloud world and how potential limitations and inhibitors can be either reduced or eliminated.

The Cloud can often been seen as a generational phenomenon. As each new generation enters the work stream, they bring with them expectations of how they wish to use IT services. The current generation is not used to the concepts of long lead times, development and procurement cycles and rigid, structured hardware. If they want email services, they get them from Google. If they want a storage solution for music, photos or video they go to the internet and rent a service. If they want a new app, they download it. This is a world apart from the traditional IT department's view of things, and has led to numerous standoffs, with devices being banned and access to social networking sites restricted. Security and governance concerns are the most cited reasons for this reticence, but when Cloud based solutions are seen to be able to add value to the business (corporate communications via Twitter, viral marketing campaigns via Facebook for example), they can no longer be ignored.

This generational trend reinforces all current research that suggests that the demand for Cloud based services is going to increase rapidly, as will the requirement for shared infrastructure. Organisations in both the public and private sectors have seen the cost benefits that can be achieved with pooled server resources and the creation of Virtual Machines. However, sharing server facilities is far easier than sharing storage. Let's look at some of the problems with shared storage in a Cloud environment.

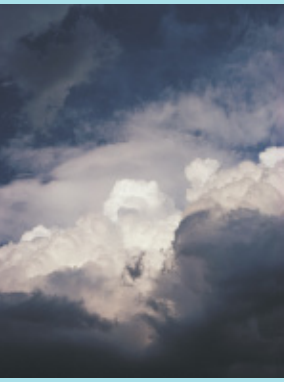

Dark Clouds?

Virtualisation of servers has produced computing capability that is powerful, manageable, scalable and cost effective; solutions have been developed that address the whole breadth of security concerns including access, compliance and governance. This leaves the question of storage as the key remaining "blocker". Why? Here are some of the reasons:-

- Time to deploy. It is recognised within most IT departments that storage is cumbersome, complex and slow to deploy. Slow deployment equates with delayed return on investment and means that it is impossible to respond quickly to changing circumstances. This is the exact opposite of what Cloud users need that is rapid implementation and the ability to overcome the limitation within a given infrastructure, the storage element always takes the longest to get right.
- Inflexible. Configuring and optimising data layouts for specific workloads and applications take time. Once optimised, an additional application is added, the user must either accept sub-optimal performance or go through the entire configuration process again. In the new world of Infrastructure as a Service, that becomes an impossible demand because customers (internal or external) want to introduce new services quickly and unpredictably. Storage has to be able to meet these demands.
- Silos. Traditionally, storage environments have not only been siloed by application, but by either service level or by cost/performance criteria. Users may well have separate T1 mission critical silos, T2 for general data and T3 silos for archive, but now they want an infrastructure that is shared and virtualised. This will allow them to deliver different service levels within a single environment. This in turn, allows them to dynamically and non-disruptively migrate data between different levels of service.
- The establishment and achievement of an ILM policy is often seen as the holy grail of storage. This is complex to set up and difficult to achieve in a multi-tenanted Cloud environment with traditional disk storage technology.
- Inefficient. Traditionally, storage has been significantly over-provisioned. In many large or ageing environments a simple summation of the used data storage at the application level compared with the raw capacity of the storage systems will reveal utilisation rate as low as 10% to 15%. To perform acceptably in a Cloud world, utilisation rates must improve significantly.
- Labour intensive. Storage is typically seen as requiring a high level of staff involvement when compared with the other elements that make up Cloud infrastructure. With the emphasis on cost saving and the need to do more with less, improvements are required.

It is clear from the above that for organisations in both the public and private sectors to be able to take advantage of the Cloud based solutions the current siloed methods of storage, with their high costs of ownership, limited flexibility, poor efficiency and high management overheads, must change. But how?

It's looking brighter

As with servers, Virtualisation is key. Virtualisation allows for the separation of logical storage from physical storage, delivering enormous flexibility in the management and utilisation of the storage capacity. This delivers a number of benefits:-

- Improved utilisation. The link broken between the physical and logical storage, utilisation can be increased by pooling the total resource. There is no longer a need to allocate specific storage areas or devices to an application, storage can be assigned to wherever it is needed at any point in time.
- Dynamic Data Migration. Being able to dynamically migrate data without interrupting an application is of significant benefit to the storage administrator and delivers real benefits to the end user. Data can be moved off over utilised devices or moved to faster storage as business needs change and the implementation of ILM policies becomes easier to achieve.
- Reduced management overhead. Whilst a significant level of traditional storage management is still required the overall burden is reduced when the entire facility is presented as a single, vast, storage area that can be managed centrally.
- Multi tenancy. In a Cloud environment it is vital that multiple tenants, be they from different organisations, or sub groups of the same user, can co-exist, with the ability to manage, secure and administer their data storage in a way that is most appropriate for them, without adverse effect on other parties.
- Thin provisioning. By allowing for the on-demand allocation of data blocks, rather than traditional "up front" requests (fat provisioning), thin provisioning delivers vastly improved utilisation rates, with little administration overhead. There is a reduced requirement to acquire storage capability up front, leading to a further reduction in under utilisation and an improvement in cost effectiveness.

Summary

When considering storage, as with all aspects of Cloud development, not only are there a number of complexities that need to be considered and resolved, there are many inter-dependencies. Successful implementations benefit from the involvement of an integration partner that has the breadth of experience to be able to bring all these strands together as well as the consulting skills required gaining a deep level understanding of the business drivers.

The Mansion House
Benham Valence
Speen
Newbury
Berkshire
RG20 8LU

T: +44 (0)844 2250526
F: +44 (0)1635 568001

E: info@2e2.com