SCALING AND OPTIMIZING CLOUD STORAGE

Stefan Vlaev, Strahil Sokolov and Mihail Vukadinoff

***Abstract**: This article provides an overview of current practices for scaling the infrastructure of cloud providers. Approach is provided for optimizing the existing storage solutions in order to achieve better performance and scalability.*

***Keywords**: Cloud Storage, Optimization, Scaling, Software Defined*

***ITHEA Keywords**: B.3 Memory Structures, E.2 Data Storage Representations*

Introduction

The last trend in the IT industry is "transformation". It requires business to change the model of how they develop new application, how they deploy them, how they support them and even how they store the end user data. Faster access to the information becomes crucial in the era of cloud computing and fast mobile networks.

Storage is the most important component of the infrastructure. It represents the fundamental system element that if properly optimized can lead to performance issues, downtimes and even data loss.

To solve this problem with the already existing storage solutions they should be able to scale and be open for optimization. The scalability of a solution is vital for cloud providers.

The authors argue that scalability is an important characteristic of cloud computing. With scalability, cost is minimized by provisioning and releasing resources according to demand. Most of current Infrastructure as a Service (IaaS) providers deliver threshold based auto-scaling techniques.

Setting up thresholds for scaling with proper values to reduce cost and fulfill Service Level Agreement (SLA) is cumbersome and rather subjective task, especially with volatile workload changes. The authors proposed dynamic threshold based auto-scaling algorithms that predict the required resource capacity using Long Short -Term Memory Recurrent Neural Network and auto-scale virtual resources based on predicted values.

The rest of the paper is organized as follows: in section 2 motivation for the paper and the existing problems is presented. In section 3 there is scaling and optimization, Section 4 provides conclusion about the topic.

Motivation and problems

1. Background

The traditional storage system provided by vendors like Dell, NetApp and IBM are monolith systems. In such system a single controller or pair of controllers will know where all of the data is stored on the system. Such concept is was perfect couple years ago but with the modern flash and high capacity drives it becomes obsolete. Despite being able to save terabytes of data they are not able to scale well like the software based systems

With the demand coming from the cloud providers and enterprises the vendors started to develop products that were able to scale in mater of number of controllers. Still these solutions were with limited flexibility and with very complex architecture that was not able to respond to the customer demands. Since its creation Backblaze one of the biggest cloud storage providers have stored more than 500 Petabytes of data. Storing such capacity with traditional monolith systems will be extremely expensive.

In order to increase availability that business are requiring the vendors are implementing clustering as a solution. With this technology is possible provide both scalability, performance increase and availability. The problem in this case is that not every system supports it. However, this is not supported from all vendors. Traditional products from companies like Dell/EMC, IBM and HP are not supporting such technology in their lower-end systems. The storage hardware vendors that support this is NetApp. Still their solutions have limits and they require special interconnect between the nodes.

With the transformation of how data is stored and analyzed there is an increased demand for better performance from the existing solutions. The traditional storage access protocols like FC,iSCSI are expensive for implementation. The monolith systems that were popular before the 2012 are not flexible enough for the rapidly changing customer requirements.

1.1 Scaling

The scalability of the IT infrastructure is key component for both Cloud providers and IT organizations. When the infrastructure is planned for prolonged use the all of the components should be chosen with scalability in mind. A good storage system should be able to scale from terabytes to petabytes without increased costs and downtime for the organization. Usually systems from vendors like DELL/EMC can

scale up to certain amounts of storage space, network ports and storage protocols. Software defined products are more flexible in such cases. Storage system build around ZFS can scale from terabytes to zetabytes with only limitation being the used hardware components. Also this type of systems can scale not only in size of the data they can store but also can adapt to new storage protocols used and additional network connections.

1.2 Optimization

Optimization of existing storage infrastructure requires not only good planning but also solution that is flexible and open for changes. The monolith systems are not able to easily change the storage protocols supported and also are not open enough to use components other from the ones provided by the vendor.

Additionally, with traditional storage vendors most of the advanced controls over a system are locked from the customers. Example for this is EMC VNX a storage system that is on the market for more than 20 years. Some of the tasks on this type of systems can be automated only from the command line using specific interface.

2. Software defined future

The infrastructure of the future is a software. This mean that all of the components of the infrastructure will be defined by software. Still this software should be open for modification and with limited vendor locking.

2.1 Open hardware and software

The infrastructure at the moment is with high level of vendor locking. In order to provide increased flexibility for the cloud providers the used hardware and software should not be tied to specific vendor. It is valid for hardware infrastructure components. This will provide flexibility for the cloud companies to modify the software and to have better control over the used hardware. With increased control over the components they can be easily modified for more specific tasks.

3. Scaling and optimizing Cloud storage

3.1 Scale-UP vs Scale-Out storage

When deploying storage infrastructure for cloud provider capacity is key point. Simplicity is a key feature for cloud storage and backup. As stated in our previous article "Software defined Storage Criteria for Cloud Providers" the vendor locking is problem for the cloud providers.

The traditional storage systems are built around the concept for scale-out. This means that when additional capacity is required the cloud providers should add JBOD's and the limit is reached completely new system should be added. Usually this type of systems is very complex meaning that only handful of disk shelf's are compatible and they are locked to the specific vendor. This means that they cannot but used in future systems. When the system's capacity limit is reached the new system should be added including possible migration of data. For cloud customer's availability is key component.

The data protection in such systems is based on the use of RAID technology. This is very limiting and complex technology. Different types of disks cannot be mixed in one group. In some cases, the configuration is so complex that different types of disks cannot exist in the same JBOD. It becomes also a problem when capacity should be extended.

The scale up systems are the opposite. They are created with the idea of using common components as much as possible. This type of scaling is additional flexibility where the disks are not connected in architecture similar to Direct Attached

Storage. This means that only individual nodes will have access to the disks connected to them.

This type of configuration is capable to grow more than a traditional system with a limiting factor only the number of disks that is connected to a node. In the end the end-user and its application are not aware where the data is stored. Since RAID is not used the data protection is software based. Such technology is used by Storage solutions like Ceph.

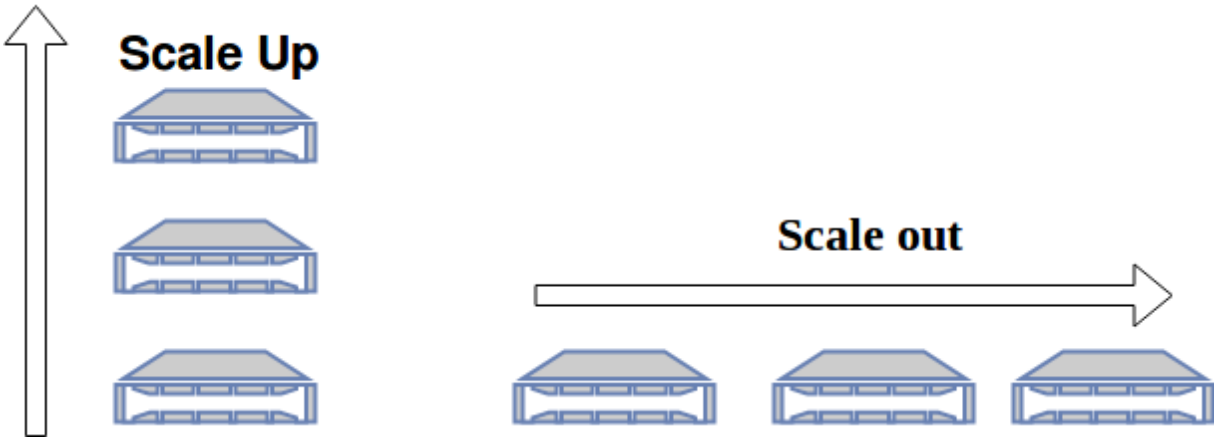

Figure 1 Scale up vs Scale out

3.2 Optimization

After a storage system is deployed and used for couple of years the performance starts to decrease, and the initial workloads changed. At this point the optimization of the system became very important task. There are several ways to perform that.

3.2.1 Pooling and abstracting storage.

Storage pooling is method where all of the available drives are added in a single storage pool. The administration in this case is not complex and also tiering can be implemented as well for performance optimization.

Another way of increasing performance can be storage abstraction. In this case multiple storage arrays can be presented as a single system. This can be performed by using both Software Defined Storage and solution like NetApp gateway. In this case all of the space of the initial system is presented as predefined LUN's to the gateway which acts like hypervisor for the client systems. Using such abstraction, it is possible to perform using only single replication license, advanced caching technologies and to add additional protocol support.

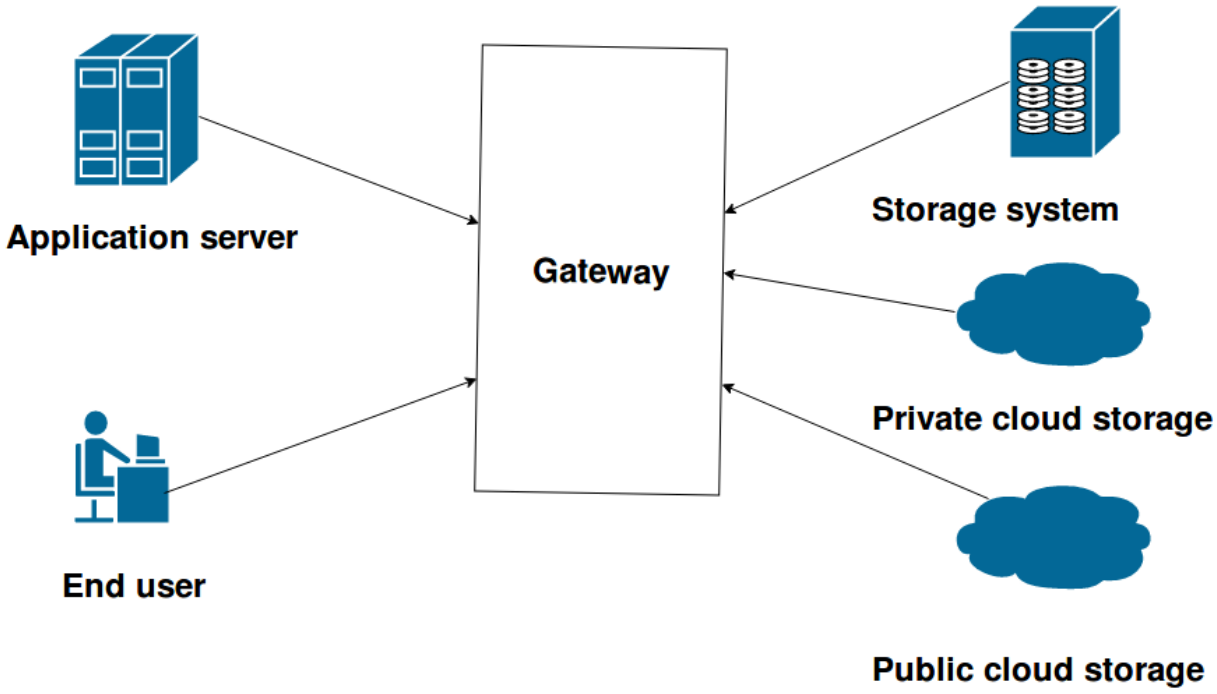

Figure 2 Abstracting storage

3.2.2 IO acceleration.

Using policies around the data used and optimizing the workloads is a way to accelerate the IO performance. For example, the use of QoS for certain workloads can not only increase the performance but also can reduce the risk of downtime.

Intelligent caching solution are another way for performance increase on most storage arrays. This valid for Software defined storage and hardware systems. Depending on the vendor there are several ways for implementation.

Adding a write cache to for workloads like backup can increase the performance and reduce the cost. On traditional storage arrays the speed of the system depends on the number of drives used. If the backups can be performed faster this can free resources for other tasks. Azure, AWS and Backblaze are providing backup as service to its customers and their customers require fast backups/

Read cache can increase the performance when read performance is a critical component. Good example for this is Databases, Big Data analytics and backup restores. The traditional way to have or to increase the read cache in system was using bigger storage controllers and expensive PCI-e cache cards. With Software Defined Storage this is avoided and now traditional SSD's can be used for a read cache. ZFS for example can add READ and Write cache to his pools at any given time. The monolith storage systems on other hand depend heavily on the vendor to allow such advanced functions.

3.2.3 vVOL.

Creating a huge single container for storing the entire contents of a VM, including all metadata can increase the performance of the system. Still this granularity can lead to more complex environment for administration. In such cases the automation of a system can be a key factor. That kind of granular control allows you to have very specific levels of policy management around the VM. This can lead to reaching the limit of LUN's that can be controlled by a single system. Already a lot of vendors are working on a solution how to avoid such scenario. Still this can be done for small deployments and will not ideal for cloud providers, where there are a lot of subscribers.

3.2.4 Data mobility.

Replicating and moving data can be a problem solver for a lot of organization. This can be implemented by both using SDS systems on site, remotely and even using cloud storage. For cloud providers the ability to move the customer workloads closer to them is key factor in some scenarios. As for the customers the ability to move the data inside the cloud can be a cheap backup solution which required a whole new storage solution before that. The ability to perform replication between storage systems without requiring expensive licenses and with smaller bandwidth limitation is very important for providers.

3.2.5 Policies, features, and built-in optimizations.

Implementing QoS policies, data compression, and deduplication can optimize the workloads significantly. Both traditional and SDS vendors are implementing large number of features both to increase the customer control over the data. Usually these built in features are not properly configured.

Tiering for example can increase performance when implemented correctly. When it is implemented the "hot data" is moved to the faster drives and when it is not need it can be moved to slower drives. When tiering is used it should be implemented after careful consideration what type of workloads will be on the system. It is not recommended to be used for heavily accessed data bases. Also if not the correct algorithms are not used is can lead to performance issues.

Deduplication can save space in some cases. Typical use case for a workload that can benefit from deduplication are backups and image libraries. For a cloud provider is typical to have a large number of preconfigured images of operating systems, databases and other. In most cases this images are access only when the customer requests them. Backups are often very static and accessed only when there is an issue on the main system. They take a lot of space, if this space can be reduced the provider can use it for other purposes. Using deduplication requires a lot of CPU resources because of that this feature should be used with caution.

3.2.6 REST APIs.

All cloud providers are looking at ways automate their storage platforms. Native APIs to perform such automation. By being able to directly integrate with products like, OpenStack, CloudStack, and even the IBM Cloud organizations get a lot more flexibility around their data.

3.2.7 Encryption and security.

With the increased data privacy concerns adding encryption to customer data is very crucial for cloud providers and government organizations. Protecting user data becomes very crucial aspect of every IT organization. Most of the vendors already support such features on their storage systems. However, this may require additional investments mostly in supported hardware. Good example for that are the encryption ready hard drives. The main difference between them and a standard hard drive is the encryption chip. This hardware module help to offload the encryption directly to the hard drive and remove the need of powerful controllers. Unfortunately, such hardware in most cases if very expensive. With the increased performance of the today CPU's and many Software defined storage providers are offering such functionality.

Acknowledgement

This work is supported by the project NIPD-10/2017 - "UNICLOUD" at the University of Telecommunications and Post, Sofia, Bulgaria.

Conclusion

With the use of software defined technologies cloud providers can have a very scalable solution on a price which cannot be achieved by using traditional monolith systems. Also using the techniques described in this article the already installed solutions can be optimized for new workloads increasing the time when the storage system can be utilized. Additionally, providing features like encryption and replication can introduce new possibilities for customers.

References

Shahin, Ashraf A. "Automatic cloud resource scaling algorithm based on long short-term memory recurrent neural network." arXiv preprint arXiv:1701.03295 (2017).

Bacblaze blog https://www.backblaze.com/blog

Datacenter knowledge http://www.datacenterknowledge.com/archives/2015/08/28/seven-ways-optimize-storage

Authors' Information

Stefan Vlaev *– is with the University of Telecommunications and Post, Sofia, 1 Acad. St. Mladenov Str, Sofia 1700, Bulgaria, E-mail:svlaev@mail.bg*

Major Fields of Scientific Research: Storage systems, Cloud computing, Software defined solutions

Strahil Sokolov *– is with the University of Telecommunications and Post, Sofia, 1 Acad.St.Mladenov Str, Sofia 1700, Bulgaria, E-mail: strahil.sokolov@gmail.com*

Major Fields of Scientific Research: Biometrics, Cloud computing, Deep learning

Mihail Vukadinoff *– is with the University of Telecommunications and Post, Sofia, 1 Acad. St. Mladenov Str, Sofia 1700, Bulgaria, E-mail: vukadinoff@gmail.com*

Major Fields of Scientific Research: Cloud computing, Big Data.