ActualTech Media

noobaa

Five Key Factors That Will Define the Future of Unstructured Data Storage

Author: Scott D. Lowe
Partner, ActualTech Media
August 2016

Introduction

Every organization leverages data in some way. Whether that data is in the form of neatly designed databases or instead consists of files randomly strewn about the data center, the fact is that every business relies on data in order to operate. However, different kinds of data present different kinds of business and technology-related challenges in the modern enterprise.

In this paper, you will learn about where traditional storage has created friction in business operations and how object storage is becoming an increasingly popular way to handle the tremendous growth in the volume and management of unstructured data.

Contents

Contents 1
Introduction 1
Defining Unstructured Data 2
A Lesson in Traditional Storage Economics 3
- Architectural Rigidity 3
- Block- and File-centricity 3
- Lockstep Scalability 4
- Hardware Lock-in 4
- Location Lock-in 4
- Cost and Complexity 5

Reshaping the Storage Environment 6
- Object Storage 6
- Frictionless Deployment 7
- Agility and Flexibility 9
- No Compromise Feature Set 10
- Improved Business Outcomes 12

Translating New Architecture to Value and Reduced TCO 12

Defining Unstructured Data

As the name implies, unstructured data consists of data that does not have a well-defined structure. Structured data, on the other hand, is neat and orderly. Relational databases, for instance, are examples of structured data. Data is organized into tables, rows, columns, and indexes. You can easily run queries against that data to glean insight from it because you know exactly what kind of data exists in the field. Given that most companies like things to be organized, you might think that structured data is the norm.

Nothing could be further from the truth. By some estimates, a *staggering 80% of enterprise data resides in unstructured formats*, such as email message body text, Word documents, spreadsheets, photographs, and much more. Making the problem worse is the continued onslaught of unstructured data from Internet of Things (IoT) devices, sensors and video cameras that constantly relay information, and even items such as medical equipment. Without careful handling, the data being created by such devices can quickly devolve into useless information. Even though "Big Data" often comes in the form tables and records, from an infrastructure perspective, it is accessed in storage like unstructured data; this is often called 'semi-structured data'. This new class of enterprise data is exacerbating already serious storage problems.

With unstructured data, organizations face two challenges:

1. Where do they store the continually increasing amount of data that's being produced?
2. How do they make sense of all of the bits and pieces of data that are being captured?

Many companies attempt to use traditional block and file storage to meet all of their storage needs. However, as you'll learn in the next section, traditional storage is often not up to the task, particularly as you consider the economics of the solution.

Structured vs. Unstructured Data

Structured data is neat and orderly and is eminently searchable. With structured data, organizations gain insight into operations.

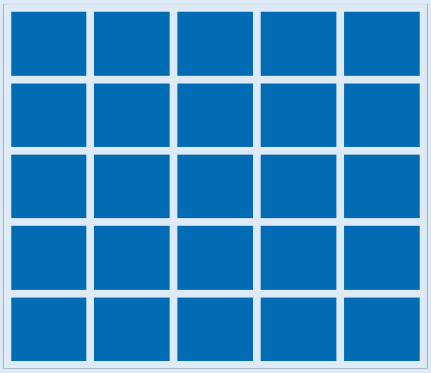

Unstructured data, on the other hand, is messy and it's hard to glean insight into it since there is little to no enforced consistency from item to item. It's also hard to find what you're looking for.

A Lesson in Traditional Storage Economics

Over the years, storage economics have become a front-and-center problem in the data center. Whether they call it storage economics or not, the term refers to the costs incurred by an organization to implement and maintain storage as well as the opportunity cost that the organization experiences as a result of challenges in the storage environment

As organizational storage needs rise exponentially, traditional storage solutions no longer make sense. They will certainly *function*, but the key is that where the marginal cost and complexity of storage needs to go down and down, these systems remain anchored in the old world, rendering them less and less capable of meeting changing business requirements

Here are six areas in which the traditional storage model breaks down.

1. Architectural Rigidity

Storage is not known for being all that flexible. In most cases, organizations are forced to make upfront decisions that impact the storage environment throughout its usable life. In a business environment that demands unprecedented levels of flexibility and the agility to enable the business to pivot at any time, many data center environments have become so complex and so rigid that any attempt to change one area results in a cascade effect that impacts the entire business.

Existing solutions require dedicated resources for storage, rack space, and networking. The truth is that most enterprises are horribly underutilized across their capacity landscape, with capacity 'stranded' in existing hosts.

2. Block- and File-centricity

When you're storing virtual machines, block and file level storage works really well. However, when you're storing millions of individual photos for your photography business, blocks and files completely break down. Traditional data center storage environments do a tremendously poor job when it comes to truly understanding what kind of data they're storing. They generally lack a means to capture *metadata* about each item being stored. Simply put, metadata is a mechanism that ascribes meaning to data. In unstructured data environments, having a means to wrap meaning around the items you're storing can be incredibly powerful. Immediately, you're able to tap the value of that data since you now understand what it is.

On the application development front—whether that's at the enterprise or web-scale level—there is an almost universal thinking about cloud paradigms. In a cloud world, file and block systems in many cases become liabilities, limiting application flexibility and scalability. As such, DevOps teams are biasing content-oriented applications more and more towards RESTful interfaces like S3.

3. Lockstep Scalability

Scaling capacity in a traditional storage environment isn't really all that difficult. You simply add more disks for new capacity and maybe some additional controllers to increase performance. But that's often where the simplicity ends.

Traditional file systems have inherent limitations that restrict the extent to which they can be scaled before an administrator has to intervene and create a new construct. Although these limits are periodically increased, the fact remains that they exist and will exist for the foreseeable future.

Furthermore, dependencies on features like file-locking fundamentally inhibit critical capabilities such as WAN flexibility and hybrid public/private configurations. Moreover, as you scale a traditional storage environment, you're forced to add resources in lockstep with one another. It's often impossible to decouple storage intelligence from capacity in any granular way, so you're often required to add resources you may not need, thereby increasing the overall total cost of ownership (TCO) of the environment.

Lockstep Scaling in Action

Being forced into lockstep scaling—that is, scaling both the data plane and the control plane—has proven to be an inefficient way to grow a storage environment. In most companies, the two planes experience different rates of growth. With a lockstep scaling paradigm, you have to deploy resources you don't need just to get the ones that you do.

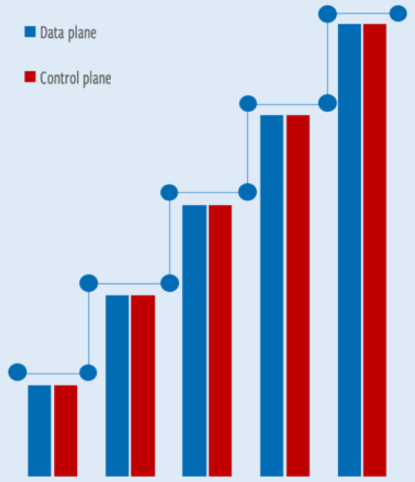

4. Hardware Lock-in

Legacy storage systems are pretty rigid in terms of the kinds of hardware that you're able to bring to bear. Typically, a vendor will require that you procure everything from them, from the shelves to the disks to the management software. And frequently, the storage vendors will charge a premium for any hardware they sell and you, as the customer, must bear those costs.

Even worse, you're stuck with whatever technology the vendor has decided to support. If your provider decides to not adopt a particular new technology—and you need that technology—you may be left with hardware that doesn't meet your needs, in which case your only option would be to rip and replace it with a solution that does meet your needs.

Essentially, with traditional storage you don't have the flexibility to tailor a solution to your company's needs, at least not in a holistic way. Further, your expanded storage must consist of what your vendor is willing to sell you and in increments that may not make sense, particularly when it comes to supporting unstructured data. As new technologies and vendors become available to the enterprise, IT teams should be able to incorporate them seamlessly.

5. Location Lock-in

Installing storage in a local data center is easy. Making it work absolutely seamlessly with your geographically distributed storage environment… not so much. Moreover, integrating with the cloud isn't always straightforward, even as the public cloud becomes a far more viable option. There are many companies that seek to adopt some kind of hybrid cloud storage solution, but their on-premises environment is holding them back in some way. Perhaps "cloud licensing" is too expensive or perhaps installing the components needed to integrate public cloud services is too complex. Regardless of the reason, these customers aren't able to leverage the benefits that the public cloud can offer.

6. Cost and Complexity

Here's a sobering fact: you probably pay for storage that you don't even use. Consider for a moment the lifecycle of your storage environment. In all likelihood, you have a 3- to 5-year replacement cycle for that environment. If you're like most organizations, that means that you buy today *all of the storage that you anticipate needing for the next five years*. In other words, let's assume that you need just 10 TB of storage today, but you anticipate needing 100 TB by the end of your replacement cycle. Many companies will buy that 100 TB today.

Frankly, in some ways, it almost makes sense for them to do so:

- **The best deals are had at inception**. Once you're locked into a platform, the storage vendor has no incentive to provide you with discounts as you need to add more capacity.
- **It can sometimes be difficult to scale traditional storage**. The process can require downtime and may require significant planning. All of this takes time and it can negatively impact the business. If you just buy it all up front, you don't need to worry about upgrading.

That said, it still hurts when you consider the fact that you're deploying storage for which you have no hope of realizing a return for possibly five more years. With this in mind, the pay-as-you-go consumption model of public cloud—a key driver of shadow IT—gives businesses a far more attractive cost model.

Reshaping the Storage Environment

It's time for storage admins and IT decision makers to begin rethinking storage. In recent years, the storage market has undergone tremendous transformation, with solutions that address some of the challenges in traditional storage. However, as we look into the future of storage, there are five key characteristics that will drive the future of storage systems that support unstructured data.

As we venture on our journey through these five characteristics, we'll briefly discuss how NooBaa's frictionless object storage environment addresses each one.

1. Object Storage

Block and file storage was and is really good in certain uses cases, such as server virtualization and employee file-shares. However, these kinds of storage devices begin to crumble under their own weight when used in large-scale unstructured-data applications such as archival and big data. As companies accumulate millions and billions of data elements, block and file storage systems demonstrate their unsuitability, particularly when those data elements are of the unstructured variety. Block and file systems don't lend themselves well to enabling organizations to understand the data that's being stored. While file names and the like provide basic descriptive capabilities, that designation paradigm does not scale well and does not enforce informational consistency.

Object storage has emerged as a way to solve the problems that become visible when companies seek to store so much data and attempt to make tha data meaningful. Data is imbued with meaning through the use of *metadata*. In short, metadata is information that describes data. If you think about a database, it's considered structured data because there is a definition of each data element. Metadata adds to unstructured data various fields of information that can be used to describe what is being stored. Each of these data elements is referred to as an *object*, hence the name *object storage*.

Whereas file information is abstracted from a file system in a file-based storage device, object storage metadata is stored alongside each object in the environment. Rather than being retrieved by a file name or a block number, items in an object store are retrieved through the use of an object ID (OID). Each unique OID is a consistent field, guartanteed to be unique, and is a part of the overall metadata record for an object. This high level of organization makes it simple to write queries that can retrieve objects based on contents in metadata fields.

An Introduction to NooBaa

NooBaa software creates a scalable storage solution in 15 minutes based on existing infrastructure, without requiring any new hardware. The agility from fast startup, elasticity, platform and network flexibility means lower IT risk for the Enterprise.

Organizations face an ever-growing demand for persisting unstructured data, but budgets are flat. Amazon S3 is agile and elastic, but imposes hidden costs and loss-of-control. Modern storage can scale, but existing solutions are way too complex, monolithic, capital-intensive, and can't offer the flexibility needed for file-storage in the cloud era.

Key Benefit of Object Storage

The core benefit of object is lower cost, easier scaling, self-healing from failure, WAN-friendly, high reliability, and, in the case of NooBaa, unprecedented agility.

2. Frictionless Deployment

Complexity is out! As companies seek to remake their storage environments, there will be less and less tolerance for solutions that are difficult to deploy or that require complex administration. Frictionless deployment of new storage solutions will become an absolute requirement.

Ease of Use

There was a time when IT could inspire awe among the user base by hiding in the back room, pushing buttons and pulling levers to make magical technology things happen. Hiring more and more IT people was seen as a simple cost of doing business. Those days are pretty much gone. Over time, the business will increasingly demand that IT people continue to morph from pure, deeply skilled, siloed technologists into hybrid technologists/business analysts.

As a result, any emerging storage environment must be eminently simple to use. Bear in mind that decision makers now compare everything IT does to what is possible with the public cloud. Executives have heard the calling of the public cloud and are enticed by its seeming simplicity and promise of unparalleled agility. Future storage for unstructured data will have to live up to the likes of Amazon Web Services when it comes to delivering a viable user experience.

Software Defined Storage Enabling Immediate Deployment

A lot of vendors used to build their own hardware, primarily due to the fact that they could create a single purpose hardware component that performed at top speed. Today, hardware has taken a back seat to software, thanks largely to the efforts of companies like Intel. With affordable CPUs on the market that contain a dozen or more processing cores, storage vendors no longer need to create their own hardware. Instead, they can dedicate a few cores of a commodity processor to storage management needs.

The ability to leverage commodity hardware has enabled software-centric storage vendors to enter the market. These *software-defined storage* (SDS) solutions make it possible for a customer to simply download a storage product and install it on hardware that they already own. This effectively unshackles these customers from storage vendors that sell overpriced hardware and enables the customer to very easily tailor a custom storage solution that meets the specific needs of their organization.

In an ideal world, software-based storage solutions should deliver agility and hardware neutrality that could deliver the holy grail of the 15-minute install. But in some cases customers have experienced the opposite. In some instances, particularly when considering open source object stores, install complexity has regressed. Further, there are bugs, instabilities, and incompatibilities to deal with before the system is operational, extending the time that it takes to become operational.

NooBaa's Frictionless Deployment features

How does NooBaa back up its claim to *frictionless deployment*? The four points below are just the beginning.

- **Immediate use**. 15 minutes from download to being operational
- **Platform agnostic**. Use any compute resource for storage
- **Network and cloud agnostic**. Cluster compute resources across any network
- **Fully automated**. Machine learning optimizes for resilience and performance with requiring constant administrator intervention

Is Your Software-Defined Storage Really Software-Defined?

Many software defined storage solutions are software-defined in name only. They carry with them narrow hardware compatibility lists (HCLs). The HCL defines exactly the hardware that the vendor is willing to support. So, while it might be an SDS solution, friction remains in the deployment process as you have to carefully match your hardware shopping list against the HCL.

Many SDS solutions have strict hardware requirements. Even some appliance vendors claim to be in the SDS business when, in reality, they are very tied to the hardware they bundle.

Fully Heterogeneous Deployment

The platform on which your storage operates really shouldn't matter. As storage systems make their way into the future and become even more disassociated from underlying hardware, it's become possible to use just about any platform as part of a storage solution. With a pure software defined storage solution, the underlying hardware should be whatever makes the most sense. In essence, you're able to eliminate friction from the platform selection process.

Physical

Although the entire world is virtual these days, physical servers still rule when it comes to raw power and expansion capabilities. In fact, physical servers are enjoying something of a renaissance thanks to the trend of adopting commodity hardware to serve critical infrastructure needs. The right software defined storage solution will allow you to leverage physical servers as part of your storage solution deployment. With these systems, you can choose your own storage types – spinning disk and/or flash – so that you can best meet the needs of your business.

Virtual

Most organizations are virtualized to a significant extent and want to leverage this environment for their storage-related servers as well. Virtual servers are simply software-based instantiations of physical servers, so deploying a storage node to a virtual server should be a fully supported operation. In fact, your storage environment should enable mixing and matching your physical servers and virtual ones.

Cloud

Companies are actively seeking ways in which they can sensibly leverage public cloud resources. The 100% OpEx-focused nature of the public cloud makes it a compelling target for certain activities. And, when you really think about it, an Amazon EC2 server is the same as an on-premises bare metal SuperMicro server, so why shouldn't it be eligible for inclusion in your fully distributed storage fabric? In fact, with companies actively looking for more flexibility in the platforms and locations on which they operate their workloads—including storage—the ability to seamlessly integrate cloud-based nodes into a storage system is a core characteristic of emerging storage solutions.

Through the combination of storage node deployments in the public cloud as well as on-premises, you're able to implement a storage environment that is hybrid in nature. This brings you the great economics of public cloud, which is perfect for some workloads, while also providing you with very fast local storage to support those workloads that demand low latency and fast throughput.

Network Flexibility

Systems need the ability to arbitrarily adapt to changing network topologies in an enterprise, and that means being able to seamlessly interoperate over the wide-area network, to and from public clouds, and the ability to shift capacity resources arbitrarily.

Complexity is Cost

Adding knobs, controls, and options just makes storage harder.

NooBaa delivers capabilities in terms of placement, tiering, and other performance optimizations, but these capabilities are automated, guided by metadata, and by better situational awareness of resources and application behavior.

NooBaa enables a fully frictionless deployment of its software defined storage solution by allowing you to deploy to a bare metal physical server, to a virtual machine, or to a cloud provider, such as Amazon Web Services. All of these target locations are able to participate alongside one another as nodes in the NooBaa architecture. Further, NooBaa hides the complexity that's inherent in linking together various nodes all over the map and provides automated tiering and disaster recovery services.

NooBaa is not constrained to single-subnet networking like traditional storage solutions. NooBaa may be distributed across subnets and can even be separated by a firewall. The distributed nature of NooBaa makes it an ideal target for ongoing storage of unstructured data.

3. Agility and Flexibility

The ability for a business to be agile and flexible is often a "lowest common denominator" function. The business can only move as fast as its slowest component. With the kinds of storage solutions that are now available on the market, there's no reason for storage to become a roadblock that makes the company less able to pivot when necessary.

Shared Hosts

Let's consider for a moment the technology phenomenon known as hyperconverged infrastructure. Hyperconverged infrastructure is a data center architectural option that effectively combines compute and storage resources into single appliances, thereby eliminating the traditional SAN from the equation. While hyperconverged infrastructure is a fantastic solution when it comes to supporting workloads such as virtualization, it's not particularly flexible in many cases and does not generally provide object storage for direct consumption by applications. Moreover, hyperconverged infrastructure vendors that sell their software bundled with appliances often have rigid hardware requirements defined with static partitioning of host resources for the very specific environments they support (such as the aforementioned VMware vSphere). So, while these environments truly aim to provide cloud services, they often only achieve a "cloud-like" environment.

Tomorrow's storage environments will represent what the hybrid cloud really should be about:

- Resources are scattered everywhere
- They exist on dedicated and shared resources
- A modern definition of "cloud" should be able to consume them to maximize utilization of IT assets and improve business agility

Dynamic Resource Environment

As you learned earlier, monolithic storage systems are a dying breed; they are being quickly replaced by nimbler solutions that can better meet the needs of the business while also significantly improving on storage economics. A software defined object-based storage system carries with it an architecture that allows IT to dial-in the perfectly balanced set of resources necessary for organizational workload support. This is achieved by decoupling the data plane and the control plane, which effectively eliminates the need to scale resources in lockstep with one another.

NooBaa Architectural Overview

Adding knobs, controls, and options just makes storage harder.

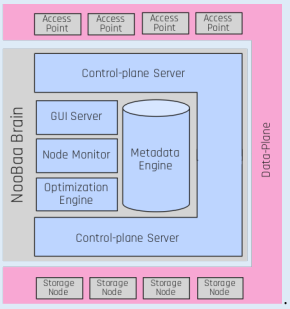

Core

NooBaa Core is the brain of the NooBaa architecture. It makes all of the decisions that are necessary to optimize the environment. NooBaa Core determines exactly where each chunked item will reside in the object store. These decisions are made based on a number of factors:

- **Resiliency**. A minimum of three copies of each object are stored in the NooBaa store. NooBaa Core makes storage decisions that ensure that each copy is isolated from the others so that the environment can withstand the loss of a disk or node.
- **Performance**. Placing a "hot" chunk—that is, one that is actively being used—on an Amazon S3 node wouldn't make much sense if there is a local node with flash storage sitting in the data center. NooBaa Core makes data placement decisions based on application performance requirements
- **Locality**. Very tightly coupled with performance is the idea of data locality. The closer that data exists to applications, the better the overall performance of the environment.

(continued on next page)

Granular Scalability

Architectures, such as the one provided by NooBaa, provide an opportunity for eminently granular scalability, a characteristic that is really important for storing unstructured data. Every workload is different. Every company is different. With an architecture that splits the control plane from the data plane, it becomes easier to grow where it makes sense. For example, if you start to push a lot of little data chunks into the system all at once, you can simply add more Access Nodes to the mix. If you're starting to push a lot of big data chunks into the environment, you can add more Storage Daemons so that those chunks have a place to live.

Intelligent Machine Learning

Software-based intelligence is all the rage these days. From IBM's Watson to Tesla's Autopilot feature, we're moving into a world in which humans will function alongside machines with virtual intelligence. Just like humans, these kinds of systems need to constantly learn and adapt their services based on what's happening *right now*. Loading a static and point-in-time set of parameters no longer cuts it. Consider Tesla's Autopilot. If you're driving a vehicle equipped with this feature, the chances are good that you'd much rather know that the feature is actively gathering information from what it's seeing as you commute to work rather than relying on information that was loaded on an SD card when the car was manufactured.

Storage-centric intelligent machine learning systems are the same. You care about what's happening in the *here* and *now* that could negatively impact the storage environment. Over time, machine learning models are used to identify trends in system health that may reflect network issues, impending failures, or general resource contention trends.

4. No Compromise Feature Set

Wanting low cost frictionless storage doesn't mean that you're suddenly deciding that enterprise-class features are unimportant. Rather, modern storage systems need to double down on high-end features in order to help organizations begin to make the case that moving to a new class of storage is both operationally and economically viable.

High Availability

Storage can't go down. Full stop. Sure, individual nodes should be able to crash or be taken out of service for maintenance, but these incidents should never result in an inability to access data.

(Continued from previous page)

Access Nodes

Access nodes are the window into the NooBaa solution. These are the nodes with which application servers will interact and they provide HTTP access to do so. As storage traffic comes into the access node:

- It is broken into chunks for storage on various nodes in the object store
- The deduplication engine fingerprints each chunk to make sure it doesn't already exist somewhere in the environment
- Each item that must be stored is compressed to reduce its storage utilization
- It's encrypted to maximize the security of your environment.

Storage Daemons

Once a chunk has been deduped and compressed by the access node and the location for placement of that chunk identified by the Core, it's sent off to a Storage Daemon node, which will be the home for that chunk. Storage Daemon nodes are the systems that contain the actual storage in the NooBaa architecture.

Data Protection/Replication (data copies/erasure coding, but won't be saying erasure coding directly since it's not supported yet)

If you lose a node or you lose a building, you shouldn't have to worry about whether or not your data is safe. Modern storage systems include comprehensive data protection capabilities and can replicate data to remote, geographically disparate data centers, enabling recovery from both local system failures as well as natural disasters.

For systems such as NooBaa, multiple copies of data are spread around different nodes – at least three – so that the storage architecture can withstand the loss of multiple nodes and remain operational.

Limitless Scalability

What if you could just keep adding new storage to your environment without having to worry about hitting a limit? This ability is a key need for any systems that are intended to support the glut of unstructured data coming into the enterprise. At a moment's notice, you should be able to deploy a brand new node that expands either the capacity of the storage environment or the brains.

Machine Learning and Object Storage

NooBaa uses machine learning to match hot data with high quality resources, to push cold data to low cost resources, and to optimize data geolocation for best application performance. In addition, the NooBaa Core actively manages resource security and trust, supports content specific optimizations, manages extended metadata, and provides a rich management interface.

Global Deduplication and Compression

You've already learned that storage capacity needs are growing at an astounding rate. Storage solutions that can scale to meet increasing capacity demands are critical. However, what if you could also store just a fraction of your data, without actually losing any data? And, what if you could squeeze the data that you do store into a smaller space?

Enter deduplication and compression. These two technologies work in concert with one another to actively reduce your capacity needs. Deduplication works by scanning each and every chunk of data that makes its way to a storage system. If that chunk of data already exists elsewhere, the incoming chunk is discarded and a small pointer is written in its place. Every environment has overlapping data of some kind. Deduplication has the capability to significantly reduce how much disk storage capacity you need.

Compression works against chunks that have already been deduplicated. Once the deduplication engine completes its work, the chunk is then squeezed down to a smaller size, kind of like when you ZIP a file.

These kinds of data services are essential to supporting unstructured data storage needs, both now and in the future. Deduplication and compression significantly reduce overall capacity needs, helping organizations save significant amounts of money on storage purchases. Further, since less hardware is required, companies are able to reduce the power and cooling costs associated with supporting storage in the data center.

Active Performance Management

We talked a bit about intelligent machine learning. With such capabilities, you shouldn't need an administrator constantly hovering over storage and turning knobs to eke out every last ounce of performance. Rather, the system should have the ability to learn what's normal and self-adjust as necessary.

5. Improved Business Outcomes

The entire goal of everything IT does is to improve overall business outcomes. This absolutely includes storage, although traditional storage has sometimes made this a challenge. As the world moves to new storage paradigms, there is an opportunity to align storage outcomes with the needs of the business in ways that have simply not been easily attainable before.

Reduced Total Cost of Ownership

When you're able to completely eliminate hardware dependencies both at the storage procurement phase and throughout the lifecycle of the solution, your ability to reduce your total cost of ownership (TCO) increases dramatically—as does your ability to leverage new technologies. Your newfound independence means that, as new technologies hit the market, you're able to leverage them without fear that they may not be supported by your storage vendor.

Flash storage is on the verge of knocking spinning disk off its price/capacity pedestal. New and innovative technologies, such as Intel's 3D XPoint, are poised to hit the market. With a storage solution that is fully decoupled from underlying hardware, you don't have to skip over technologies that may be transformational for your business just because your storage vendor won't support them.

This is a big deal and, at first, the TCO benefit might not be apparent. However, as you consider *opportunity cost* as part of your TCO story, things start to change.

Of course, *direct* costs, not opportunity costs, are generally considered staples of TCO calculations. With a solution such as NooBaa hardware and platform heterogeneity mean that you can deploy to the lowest cost hardware that is reasonable to support your use case. Your budgets are positively impacted in a few different ways:

- You're able to reduce your capital expenditures (CapEx) since you can deploy to low-cost hardware. Further, the architectural granularity of NooBaa means that you can specifically target what needs to be expanded and you're not having to add resources in lockstep. For example, if you don't need more storage capacity, but you do need more processing power, you can add just processing power. This reduces overall waste in the environment.
- Your operational expenditures (OpEx) are reduced thanks to the overall ease of use of the platform.

Translating New Architecture to Value and Reduced TCO

Over time, technology is supposed to become better, faster, and cheaper. It's also supposed to become easier to deploy, and easier to maintain. This paradigm has existed for decades and we see it in action every day. Whereas you used to wander about with a dozen devices—your compass, a camera, a cell phone, maybe a GPS unit, and a weather radio—when you took a hike, today, you just carry your smartphone, which has replaced all of those devices and is far easier to manage.

The same model is coming to the world of storage. Through the frictionless deployment of an easily-managed platform-agnostic software defined object storage solution with the ability to span clouds and data centers, you'll discover how a new storage architecture leads to the taming of unstructured data growth and brings to your business a reduced overall total cost of ownership.

To learn more about the future of unstructured data storage, visit www.noobaa.com.