The Software-Defined Data Center of Tomorrow Starts Today

Scott D. Lowe
Partner and Co-Founder
ActualTech Media

November 2015

Table of Contents

ActualTech Media

Introduction

As is the unfortunate case for many buzzword-based trends, "software-defined" data center resources are in the midst of going through their own hype-cycle – with the outcome, at least temporarily, of confusion and, until some semblance of reality kicks in. For many, "software-defined" is a meaningless term, hijacked by corporate marketing machines and twisted to meet the current product set for a specific vendor. For others, there is simply a lack of full understanding of what the phrase really means and, more importantly, what it will mean for the data center of tomorrow.

In this paper, you will learn about the software-defined data center – with a focus on storage. You will discover the key challenges that are giving rise to software-defined storage. Finally, you will also learn about how software-defined resources operate and why this seemingly overhyped technology is just now becoming a major force for data centers that are ready to meet the business needs of today and tomorrow.

First, There Was Hardware

Storage is hardware. You save your bits and bytes to some kind of storage device. When people refer to non-software-defined storage, they are generally referring to more traditional storage arrays. These storage arrays stand-alone and ship as a hardware chassis filled with some combination of spinning and solid state disks along with an operating system that exposes that storage.

This reveals a key fact that isn't always discussed: storage has always been software-defined. No matter what the underlying hardware looked like, there is always a software layer running atop that hardware that provides an interface to other systems and applications. One common misperception around software-defined storage is that the hardware no longer matters.

Nothing could be further from the truth.

Hardware remains critical even in a software-defined world. You still need to take care selecting hardware. Software alone won't make a spinning disk read data faster than a solid state disk. Software alone won't transform your five-year-old SAN into a state of the art all flash array.

However, software can imbue even somewhat older storage with capabilities not originally provided with that hardware. For example, software can give thin provisioning capabilities to storage systems that lack them. With software-defined storage, the underlying hardware can gain some attributes beyond those that shipped with it.

Software-Defined Storage Origins

In its earliest form, software-defined storage was devised to help customers focus less on underlying hardware and focus more on mixing and matching different kinds of storage. More specifically, it was a way for storage vendors to de-emphasize the need for specific supporting hardware. Software-defined storage enables customers to choose the appropriate combination of storage hardware that they need in order to support their workloads. Rather than a vendor putting together a hardware and software configuration for you, you get to put together your own hardware stack and apply the software-defined storage solution to manage it.

In a software-defined storage system, unlike storage devices are aggregated together and pooled under a common management paradigm. Then, various data services – such as provisioning, data protection, and tiering – are applied to that storage pool.

Paving a Road to the Future

Storage is undergoing a major transformation. The storage architecture of yesterday is no longer appropriate as companies seek more agility and flexibility in their data center operations. Perhaps more importantly, companies are seeking ways to simplify data center operations and reduce the overall cost of procuring and managing the data center environment.

Watching what's happening in the broader storage environment, it's readily apparent that an era of simplicity is upon us. Just look at the move from direct attached storage to SANs. Even back then, it was about ease of use. Rather than managing dozens of separate storage silos on individual servers, SANs enabled administration from a central location. Over time, however, SANs have become among the most expensive and complex data center elements, so the market is reacting through the introduction of new solutions, such as hyperconverged infrastructure – an architecture that is getting a lot of attention.

At the core of hyperconvergence, and available as a data center architecture in its own right is a whole new generation of software defined storage. Many traditional storage approaches – even some original software defined storage systems – hit limits on what could be accomplished. Modern software defined storage systems seek to move past these limits in order to embrace new potential.

Consider this: corporate data growth continues to increase at levels not matched by IT budgets, which may actually be shrinking. In fact, data growth outstrips budget growth by up to a factor of 10. The ability to leverage commodity storage devices is a hallmark for modern storage architectures. The word commodity in this context shouldn't be construed as indicating that hardware is an afterthought or in some way unimportant. Nothing could be further from the truth. Rather, with modern software defined storage systems, there is an ability to leverage whatever kinds of hardware make sense for the organization's needs. The solution is led through addressing outcomes, not by attempting to cobble together predefined hardware configurations into something semi-workable. Without the need to buy custom-developed hardware – by leveraging commodity hardware – costs can be significantly decreased, often by as much as 60%.

ioFABRIC Vicinity

ioFABRIC Vicinity is a next-generation software defined storage solution intended to meet even the most demanding organizational data storage needs.

Vicinity Deployment Model

ioFABRIC can be described as a "Distributed Server SAN," using modern parlance. By running an ioFABRIC agent on a Linux server with direct-attached storage, you effectively turn that Linux server into a storage node that can participate in a larger cluster of storage resources. Additionally, you can choose to run in what is called "Converged Mode." A converged mode cluster node runs both the ioFABRIC agent as well as your other workloads, which makes the solution operate much like a hyperconverged infrastructure system. In converged mode, ioFABRIC presents storage as a virtual disk or SCSI device.

The agents installed on each node are responsible for managing all of the local storage on that node, which includes a *profiling* process (which is described later) to look for all visible storage, which can include spinning disks, flash media, and even DRAM.

At present, ioFABRIC exposes managed storage via iSCSI. As such, any system capable of connecting to a storage system via iSCSI can consume storage from an ioFABRIC-managed platform.

Comprehensive Capacity Utilization and Seamless Expansion

First, Vicinity makes storage capacity provisioning seamless if there is capacity anywhere in the organization that can be leveraged. If there is capacity available, it should be used. If capacity is not available, a storage administrator should be made aware. If you need more capacity, you simply add it and Vicinity uses it. In this sense, Vicinity is more like delivering storage as a service to the company.

One of the key capabilities in Vicinity is its ability to support both brownfield and greenfield deployments. You don't need to have brand new equipment in order to deploy Vicinity. You can incorporate your existing storage environment in an ioFABRIC-driven solution. For example, you may like to deploy a software-led storage fabric that incorporates newer architectures like hyperconverged infrastructure and flash media, but also legacy devices. ioFABRIC can non-disruptively embrace and extend your current investment. It discovers the storage in the infrastructure, profiles it – gains insight into its performance, capacity, and data protection characteristics – and then uses it all to deliver the quality of service needs to applications.

ioFABRIC allows you to grow storage in multiple ways. You can scale *up* by adding more storage to an existing ioFABRIC node. Or, you can add complete nodes with new storage – this is the scale *out* method. Each method has its own benefits. By scaling up, you don't need to worry about adding any hardware except storage; you basically just add more disks to an existing server. However, you do place a bit of additional stress on the shared components in the server. When you scale out, you are adding all of the underlying resources that storage needs in order to operate, but the expansion can be a bit more expensive.

A Simplified Storage Experience

Simplification is becoming a major trend in all things IT and nowhere is this more evident than in the data center. Administrators have made a valiant effort to cope with the increasing complexity and the need to continually do more with less, but the battle is being lost. Costs keep rising; complexity has continued to increase – which also increases costs. Worse, with traditional systems, provisioning times continue to increase, as does downtime. None of these trends are palatable and none make for a sustainable environment.

The industry needs to make things better in a way that's easy and non-disruptive. The storage should just be there and should just work. This is exactly the outcome promised by ioFABRIC Vicinity. Figure 1 provides you with a look at Vicinity's very clean and accessible administrative experience.

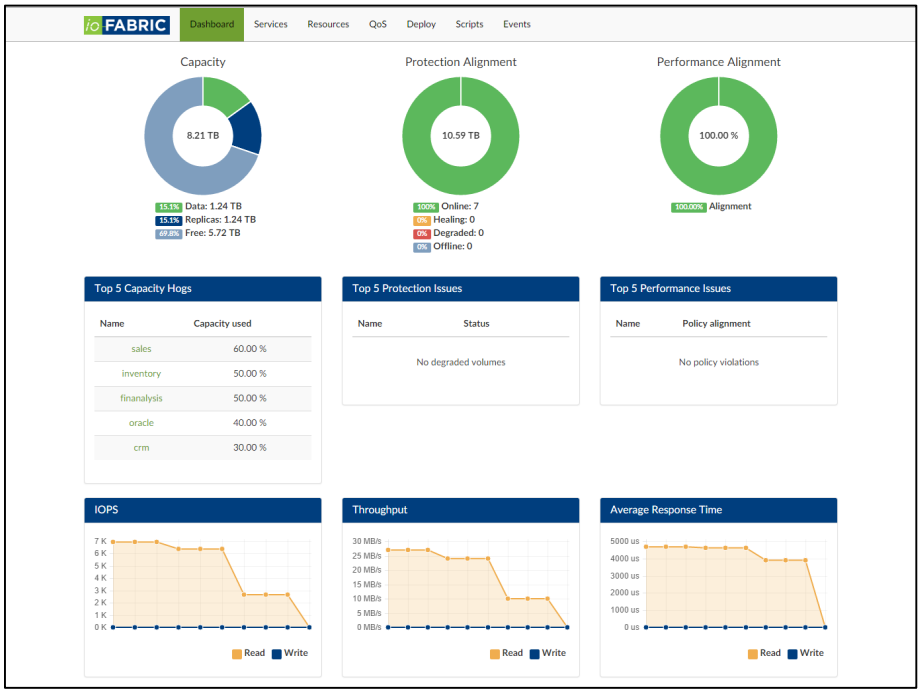

Storage Efficiency Equates to Cost Savings

By improving capacity utilization and reducing complexity – and, thus, administrative burden – Vicinity has the potential to help organizations reduce their storage spend, particularly when viewing the total cost of ownership for the storage service. However, there are some additional features that Vicinity brings to bear that can further reduce the overall cost of storage.

Human nature seems to force us to always demand more "room to grow." It starts with buying kids' clothes that are a size too big, but it also transfers to data storage, where storage administrators seek to ensure continual sufficient capacity by overprovisioning storage. In fact, in many places, storage maxes out at 60% utilization due to overprovisioning. Vicinity brings global thin provisioning across all storage assets to bear to help administrators make more efficient use of capacity. With thin provisioning, administrators can enjoy 90% and higher utilization, which can mean the difference between having to spend more money now to add capacity or being able to defer that spend for up to two years.

Donald Lopez
IT Manager
InReach

"ioFABRIC Vicinity allowed me to simplify an overcomplicated NAS storage environment and utilize the available storage capacity with far greater efficiency than I could ever hope for from our appliances.

With Vicinity automating management of our infrastructure at the service level, we were able to optimize performance, reduce our storage costs and extend the life of our NAS units, undoubtedly for several more years."

Proactive Performance

In recent years, storage performance has emerged as one of the primary challenges faced in the data center. ioFABRIC Vicinity boasts a number of capabilities intended to help organizations handle storage performance challenges without having to hire expensive PhD-level storage engineers.

Profiling

First is the previously mentioned storage profiling process that Vicinity uses to determine overall storage capabilities. This profiling process is not a one-time operation. In fact, it's an ongoing process that takes into consideration real-time events that may be impacting performance and availability. For example, suppose you have three network interfaces on a storage array. Vicinity will actively manage all three paths in order to ensure that all are being used to maximum effectiveness and will work around problems – e.g. one of the interfaces fails – in order to maintain optimal performance and availability. Not that long ago, such proactive measures would have required the constant attention of a storage engineer.

The profiling process happens in real-time and is not simply an algorithm that looks at yesterday's statistics in an attempt to determine what might happen today. Across all managed storage resources – both local and cloud-based – Vicinity keeps a watchful eye on things in order to determine what is happening at the present time.

Action

But performance goes beyond just profiling storage capabilities. It also requires action based on the results of the profiling process. This is where Vicinity's Intelligent Data Placement – sometimes referred to as micro-tiering – comes into play. Vicinity's Intelligent Data Placement service leverages the results of the profiling process by automatically moving active data to appropriate media types based on performance requirements. Operating using a small block size of 256K, this service will actively move in-use data to higher-performing media while inactive data is moved to lower performance media.

Flash and DRAM Usage

By leveraging different media types in this proactive way, Vicinity helps you to uncover the real benefits of flash and remove barriers to entry. If you're running an all spinning disk storage environment and decide to add some flash, Vicinity will find it and start to move high-use data to that fast storage, enabling immediate and automatic value for the flash media.

Maybe you don't have any flash disks in your storage environment. In order to maximize performance, Vicinity also pools together DRAM resources across all storage nodes and treats that pool as a storage location. By granting the ability to use ultrafast DRAM as a part of your storage environment, you can gain immediate performance benefits.

Application-Centric Quality of Service

You buy storage to meet the needs of your business applications, which might include file sharing, email and collaboration, databases, ERP, and much more. Each of these applications has unique needs when it comes to storage and these needs are a combination of capacity, performance, and data protection.

For example, for a database application, you may need to have very fast storage with high levels of data protection, but relatively small space needs. Alternatively, for an archive, you may need a lot of capacity, but performance and data protection may not be as critical. Vicinity enables administrators to define what level of quality of service is needed on a per-application basis. Vicinity dynamically maintains service levels as the infrastructure changes and as the company's performance and protection needs change. By operating in this way, Vicinity effectively transitions policy from the storage layer down to the data level, operating more granularly than many other available tools. As shown in Figure 2, administrators can easily add new storage policies to the platform.

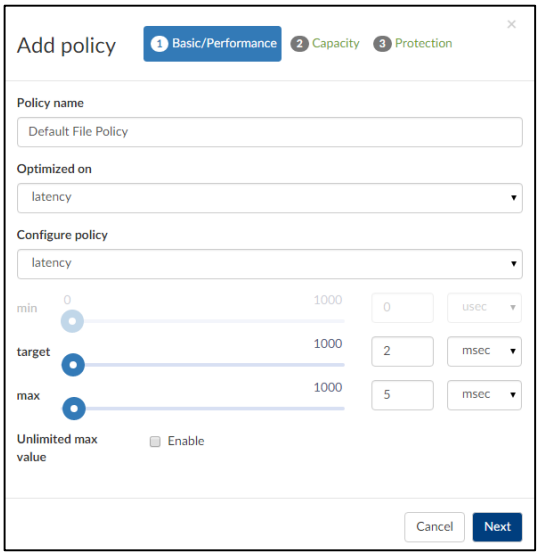

Vicinity's quality of service capabilities are deeply tied to the profiling service. As environment conditions change and as administrators make updates to QoS settings, Vicinity moves data to appropriate locations to maintain current QoS.

Bruno Barros
IT Administrator
Oforma

Vicinity is a great tool for anyone looking to maximize their storage and make it more efficient. The ability to bring together practically any type of storage into a pool is good already, but the fact that it then optimizes your bandwidth based on rules you can customize makes it that much better. The interface is extremely easy to use and understand and I love the overview you get on the dashboard. A must for any IT with limited funds.

Data Protection and Availability

No modern storage system is complete without data protection capabilities and Vicinity is no exception. To address critical data protection needs, Vicinity includes load balancing, replication and fault tolerance capabilities as well as the ability to snapshot and clone volumes. Data protection is ingrained into the quality of service features, which can be set per-application. Vicinity will ensure that data protection requirements are met by creating dynamic replicas in appropriate locations and will alert administrators if something goes awry.

Summary

Any storage. Easy to manage. That's the hallmark of the data center of tomorrow. With solutions that allow you to focus more on storage outcomes rather than inputs, ioFABRIC Vicinity enables companies to leap ahead in their journey to the software defined data center.