

The three must-haves for modern scale-out

Contents

Introduction—why scale? 1

Scale-up vs. scale-out 1

The three persistent pains of scale 2

The three must-haves for modern scale-out 3

1. Designed for scale-out 3

2. Built for the applications you need to scale 4

3. Intelligent and analytics-driven 4

Checklist 7

Summary 7

Introduction – why scale?

If you work in the data center, you've got a full plate. Your leadership pushes for a hybrid-cloud strategy. Your colleagues call for new services and projects. And you have to keep your existing (growing) virtual footprint up and running. Besides pulling you in multiple directions, the one thing all these items have in common is the requirement to scale.

And so over the next few pages we'll explore the imperative to scale. We'll identify the pain points associated with conventional scaling methods and define the three must-haves for modern, simple storage scale-out.

Scale-up vs. scale-out

Conventional scale typically starts with an enterprise-grade storage array containing two controllers and a lot of disks. To increase the capacity of the array, more shelves full of disks are added to the same original storage controllers. This is called storage scale-up.

This means of reaching scale has a major constraint: flash disks now have so much performance by themselves that the controller determines array performance and oftentimes is the performance bottleneck. Adding additional shelves only adds capacity and does nothing to alleviate the controller bottleneck—unlike the days of spinning disks, when adding more disks via shelves also added performance. And by piling up shelves of capacity you're creating a large failure domain—if you lose the controllers you lose the data on ALL the shelves attached to it.

The alternative approach is storage scale-out, which we'll focus on for the remainder of this paper. Instead of adding shelves of disks, scale-out involves interconnecting arrays to join their controllers together—adding both capacity and performance. This adds redundancy (of controllers) and increases your ability to put all that performance to effective use.

But, this also results in a far more rigid and complex arrangement. Often, you have to scale-out with costly proprietary hardware interconnects—you need a high-speed network for your multiple nodes to communicate. Furthermore, since this type of conventional scale-out is not designed for virtualized environments and cloud, you have to build out a mass of LUNs and volumes to manage your footprint, and a team of storage experts to manage it all.

Consider how different it is to scale compute vs. storage. If you need more compute resources for your virtualized workloads, you simply install a new virtualized server, add it to the resource pool and automatic live migration optimizes the VMs across the pool. What if you could scale storage as easily as you scale compute?

The three persistent pains of scale

Conventional storage scale-out is associated with a number of pains:

1. Inflexibility

At the outset of your scale-out efforts you generally have to purchase expensive storage hardware, proprietary hardware interconnects and even additional software or a different operating system than the non-scale-out systems from the same vendor. Plus, you need to staff one or more storage PhDs to make it all work together. And so, it's difficult to start small with scale-out.

Even more frustrating, if you bite the bullet and invest big in infrastructure and talent, there's no guarantee that you can scale much bigger. Some platforms that are billed as scale-out can only grow to 8 or 12 nodes, often in pairs or other lockstep increments, which is both a low ceiling and an inflexible growth path for a fast growing virtual footprint.

2. Complexity

Conventional, general-purpose storage has to account for legacy, physical workloads. Its foundation of LUNs and volumes was not designed with virtualized workloads in mind. Scale-out from hundreds to thousands of virtual machines is completely different from scale-out of one enormous database or physical workload. And so the inevitable happens—a proliferation of LUNs and volumes. Now, when virtual machines start to struggle with performance, you've got to intervene and shuffle them between LUNs and around your scaled footprint—but where to place them? What and how much to add to improve performance? How soon will you run out of capacity? LUNs obscure information about their behaviors and needs, so you're making your best guess and trying to track it all in spreadsheets. Who's spending their time keeping tabs on this very tactical mess? Your most skilled and expensive storage specialists who would much rather work on more strategic projects.

3. Cost

What's the result of inflexible and complex storage? It's high cost! It's easy for your investment in storage scale-out to spiral out of control.

In terms of capital expense, you have huge outlays for hardware, software and networking. In order to buffer the performance of your growing footprint, there's a good chance you'll be over-buying all-flash or over-provisioning storage in general. According to Gartner, the average organization only uses 58% of their storage footprint.¹

And operating expense has been growing even faster. Over the past decade, organizations have grown their operating expenses 4x faster than capital expenses.² Why? The average storage expert can manage only 344 terabytes of storage.¹ As you look ahead to your storage future, and potentially petabytes of storage, how much expertise can you find, afford and retain?

What are organizations doing to address all this storage pain? Solve for the root cause. Here's what they (and you) need to consider to source storage that enables a modern, easy-to-manage approach to scale-out.

¹ Gartner. IT Key Metrics Data 2016: Key Infrastructure Measures: Storage Analysis: Current Year. Linda Hall, Eric Stegman, Shreya Futela, Disha Gupta. 14 December 2015. G00291389

² IDC. Trends and Technologies Impacting Today's Storage Infrastructure Market. Eric Sheppard. Presentation dated August 30, 2015.

The three must-haves for modern scale-out

Storage scale-out can be as simple as scaling compute. You can scale-out and manage tens of petabytes and hundreds of thousands of virtual machines with a single employee—who does not need to possess any storage expertise. A number of organizations have proven that a modern scale-out strategy can be a source of immense time and money savings, and also a clear competitive advantage. But to keep scale-out simple and join their ranks, your scale-out platform must have the following:

1. Designed for scale-out

While plenty of conventional storage platforms claim to support scale-out, they lack the hallmarks of a platform purpose-built for scale-out:

- **Easy to start.** Rather than require a huge upfront investment, true scale-out storage can start with just two systems. There's no expensive networking needed, and anyone can manage it—because when you're starting out it's rare to have deep storage expertise on staff.
- **Easy to expand.** Whether you start with two devices or many more, it should be easy to grow that footprint to meet your needs. That means your scale-out platform needs to be:
 - **Federated.** You must be able to manage all your devices as one pool. That means that all-flash and hybrid, existing and future systems need to share a common architecture and operating system.
 - **Loosely coupled.** The individual nodes should be loosely coupled to separate control flow from data flow, ensuring low latency and scalability to a very large number of storage nodes. This approach allows you to expand well beyond the conventional limitations of 8 or 12 nodes to accommodate thousands of virtual machines.
- **Easy to manage.** Any scale-out solution that relies heavily on storage experts to manually plan, implement, and maintain the storage solution is inherently not scalable. Organizations increasingly turn to IT generalists aligned to needs of the business and its applications. You need to be able to do this too even if your storage needs to scale-out to meet those needs.

And since not every scale-out effort starts with massive investment or intention, you've got to be able to start small before you grow big. Most scale-out solutions require you to start big in terms of capacity, performance, and investment, and to fully commit to a scale-out platform into the future. Your scale-out infrastructure needs to be flexible to grow with your needs:

- **Grow from TB to PB.** Start with a single array, and add nodes as needed, with the right combination of capacity and performance for your needs.
- **Investment protection and future-proof.** If you're careful to select a storage platform that is both intelligent and designed for scale-out, then you've got investment protection. That's because you can grow a federated pool of storage that scales flexibly with both all-flash and hybrid, works with existing and future systems, across multiple hypervisors, and can be easily managed by non-storage specialists.
- **Scale with less.** Modern scale-out allows your storage footprint to grow according to your needs, meaning you can get higher utilization of each storage node because you can optimize for both capacity and performance independently. And you can manage it all with a small team of non-storage specialists.

2. Built for the applications you need to scale

Conventional storage arrays are “general purpose.” Which means that they were not built specifically for virtualized and cloud environments. And just because a vendor is all-flash does not mean that they are not conventional. Most of the storage vendors on the market today claim to do everything, which usually means that they do not do well with virtualized and cloud applications.

An important question to ask: what are the applications and services that drive your organization? Most organizations today would answer that their virtualized and cloud applications grow their business the most. They would also say that it is these very applications that are the hardest to design scale-out storage environments for. Every VM has its own SLA and performance requirements, but it is difficult to maintain those requirements even on a single conventional storage array using LUNs and volumes, never mind a complex scale-out solution with multiple controllers and LUNs and volumes spanning those controllers.

It is important to have VM-aware storage arrays to provide the performance, analytics and quality of service that each VM needs on an individual array. And it is just as important that the scale-out solution you choose fulfills the needs of the virtualized and cloud applications running on your infrastructure, not on LUNs or volumes.

3. Intelligent and analytics-driven

Today you might manage hundreds or thousands of virtual machines. As you scale-out, that management responsibility could expand to tens of thousands of virtual machines or even more. For that reason, your scale-out platform must have the intelligence to automate previously manual tasks and the analytics to inform your most important decisions.

- **Intelligence.** Your storage scale-out platform should work for you, not the other way around. Here are a couple of examples of how scale-out intelligence massively reduces the complexity of managing your storage footprint:
 - **Distribution.** An intelligent platform optimizes virtual machine distribution across the pool constantly gathering multi-dimensional data about the behavior of each virtual machine, taking into account space savings, resources required by each VM, and the cost in time and data to move VMs. Algorithms use this data to determine the optimum distribution of virtual machines—factoring in both capacity and performance, and making “least-cost” decisions based on a complete understanding of each VM’s needs and each array’s capabilities. This is a virtuous cycle that greatly simplifies management, and enables a single non-storage specialist to handle a storage footprint in excess of ten petabytes.
 - **Policies.** Applying policies to a growing footprint is not a trivial task. With conventional storage it’s highly manual, and as you move virtual machines between LUNs or arrays you typically have to manually re-apply policies. Intelligent scale-out helps you group your virtual machines into like groups, and apply quality of service and/or data protection policies to the entire group. As individual virtual machines are moved within the scale-out environment, they retain their policies automatically. The policies you set when you have 500 virtual machines stay in place as you grow to 5,000.

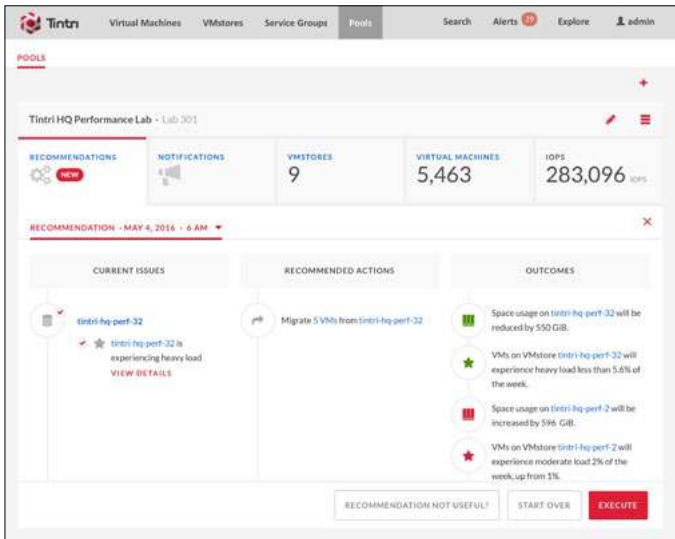
Green Cloud Technologies knows scale-out

Green Cloud Technologies was just named the fastest growing business in the state of South Carolina (2015). They are a cloud service provider, offering a set of services to hundreds of resellers and thousands of end customers. Over the past three years they have grown their business 4000% and their storage footprint 400%. Despite this rapid growth, they have no storage experts on staff, and spend less time managing storage now than they did three years ago. How? They’ve followed the three A’s of scale-out.

“Storage isn’t the focus of our team meetings any longer. Instead we spend our time on strategic projects, and finding ways to better serve our customers.”

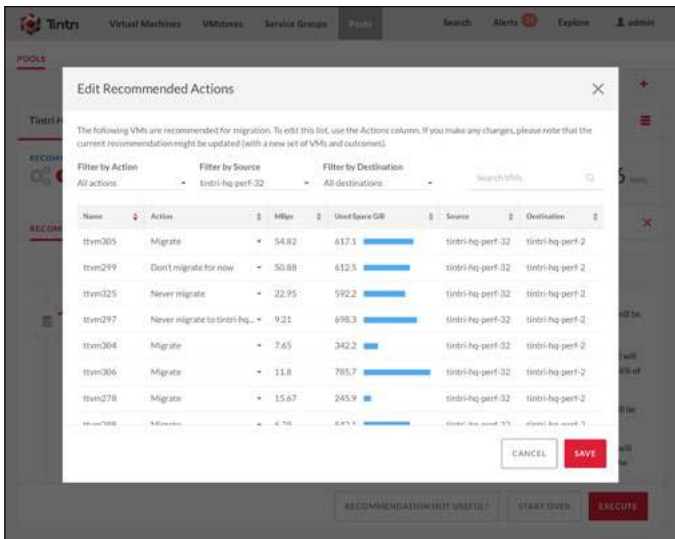
Eric Hester,
VP Engineering and Operations, Green Cloud

How automated placement recommendations work:



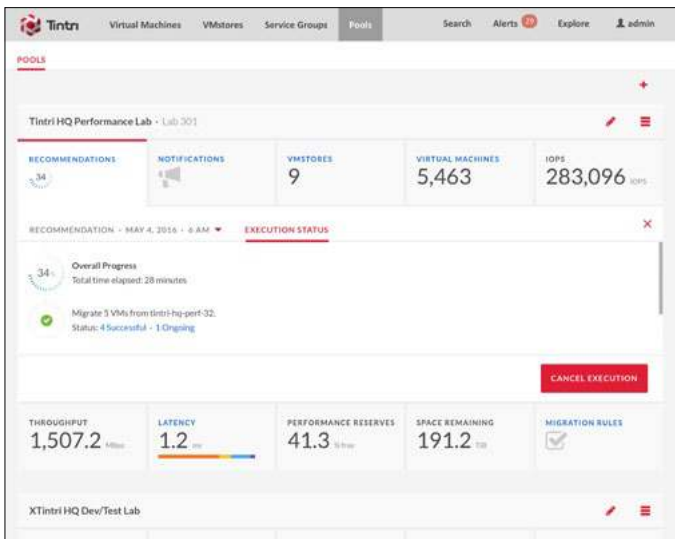
Step 1

The recommendation engine identifies an array that is in danger of running out of space in a week based on current growth patterns.



Step 2

Algorithms highlight 20 VMs that have been selected for migration to another system based on—rate of capacity growth, I/O load, flash working set required, presence of native snapshots and the impact on the overall health of the pool. The software calculates the outcome of this change.



Step 3

The admin executes the recommendation and the scale-out software completes the migration of these 31 VMs with no further manual intervention—moving all snapshots and policies.

- **Analytics.** All storage platforms generate analytics of one form or another. The difference is that some of these analytics are merely informative, while others drive decisions. The latter is a scale-out must-have, and so you need to look for:
 - **VM-level.** Most conventional storage relies on LUNs and volumes as the unit of management. Each LUN or volume may contain tens or hundreds of individual virtual machines. And so analytics presented at the LUN or volume level are not particularly useful or actionable. When virtual machines are the unit of management, you can take action at that level of granularity.
 - **Real-time.** The definition of real-time is a gray area. Some storage analytics that claim real-time delivery are actually delayed 1-24 hours. If you're heading outdoors now, it doesn't help much to know what the weather was like hours ago. As you review the analytics capabilities of storage scale-out platforms, be sure you understand exactly how the provider defines real-time. When you can see the root cause of latency, I/O consumption and other data about a single virtual machine right now, you can make better decisions.
 - **Application profiling.** Modern scale-out also includes the ability to profile applications—an input for analytics modeling. This involves grouping together types of applications (for example, SQL servers), so you can look at the behaviors of this group. You can see how many resources the application type is consuming, and drill into the member virtual machines to see which outliers need attention. If you're asked to add another 10 SQL Servers, quick what-if analysis of the application profile produces an immediate answer. (SEE FIGURE 1)
 - **Predictive planning.** Conventional, LUN-based storage can only estimate when you will run out of storage capacity. But when you know the precise behavior of all your virtual machines you can make very accurate predictions about your future need for capacity AND performance. Modern scale-out analytics allow you to trend consumption of capacity, I/O and working set, so you can see which resource is on the critical path, and project exactly when you will need to acquire additional resources, and what the balance of capacity and performance should be on those new arrays. (SEE FIGURE 2)

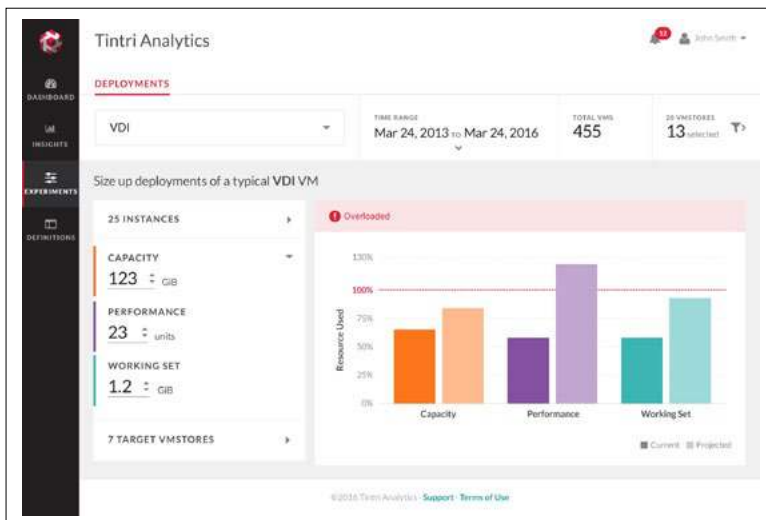


Figure 1: What-if analytics. Model the impact of application changes (number of instances, capacity, performance, working set) on your environment

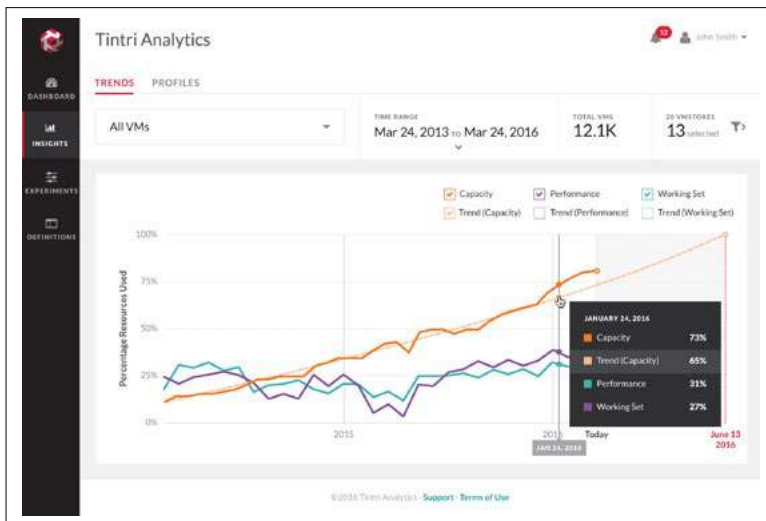


Figure 2: Predictive analytics. Forecast exactly when one device or a pool of devices will run out of capacity, performance or working set.

To summarize the key points above and help you find a platform that's a protected investment we've prepared the following checklist. Every "NO" answer is a strike against the platform.

Checklist:

The following checklist is intended to help you find a storage platform that can scale-out effectively. As you apply it, every “No” response is a strike against the platform. Three strikes and that platform is ill-suited for scale-out.

	Statement	YES	NO
ARCHITECTURE	The scale-out platform uses virtual machines as the unit of management instead of LUNs or volumes		
	The scale-out platform uses NO custom hardware interconnects between storage systems		
	The scale-out platform is backward compatible with the vendor’s existing storage systems		
	The scale-out platform is forward compatible with the vendor’s future storage systems		
	The scale-out architecture supports both all-flash and hybrid systems in the same pool		
	The architecture can start with a few terabytes and a couple hundred virtual machines		
	The architecture supports more than 1 petabyte of virtual machine data		
	The architecture supports more than 100,000 virtual machines		
	Storage and server can be scaled independently		
	The architecture requires no additional copies of data on different arrays as a redundancy measure		
	The scale-out pool can contain VMs from multiple hypervisors concurrently, without partitioning on the basis of hypervisor format		
	The scale-out pool can be partitioned into smaller pools that are managed independently		
AUTOMATION	Storage policies can be applied to groups of virtual machines		
	Storage policies travel with virtual machines as they move between arrays		
	The scale-out platform makes recommendations that optimize the distribution of individual virtual machines		
	The scale-out platform executes recommendations with no further manual intervention		
	The scale-out platform includes APIs and other tools that allow for custom scripting and automations		
ANALYTICS	The scale-out platform provides analytics at the virtual machine level		
	The scale-out platform offers real-time VM-level analytics		
	The scale-out platform allows groups of virtual machines to be analyzed in aggregate as an application profile		
	Analytics capabilities include organic growth modeling and what-if analysis		
	Analytics include the ability to forecast storage capacity requirements		
	Analytics include the ability to forecast storage performance requirements		

Summary

Throughout this document we have explored the elements of modern scale-out—the platform design and intelligence required to scale with simplicity. And there’s an undercurrent to this entire storyline: modern scale-out uses none of the LUNs and volumes that weigh down conventional storage. A LUN-based architecture was relevant when workloads were physical, but has become an irrelevant unit of management now when more than 75% of workloads are virtualized. LUNs and volumes severely compromise scale-out strategy—they obscure access and visibility into your virtual machines.

And so a prerequisite for modern scale-out is VM-aware storage—an operating system and software that operates entirely at the virtual machine level. If scale-out is (or will be) an imperative for your organization, it’s imperative that you build on a VM-aware storage foundation.